

From Molecules to the Masses

Visual Exploration, Analysis, and Communication of Human Physiology



Laura Ann Garrison

Thesis for the degree of Philosophiae Doctor (PhD)
University of Bergen, Norway
2022

UNIVERSITY OF BERGEN



From Molecules to the Masses

Visual Exploration, Analysis, and Communication of
Human Physiology

Laura Ann Garrison



Thesis for the degree of Philosophiae Doctor (PhD)
at the University of Bergen

Date of defense: 16.09.2022

© Copyright Laura Ann Garrison

The material in this publication is covered by the provisions of the Copyright Act.

Year: 2022

Title: From Molecules to the Masses

Name: Laura Ann Garrison

Print: Skipnes Kommunikasjon / University of Bergen

Scientific Environment

The work presented in this thesis was conducted as part of my PhD research in the Department of Informatics at the University of Bergen, where I was also enrolled in the Research School in Information and Communication Technology (ICT). In addition, parts of this work have been carried out in affiliation with the Mohn Medical Imaging and Visualization Centre (MMIV) at Haukeland University Hospital in Bergen, and the Center for Data Science (CEDAS) at the University of Bergen. Lastly, I have had the great fortune to collaborate with a number of external institutions during my PhD. These include Otto-von-Guericke University in Magdeburg, Germany, where I completed two research stays, one physical with the Medicine and Digitization Research Group in the Department of Neurology, and one virtual with the Simulation and Graphics Group in the Department of Computer Science. Additional collaborative institutions for research conducted in the course of my PhD include Johns Hopkins University in Baltimore, USA; TU Wien in Vienna, Austria; King Abdullah University of Science and Technology (KAUST) in Thuwal, Saudi Arabia; and Rainfall AS in Bergen, Norway.

Acknowledgements

This thesis is the result of a rather spontaneous leap into a different career path and future than I had imagined for myself. It has introduced me to several wonderful and inspiring people, many of whom I have had the fortune to befriend and collaborate with over these years.

My interest in combining visualization with the biological sciences grew from conversations with William Hamilton when I was halfway through my bachelor's degree and realizing perhaps I didn't like plants enough to make their ecology into a profession. His encouragement led me to pursue a career in medical illustration, and his support and friendship since have been invaluable. I'm indebted to Lindsay McHutchion, Evan Molinelli, Tarek Sherif, and Krystle Quinn, colleagues and friends from my time as a professional medical illustrator who encouraged me when I first considered pursuing a PhD and who have supported me throughout this journey.

For coming to Norway, I am indebted first and foremost to my supervisor, Stefan Bruckner, for taking a chance on my non-traditional background for a computer science PhD. I cannot understate the value of his guidance and support, his weird sense of humor, and for allowing me the freedom to develop and pursue my own scientific questions. Thanks also to my co-supervisor, Renate Grüner, for her advice and encouragement, particularly on all things MRI-related. I am grateful to Helwig Hauser for his sage advice and insightful discussions, not only on the scientific questions I'd pepper him with but also for meta-questions about academia. I'm indebted to Noeska Smit, who always made time to brainstorm, discuss questions in medical visualization and inspired me to rethink my deep dislike of knitting.

Proving that academic parents can convince their kids to be friends, I'm especially thankful to Steffen Oeltze-Jafre, Stefan Bruckner, and Helwig Hauser for connecting me to Juliane Müller-Sielaff in Magdeburg. I had the immense fortune to collaborate with Juliane on an incredibly fun and fruitful year-long research project, and she introduced me to the joys of D3 and 80's-inspired electronic music. Also in Magdeburg, I'm grateful to have connected with Bernhard Preim and Monique Meuschke. I have truly appreciated their mentorship and enthusiastic welcome into new research directions merging scientific storytelling into medical visualization.

I am indebted to past and current colleagues in the VisGroup: Sergej Stoppel, Fabian Bolte, Juraj Pálenik, Sherin Shugathan, Eric Mörth, Thomas Trautner, Fourough Gharbalchi, and Yngve Kristiansen, for our scientific discussions, knowledge sharing, and social gatherings that helped to make these years truly enjoyable. I am grateful as well to colleagues in the Mohn Medical Imaging and Visualization Centre, including Frank Riemer, Hauke Bartsch, Karsten Specht, Sathiesh Kaliyugarasan, Heidi Espedal, and many others for their various insights on medicine and MRI protocols, scripting tricks, and always being up to talk bikes, hikes, and the occasional burger and beer. I par-

ticularly am grateful to Alex Craven for his patience and help with MR spectroscopy data, his mentorship as I learned programming, and manuscript-proofreading over an embarrassing volume of coffee.

I am grateful to all of those who kindly participated in the surveys and focus groups in my research—in particular, Jeni Fairman, David Ehlert, Veronica Falconieri Hays, Jodie Jenkinson, Nick Woolridge, Jason Sharpe, Kevin Millar, Alison Burke, Cassio Lymn, Leah Leibowicz, and Bara Kozlíková. Without their time and willingness, much of this work never would have made it off the ground.

Finally, I can't thank enough my friends and family scattered across North America and Europe who I haven't already mentioned, in particular Anna Eilertsen for sharing her love of climbing, tea, and deep conversation; Kenneth Fossen for introducing me to the concept of Norwegian *friluftsliv*; Anita Nordal for wine and Dirty Vegan evenings; Neil Campbell for being a well of support and wit. To my family: Deanna, Kim, and Rod Garrison. Without their support, particularly in the depths of the pandemic, I would not have made it to this point. And, last but not least, to Ollie, my four-legged partner-in-crime and honorary VisGroup mascot, who is just as ready for monster hikes as he is for lazy sofa days.

Abstract

The overarching theme of this thesis is the cross-disciplinary application of medical illustration and visualization techniques to address challenges in exploring, analyzing, and communicating aspects of physiology to audiences with differing expertise.

Describing the myriad biological processes occurring in living beings over time, the science of physiology is complex and critical to our understanding of how life works. It spans many spatio-temporal scales to combine and bridge the basic sciences (biology, physics, and chemistry) to medicine. Recent years have seen an explosion of new and finer-grained experimental and acquisition methods to characterize these data. The volume and complexity of these data necessitate effective visualizations to complement standard analysis practice. Visualization approaches must carefully consider and be adaptable to the user's main task, be it exploratory, analytical, or communication-oriented. This thesis contributes to the areas of theory, empirical findings, methods, applications, and research replicability in visualizing physiology.

Our contributions open with a state-of-the-art report exploring the challenges and opportunities in visualization for physiology. This report is motivated by the need for visualization researchers, as well as researchers in various application domains, to have a centralized, multiscale overview of visualization tasks and techniques. Using a mixed-methods search approach, this is the first report of its kind to broadly survey the space of visualization for physiology. Our approach to organizing the literature in this report enables the lookup of topics of interest according to spatio-temporal scale. It further subdivides works according to any combination of three high-level visualization tasks: exploration, analysis, and communication. This provides an easily-navigable foundation for discussion and future research opportunities for audience- and task-appropriate visualization for physiology. From this report, we identify two key areas for continued research that begin narrowly and subsequently broaden in scope: (1) exploratory analysis of multifaceted physiology data for expert users, and (2) communication for experts and non-experts alike.

Our investigation of multifaceted physiology data takes place over two studies. Each targets processes occurring at different spatio-temporal scales and includes a case study with experts to assess the applicability of our proposed method.

At the molecular scale, we examine data from magnetic resonance spectroscopy (MRS), an advanced biochemical technique used to identify small molecules (metabolites) in living tissue that are indicative of metabolic pathway activity. Although highly sensitive and specific, the output of this modality is abstract and difficult to interpret. Our design study investigating the tasks and requirements for expert exploratory analysis of these data led to *SpectraMosaic*, a novel application enabling domain researchers to analyze any permutation of metabolites in ratio form for an entire cohort, or by sample region, individual, acquisition date, or brain activity status at the time of acquisition.

A second approach considers the exploratory analysis of multidimensional physiological data at the opposite end of the spatio-temporal scale: population. An effective exploratory data analysis workflow critically must identify interesting patterns and relationships, which becomes increasingly difficult as data dimensionality increases. Although this can be partially addressed with existing dimensionality reduction techniques, the nature of these techniques means that subtle patterns may be lost in the process. In this approach, we describe *DimLift*, an iterative dimensionality reduction technique enabling user identification of interesting patterns and relationships that may lie subtly within a dataset through *dimensional bundles*. Key to this method is the user's ability to steer the dimensionality reduction technique to follow their own lines of inquiry.

Our third question considers the crafting of visualizations for communication to audiences with different levels of expertise. It is natural to expect that experts in a topic may have different preferences and criteria to evaluate a visual communication relative to a non-expert audience. This impacts the success of an image in communicating a given scenario. Drawing from diverse techniques in biomedical illustration and visualization, we conducted an exploratory study of the criteria that audiences use when evaluating a biomedical process visualization targeted for communication. From this study, we identify opportunities for further convergence of biomedical illustration and visualization techniques for more targeted visual communication design. One opportunity that we discuss in greater depth is the development of semantically-consistent guidelines for the coloring of molecular scenes. The intent of such guidelines is to elevate the scientific literacy of non-expert audiences in the context of molecular visualization, which is particularly relevant to public health communication.

All application code and empirical findings are open-sourced and available for reuse by the scientific community and public. The methods and findings presented in this thesis contribute to a foundation of cross-disciplinary biomedical illustration and visualization research, opening several opportunities for continued work in visualization for physiology.

Abstract in Norwegianian

Det overordnede målet med denne avhandlingen er tverrfaglig anvendelse av medisinske illustrasjons- og visualiseringsteknikker for å utforske, analysere og formidle aspekter ved fysiologi til publikum med ulik faglig nivå og bakgrunn.

Fysiologi beskriver de biologiske prosessene som skjer i levende vesener over tid. Vitenskapen om fysiologi er kompleks, men samtidig kritisk for vår forståelse av hvordan levende organismer fungerer. Fysiologi dekker en stor bredde romlig-temporale skalaer og fordrer behovet for å kombinere og bygge bro mellom basalfagene (biologi, fysikk og kjemi) og medisin. De senere årene har det vært en eksplosjon av nye, avanserte eksperimentelle metoder for å detektere og karakterisere fysiologiske data. Volumet og kompleksiteten til fysiologiske data krever effektive strategier for visualisering for å komplementere dagens standard analyser. Hvilke tilnærminger som benyttes i visualiseringen må nøye balanseres og tilpasses formålet med bruken av dataene, enten dette er for å utforske dataene, analysere disse eller kommunisere og presentere dem.

Arbeidet i denne avhandlingen bidrar med ny kunnskap innen teori, empiri, anvendelse og reproduserbarhet av visualiseringsmetoder innen fysiologi. Først i avhandlingen er en rapport som oppsummerer og utforsker dagens kunnskap om muligheter og utfordringer for visualisering innen fysiologi. Motivasjonen for arbeidet er behovet forskere innen visualiseringsfeltet, og forskere i ulike anvendelsesområder, har for en sammensatt oversikt over flerskala visualiseringsoppgaver og teknikker. Ved å bruke søk over et stort spekter av metodiske tilnærminger, er dette den første rapporten i sitt slag som kartlegger visualiseringsmulighetene innen fysiologi. I rapporten er faglitteraturen oppsummert slik at det skal være enkelt å gjøre oppslag innen ulike tema i rom-og-tid-skalaen, samtidig som litteraturen er delt inn i de tre høynivå visualiseringsoppgavene data utforskning, analyse og kommunikasjon. Dette danner et enkelt grunnlag for å navigere i litteraturen i feltet og slik danner rapporten et godt grunnlag for diskusjon og forskningsmuligheter innen feltet visualisering og fysiologi. Basert på arbeidet med rapporten var det særlig to områder som det er ønskelig for oss å fortsette å utforske: (1) utforskende analyse av mangefasetterte fysiologidata for ekspertbrukere, og (2) kommunikasjon av data til både eksperter og ikke-eksperter.

Arbeidet vårt av mangefasetterte fysiologidata er oppsummert i to studier i avhandlingen. Hver studie omhandler prosesser som foregår på forskjellige romlig-temporale skalaer og inneholder konkrete eksempler på anvendelse av metodene vurdert av eksperter i feltet.

I den første av de to studiene undersøkes konsentrasjonen av molekulære substanser (metabolitter) ut fra data innsamlet med magnetisk resonansspektroskopi (MRS), en avansert biokjemisk teknikk som brukes til å identifisere metabolske forbindelser i levende vev. Selv om MRS kan ha svært høy sensitivitet og spesifisitet i medisinske an-

vendelser, er analyseresultatene fra denne modaliteten abstrakte og vanskelige å forstå også for medisinskfaglige eksperter i feltet. Vår designstudie som undersøkte oppgavene og kravene til ekspertutforskende analyse av disse dataene førte til utviklingen av SpectraMosaic. Dette er en ny applikasjon som gjør det mulig for domeneeksperter å analysere konsentrasjonen av metabolitter normalisert for en hel kohort, eller etter prøveregion, individ, opptaksdato, eller status på hjernens aktivitetsnivå ved undersøkelsestidspunktet.

I den andre studien foreslås en metode for å utføre utforskende analyser av flerdimensjonale fysiologiske data i motsatt ende av den romlig-temporale skalaen, nemlig på populasjonsnivå. En effektiv arbeidsflyt for utforskende dataanalyse må kritisk identifisere interessante mønstre og relasjoner, noe som blir stadig vanskeligere når dimensjonaliteten til dataene øker. Selv om dette delvis kan løses med eksisterende reduksjonsteknikker er det alltid en fare for at subtile mønstre kan gå tapt i reduksjonsprosessen. Isteden presenterer vi i studien DimLift, en iterativ dimensjonsreduksjonsteknikk som muliggjør brukeridentifikasjon av interessante mønstre og relasjoner som kan ligge subtile i et datasett gjennom dimensjonale bunter. Nøkkelen til denne metoden er brukerens evne til å styre dimensjonalitetsreduksjonen slik at den følger brukerens egne undersøkelseslinjer.

For videre å undersøke kommunikasjon til eksperter og ikke-eksperter, studeres i neste arbeid utformingen av visualiseringer for kommunikasjon til publikum med ulike nivåer av ekspertnivå. Det er naturlig å forvente at eksperter innen et emne kan ha ulike preferanser og kriterier for å vurdere en visuell kommunikasjon i forhold til et ikke-ekspertpublikum. Dette påvirker hvor effektivt et bilde kan benyttes til å formidle en gitt scenario. Med utgangspunkt i ulike teknikker innen biomedisinsk illustrasjon og visualisering, gjennomførte vi derfor en utforskende studie av kriteriene som publikum bruker når de evaluerer en biomedisinsk prosessvisualisering målrettet for kommunikasjon. Fra denne studien identifiserte vi muligheter for ytterligere konvergens av biomedisinsk illustrasjon og visualiseringsteknikker for mer målrettet visuell kommunikasjonsdesign. Særlig beskrives i større dybde utviklingen av semantisk konsistente retningslinjer for farging av molekylære scener. Hensikten med slike retningslinjer er å heve den vitenskapelige kompetansen til ikke-ekspertpublikum innen molekylær visualisering, som vil være spesielt relevant for kommunikasjon til befolkningen i forbindelse med folkehelseopplysning.

All kode og empiriske funn utviklet i arbeidet med denne avhandlingen er åpen kildekode og tilgjengelig for gjenbruk av det vitenskapelige miljøet og offentligheten. Metodene og funnene presentert i denne avhandlingen danner et grunnlag for tverrfaglig biomedisinsk illustrasjon og visualiseringsforskning, og åpner flere muligheter for fortsatt arbeid med visualisering av fysiologiske prosesser.

List of Publications

The following publications are included in this thesis:

- (A) **L. Garrison**, I. Kolesár, I. Viola, H. Hauser, and S. Bruckner, "Trends & Opportunities in Visualization for Physiology: A Multiscale Overview," *Computer Graphics Forum*, vol. 41(3), pp. 609-643, 2022.

Laura Garrison completed the majority of the literature search, taxonomy concept, and the writing of the report. Ivan Kolesár contributed to the literature search, taxonomy concept, an early paper draft, and implemented the visual tool that accompanies the report. Ivan Viola, Helwig Hauser, and Stefan Bruckner provided guidance and feedback over the course of the research.

- (B) **L. Garrison**, J. Vašíček, A. R. Craven, R. Grüner, N. Smit, and S. Bruckner, "Interactive Visual Exploration of Metabolite Ratios in MR Spectroscopy Studies," *Computers & Graphics*, vol. 92, pp. 1–12, 2020.

Laura Garrison contributed to the task and requirement characterization, tool design, case studies, and paper write-up. Jakub Vašíček provided the implementation of the visual tool in partial fulfillment of his Master's degree. Alex Craven provided implementation for the upstream processing steps, helped acquire and anonymize data, and advised over the course of the research. Renate Grüner and Noeska Smit provided guidance and feedback over the course of the project. Stefan Bruckner provided extensive guidance in the ideation, development, and writing phases of the research.

- (C) **L. Garrison**, J. Müller, S. Schreiber, S. Oeltze-Jafra, H. Hauser, and S. Bruckner, "DimLift: Interactive Hierarchical Data Exploration through Dimensional Bundling," *IEEE Transactions on Visualization & Computer Graphics*, vol. 27(6), pp. 2908-2922, 2021.

Laura Garrison and Juliane Müller contributed equally to the ideation of the method presented in this paper. Juliane provided the implementation of the application, and Laura developed the visual design, case studies, and write-up of the paper. Stefanie Schreiber advised on medical content, provided the CSVD clinical dataset for use, and participated in the CSVD case study. Steffen Oeltze-Jafra, Helwig Hauser, and Stefan Bruckner provided guidance and feedback over the course of the research.

*Paper B and Paper C together received an **Honorable Mention** for the **2021 MedVis (Karl Heinz Höhne) Award**, awarded by the professional group Visual Computing in Biology and Medicine, a section of the German Society for Computer Science, and sponsored by Brainlab AG in Munich.*

- (D) **L. Garrison**, M. Meuschke, J. Fairman, N. Smit, B. Preim, and S. Bruckner, "An Exploration of Practice and Preferences for the Visual Communication of Biomedical Processes," *Proceedings of EuroGraphics Workshop on Visual Computing for Biology & Medicine (VCBM)*, pp.1–12, 2021.

*Received **Honorable Mention for Best Paper** at the VCBM 2021 Workshop in Paris, France.*

Laura Garrison contributed the survey concept and design, survey assets, focus groups, analysis, and paper write-up. Monique Meuschke contributed the dataset for the blood flow and aneurysm visualization assets and advised on their production. Jennifer Fairman advised on production of medical illustration survey assets, survey design, and provided guidance and feedback on the paper. Noeska Smit, Bernhard Preim, and Stefan Bruckner provided guidance and feedback through all phases of the research.

- (E) **L. Garrison** and S. Bruckner, "Considering Best Practices in Color Palettes for Molecular Visualizations," to appear in *Design X Bioinformatics Special Issue of Journal of Integrative Bioinformatics*, 2022.

Laura Garrison developed the concept, and conducted the small study and write-up of this work, with guidance and feedback from Stefan Bruckner.

The following publications are also related to this thesis:

- (1) **L. Garrison**, J. Vašíček, R. Grüner, N. Smit, and S. Bruckner, "A Visual Encoding System for Comparative Exploration of Magnetic Resonance Spectroscopy Data," *Proceedings of EuroVis (Posters)*, 2022.

*Awarded **Best Poster** at EuroVis 2019 Conference in Porto, Portugal.*

***Paper B** and **Paper 2** developed from the ideas presented in this poster.*

- (2) **L. Garrison**, J. Vašíček, R. Grüner, N. Smit, and S. Bruckner, "SpectraMosaic: An Exploratory Tool for the Interactive Visual Analysis of Magnetic Resonance Spectroscopy Data," *Proceedings of EuroGraphics Workshop on Visual Computing for Biology & Medicine (VCBM)*, p. 1–10, 2019.

***Paper B** is an extension of this work.*

- (3) H. Bartsch, **L. Garrison**, S. Bruckner, A. Wang, S. F. Tapert, and R. Grüner, "MedUse: A Visual Analysis Tool for Medication Use Data in the ABCD Study," *Proceedings of EuroGraphics Workshop on Visual Computing for Biology & Medicine (VCBM) (Short Papers)*, p. 97–101, 2019.

- (4) J. Müller, **L. Garrison**, P. Ulbrich, S. Schreiber, S. Bruckner, H. Hauser, and S. Oeltze-Jafra. "Integrated Dual Analysis of Quantitative and Qualitative High-Dimensional Data," *IEEE Transactions on Visualization & Computer Graphics* vol. 27(6), pp. 2953-296, 2021.

- (5) Y. S. Kristiansen, **L. Garrison**, and S. Bruckner, "Semantic Snapping for Guided Multi-View Visualization Design," *IEEE Transactions on Visualization and Computer Graphics* vol. 28(1), pp. 43-53, 2022.

- (6) M. Meuschke, **L. Garrison**, N. N. Smit, B. Bach, S. Mittenentzwei, V. Weiß, S. Bruckner, K. Lawonn, and B. Preim, "Narrative Medical Visualization to Communicate Disease Data," in review to *Computers & Graphics*, 2022.
- (7) Y. S. Kristiansen, **L. Garrison**, and S. Bruckner, "Content-Driven Layout for Visualization Design," in review to *International Symposium on Visual Information Communication and Interaction (VINCI)*, 2022.

Contents

Scientific Environment	i
Acknowledgements	iii
Abstract	v
Abstract in Norwegian	vii
List of Publications	ix
I Overview	1
1 Introduction	3
1.1 Problem Statement	3
1.2 Background	5
1.2.1 Physiology Primer	5
1.2.2 Visualization Tasks Primer	6
1.2.3 Color Primer	8
1.3 Scope and Contributions	9
1.4 Thesis Structure	10
2 Research Context	13
2.1 Visualization of Physiology	13
2.2 Discovery-Centric Visualization	15
2.3 Communication-Centric Visualization	18
3 Contributions	23
3.1 Visualization for Physiology	24
3.1.1 Mixed Methods Literature Search and Classification	25
3.1.2 Challenges and Opportunities	28
3.2 Exploratory Analysis of Multifaceted Data	30
3.2.1 Exploratory Analysis of MRS Data	30
3.2.2 Exploratory Analysis of High Dimensional Clinical Cohort Data	40
3.3 Visual Communication of Physiology	48
3.3.1 Practice & Preferences for Experts and Non-Experts	48
3.3.2 Color Semantics in Molecular Visualization	57
3.4 Research Replicability	60

4 Discussion & Outlook	63
II Scientific Results	67
A Trends & Opportunities in Visualization for Physiology: A Multiscale Overview	69
A.1 Introduction	70
A.2 Physiology Background	73
A.3 Scope and Methodology	74
A.4 Taxonomy and Overview	77
A.4.1 Spatio-Temporal Organization	78
A.4.2 High-Level Visualization Task	79
A.5 Molecular Function	81
A.5.1 Molecular Dynamics	82
A.5.2 Molecular Interactions	84
A.5.3 Molecular Pathways	85
A.6 Cellular Function	87
A.6.1 Cellular Dynamics	88
A.6.2 Cellular Interactions	91
A.7 Tissue Function	92
A.7.1 Tissue Dynamics	94
A.7.2 Tissue Interactions	95
A.8 Organ Function	97
A.8.1 Blood Flow	98
A.8.2 Heart Function	100
A.8.3 Lung Function	102
A.8.4 Brain Function	103
A.8.5 Other Organ Function	105
A.9 True Multiscale and Beyond	106
A.10 Discussion	108
A.11 Research Outlook	110
A.12 Conclusion	111
A.13 Cited Literature in Report	112
B Interactive Visual Exploration of Metabolite Ratios in MR Spectroscopy Studies	115
B.1 Introduction	116
B.2 Related Work	117
B.3 Background	118
B.4 Task and Requirement Analysis	119
B.4.1 Task Analysis and Abstraction	120
B.4.2 Design Requirements	122
B.5 SpectraMosaic Workflow and Interface	123
B.6 Spectral Ratio Heatmap	124
B.7 Implementation	129
B.8 Case Study	129

B.9	Discussion and Limitations	133
B.10	Conclusions and Future Work	134
C	DimLift: Interactive Hierarchical Data Exploration through Dimensional Bundling	135
C.1	Related Work	137
C.2	DimLift Approach	139
C.2.1	Creating Expressive <i>Dimensional Bundles</i>	141
C.2.2	Visual Encodings	143
C.2.3	Lifting Expressive Dimensions	146
C.2.4	Handling Mixed and Missing Data	150
C.3	Case Studies	151
C.3.1	Plant Ecology	152
C.3.2	USDA National Nutrient Data	153
C.3.3	Clinical Cohort: Cerebral Small Vessel Disease	154
C.4	Discussion	157
C.5	Conclusions	159
D	An Exploration of Practice and Preferences for the Visual Communication of Biomedical Processes	161
D.1	Introduction	162
D.2	Related Work	162
D.3	Abstraction Constructs	164
D.4	Study Design	165
D.4.1	Survey Scenarios	167
D.4.2	Survey Visual Assets	167
D.4.3	Survey Design Structure	172
D.4.4	Survey Recruitment	172
D.5	Study Findings	172
D.6	Discussion	176
D.7	Research Opportunities	179
D.8	Conclusion	179
E	Considering Best Practices in Color Palettes for Molecular Visualizations	181
E.1	Introduction	182
E.2	Related Work	185
E.3	Color Choices in Molecular Visualization	188
E.3.1	Color Considerations	188
E.3.2	Color Strategies: A Contemporary Sample	189
E.4	Considerations for Molecular Coloring Best Practices	192
E.5	Limitations	194
E.6	Challenges & Outlook	195
E.7	Conclusion	196
	Bibliography	199
	Appendix A	235

Appendix D

239

Part I
Overview

Chapter 1

Introduction

That thing the nature of which is totally unknown to you is usually what you need to find, and finding it is a matter of getting lost.

Rebecca Solnit

Physiology describes the interconnected biochemical processes that sustain life. Life is complicated—millions of reactions occur within our bodies at any given time, allowing us to extract energy from the food that we consume, breathe, move, and more. Understanding more about these processes in humans can lead to, for example, improved communication between experts and general audiences, promote drug development that is more targeted and effective, and advance the quality of healthcare and individual health management. The study of physiology unifies the basic disciplines of biology, chemistry, and physics with medicine and numerous other connecting disciplines, such as systems biology. Advances in hardware and software have led to new imaging paradigms at higher resolution, more sophisticated simulations, and new experimental methods able to capture processes in real time *in situ*, such as spatial transcriptomics [474]. The volume and complexity of these data necessitate new methods for both experts and non-experts to explore, analyze, and communicate salient information. The work in this thesis is motivated by the profusion of multifaceted, multidimensional physiological data that necessitate new visual approaches for their interpretation and communication from a cross-disciplinary visualization and biomedical illustration perspective. This introductory chapter describes the motivation for this thesis and provides relevant background on physiology, visualization tasks, and a brief primer on color theory. A detailed outline of the scope and contributions follows, along with an overview of the structure of the remainder of this thesis.

1.1 Problem Statement

New experimental and computational techniques are enabling researchers to map the spatio-temporal organization of the body and dynamic behaviors of molecules, cells, tissues, and organs in unprecedented detail [30, 137, 211, 498, 527]. These techniques

and research initiatives span multiple disciplines. Consequently, identifying and bringing together the myriad visualization opportunities that these new methods present requires a multidisciplinary perspective.

Data describing aspects of physiology are complex due to their often multifaceted, high-dimensional nature. Experts need an effective means to **discover** interesting features in these data in order to perform tasks, such as identifying new biomarkers for disease. Finding such interesting features or patterns can be particularly challenging as the dimensionality of the data increase. Keeping track of different, possibly interesting, aspects is also challenging, e.g., comparing and aligning qualitative and quantitative data features. While already difficult in the case of an individual, these challenges are compounded in the case of a group, as occurs in clinical cohort data. Visual analysis approaches that integrate statistical methods with human interaction are useful to reduce the analysis space, but maintaining a meaningful connection to the input data and balancing user agency and insights with data complexity remains an ongoing challenge.

Communicating the key insights of physiology data to different application domains, e.g., clinical research, or to the public is yet another grand challenge in the visualization of physiology. Effective communication necessitates finding an appropriate level of abstraction to convey information in a way that is comprehensible to the intended user or audience. Limited visualization research addresses the practice of visual communication for physiology, often falling to disciplines outside of the core visualization field, e.g., biomedical illustration [148]. Better understanding of these practices and preferences is necessary in visualization research to advance our efforts in communicating the complexity of physiology to different stakeholders.

Biomedical illustrations are ubiquitous in visual communications for their tuneability to the degree of abstraction that is most appropriate for conveying the intended message to the target audience. They are furthermore employed when the source data are absent, of poor quality, or otherwise inadequate to communicate the desired information. However, in some cases illustrations can be too abstract due to their lack of direct connection to the source data, or lack consistent semantic meaning in favor of exercising creative license. While conventional visualization approaches often tackle a problem from a technical perspective driven by data characteristics and user task demands, approaches from biomedical illustration more often frame the problem from a visual design-centered perspective. The latter perspective prioritizes visual composition and storytelling methods that are beneficial for communication-based tasks. While illustrative visualization draws inspiration from illustration techniques in its approach [412], further work exploring the relationship and integration of data- and illustration-driven techniques is necessary. Such work can improve semantic consistency and create visual narratives that better facilitate the communication objective of a given visualization.

Our cross-disciplinary perspective in this thesis enables us to address these challenges and gaps between conventional visualization approaches and biomedical illustration techniques in visualization for physiology. This perspective is enabled in part by the increased recognition of physiology as a multidisciplinary field of study, and by the increased societal demands for compelling and understandable tools to craft visualizations of physiology that have begun to span disciplines in new ways.




Figure 1.1: Schematic of spatio-temporal scale ranges in which physiological processes occur (from Paper A).

1.2 Background

1.2.1 Physiology Primer

Human physiology requires a careful balancing act, known as homeostasis, of a multitude of processes occurring over a broad span of time and space, as summarized in Fig. 1.1. The **cell** is the smallest entity in the human body with the functional characteristics for independent life. Cells contain inorganic **molecules**, such as water and ions, as well as organic molecules, such as proteins, that are involved in the processes necessary for survival. Genes are the basic unit of heredity in cells that are made up of DNA and which encode the synthesis of RNA, which then is translated to build proteins. Genes, proteins, and other molecules interact through sequences of reactions that form pathways, which ultimately form networks contributing to specific cellular functions. For example, a metabolic pathway is any set of chemical reactions that converts food into usable energy and removes waste from the cell [180]. The molecular reactants, products, and intermediates of this type of pathway are called metabolites. Metabolites are the main focus of one of the visualization approaches, *SpectraMosaic* (Paper B), that we discuss in this thesis. Molecular structures are typically nanometers to hundreds of nanometers in size, and **molecule-scale processes** can occur over a broad temporal range from femtoseconds, e.g., bond vibrations between atoms within a molecule, to seconds, e.g., global motions or the chain of reactions in a pathway [193], to minutes and hours in the case of large-scale pathways, e.g., gene expression, signal transduction, and metabolism [31].

Cells contain specialized cellular structures called organelles that perform specific functions within the cell. For example, the nucleus houses the cell's genes, while the mitochondrion is often referred to as the powerhouse of the cell [180]. Specialized molecules, known as receptors, as well as channels and other structures facilitate com-

munication and exchange of nutrients between the cell and the outside world through a semipermeable membrane. Cells vary widely in size, but average around tens of microns with **cell-scale processes** occurring over a temporal range of milliseconds, e.g., the generation of an action potential, to a day for a complete cycle of human cell growth and division.

Tissue is formed from groups of specialized cells with shared properties and functions. There are four primary tissue types in the body, each with unique physiological properties reflective of their function: muscle, epithelial, connective, and nerve tissue. For example, connective tissue supports and protects the various structures of the body, in addition to storing fat and assisting in structural repairs. Abnormal tissues arise when cells take on different characteristics than the surrounding tissue in which they occur, as is the case in cancerous tumors. **Tissue-scale processes** span hundreds of microns to millimeters in size and range in time from milliseconds, e.g., signal propagation, to weeks or months, as in tissue growth and development.

Organs are composed of a collection of tissues that perform specific function(s) in the body. For example, the main function of the lungs is gas exchange, whereby freshly-oxygenated air is brought into the body and deoxygenated air is removed. Most human organs measure in the range of centimeters. **Organ-scale processes**, such as a complete heartbeat or a complete breath cycle, occur in roughly one second, but the lifespan of an organ may span the years the organism itself survives.

A **system** is a group of organs that work in concert to perform one or more functions in the body. For example, the respiratory system includes the lungs, airways, and accessory muscles that bring air in and out of the body. Our organ systems are interdependent. The respiratory system cannot function without the cardiovascular system, which includes the blood vessels which carry oxygen throughout the body. The normal functioning of an **organism** depends upon the systems of the body working in harmony.

Physiology is valuable to study at the **population** scale to understand general patterns of processes across larger groups, or to follow evolutionary processes over even longer time scales. This thesis touches on small-scale population physiology in two applications that we discuss in more detail in Chapter 3: in the study of metabolite ratios of small cohorts in *SpectraMosaic* (Paper B) and in larger cohorts with a broader range of data types in *DimLift* (Paper C).

1.2.2 Visualization Tasks Primer

Visualization is a powerful tool used to derive insights from data. These data sources come from, and can be used in, a number of different ways. While computer science tends to think of visualization as a means to augment human capabilities through visual representations of a dataset [354], these representations may also be *data-inspired*, often to convey a clear message to a chosen audience. The latter is true for much of the work from the centuries-old field of biomedical illustration. Just as in visualization, advances in technology have introduced new opportunities for biomedical illustrators to use data in their craft. While many biomedical illustrators rely on observation, just as Renaissance artists like Andreas Vesalius and Leonardo Davinci did to illuminate the structure and function of the human body, modern-day biomedical illustrators can incorporate data more directly into their workflow, e.g., modeling and animating a beat-

ing heart from CT data. This thesis investigates the overlap and respective influences of visualization and biomedical illustration techniques to address user tasks in visualizations for physiology.

Effective visualizations are those designed to solve a particular user task or set of tasks for a given user and data type. Adapting from Brehmer & Munzner's multilevel typology of abstract visualization tasks [66], these tasks can be categorized at a high-level as **exploratory**, **analytical**, or **communication-oriented**. A single visualization may address these three tasks to varying degrees.

Exploratory tasks often occur when the user is unsure of the contents of the data. In the context of the visualization pipeline, the user, often an expert, typically wishes to minimally abstract the data to produce a visual mapping that remains close or is easily trackable to the source format. Through exploration the user can familiarize themselves with the characteristics and features of the data prior to a more detailed analysis. Exploratory visualization strategies are often direct visualizations, e.g., direct volume rendering or time-lapse image sequences. Interaction techniques allow users to browse the data for interesting features, and focus+context techniques enable users to explore specific aspects of the data within the context of the overall dataset.

Analysis tasks can arise when the user has familiarized themselves with the intrinsic characteristics of the data and now wishes to organize and extract meaning from these data. Analysis and exploration are closely linked through the concept of discovery. A user may begin with an exploratory approach to generate a hypothesis, then perform low-level analytical tasks incorporating statistical methods to verify their hypothesis. This workflow shows the human as an integral part of the analysis process, where insights are generated by human-machine cooperation. In the visualization pipeline, analysis may produce new artefacts through transformation, derivation, and abstraction of the data [10]. Visualization strategies to accomplish analysis-based tasks often employ interaction techniques, such as brushing and filtering, to sift through the data and identify features of interest. Linked multi-view visualizations that include common chart types such as bar and scatter plots are common approaches to facilitate detailed analytical tasks.

Key data concepts can be emphasized or summarized through **communication tasks** such as presentation, education, or enjoyment. These tasks serve a particular audience or user group, e.g., experts between domains or the general public. Visualizations developed for communication are often further abstracted from the data than in analysis- or exploration-oriented tasks, often incorporating cinematic or storytelling elements to convey the author's interpretation of the data. While nearly all publications include figures to communicate scientific results, for the work presented in this thesis we identify uses of visualization for communication beyond what is achievable with standard, out-of-the-box tools. Illustrative visualization techniques, inspired by biomedical illustration, are accessible strategies for communication as they can tell a simpler story than, e.g., a complete volume rendering strategy may tell.

Many visualizations cannot be defined through one of these tasks alone, and instead often aim to achieve a combination of tasks. This thesis describes visualization approaches and findings from case studies that achieve a combination of user tasks, e.g., exploration and analysis.

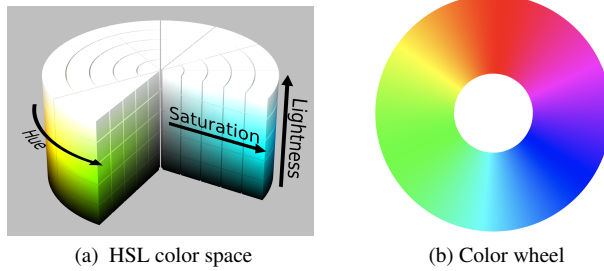


Figure 1.2: Color. (a) Source: https://en.wikipedia.org/wiki/HSL_and_HSV. (b) Adapted from: https://commons.wikimedia.org/wiki/File:RGB_color_wheel_360.svg (from Paper E).




Figure 1.3: Three common color harmony rules using blue as the base color. Created in Adobe Color [6]: <https://color.adobe.com/create/color-wheel> (from Paper E).

1.2.3 Color Primer

Color can help users achieve their desired tasks more easily [354]. While color can be represented through numerous models, such as RGB (red, green, blue), CMYK (cyan, magenta, yellow, key: black), or HSL (hue, saturation, lightness) [417], for our purposes the HSL color model is the most useful and intuitive way to consider color. The 3D cylinder in Fig. 1.2a depicts the three parameters of this model. Hue describes a base color, e.g., green, that is localized by angle around the color wheel shown in Fig. 1.2b. Saturation specifies the purity of a hue, and range from no saturation, i.e., grey, at the center of the cylinder to complete saturation, e.g., pure green, the cylinder's outer perimeter. Lightness defines the brightness of a color, ranging from no light, i.e., black, at the base of the cylinder to fully light, i.e., white, at the top of the cylinder. Blending black into a color produces a shade, while mixing a color with white produces a tint.

The combination of colors used to design a visualization is called a color palette. Several color harmony rules aid in creating aesthetically pleasing palettes. Drawing from Itten's seven models of color contrast [222], harmony rules may be *monochromatic*, *analogous*, or *complementary*, among others. *Monochromatic* palettes are created from tints and shades of a single color, as shown in Fig. 1.3a. *Analogous* palettes consist of adjacent colors on the color wheel, as in Fig. 1.3b. *Complementary* palettes

consist of opposite colors on the color wheel, as in Fig. 1.3c. Colors in a palette can be used strategically to, e.g., establish a visual hierarchy or to guide the eye through a narrative.

1.3 Scope and Contributions

The overarching theme of this thesis is the cross-disciplinary application of biomedical illustration and visualization to address challenges in exploring, analyzing, and communicating aspects of physiological processes to audiences and users with different degrees of expertise. We restrict our scope according to the following **field**, **actors**, and **cases**:

1. Our work is restricted to visualization primarily for the **field** of human physiology, although some instances noted in our state-of-the-art report come from model organisms.
2. The **actors** we focus on are domain experts and non-experts, i.e., individuals who either classify as a lay audience or who lack expert-level knowledge of the particular process being visualized.
3. The **cases** this thesis deals with are restricted to high-level user tasks, predominantly exploratory analysis or communication of physiology through visualization. The theory, empirical findings, methods, and applications that this thesis presents are developed from these high-level tasks or combinations of tasks.

This thesis makes several contributions in the areas of **theory**, **empirical findings**, **method**, **application**, and **research replicability**. We outline individual contributions according to these areas below:

Theory:

1. A novel classification system that identifies visualizations according to the high-level visualization task they serve, and to the spatio-temporal scale of the physiological process(es) that they cover, as displayed in Fig. 3.1 (Paper A).
2. The groundwork for a set of guidelines and best practices for semantically-meaningful coloring of molecular visualizations (Paper E).

Empirical Findings:

3. A state-of-the-art report that surveys both within and beyond the core visualization venues to provide a view into mature and open opportunities in visualization research for physiology according to spatio-temporal scale and high-level visualization task (Paper A).
4. A detailed review of magnetic resonance spectroscopy (MRS) data characteristics and abstraction of spectral analysis tasks identified from domain expert collaboration (Paper B).

5. A short study on the role of algorithm choice in data imputation for high-dimensional, mixed-type datasets in the context of our *DimLift* method (Paper C).
6. An asset library cataloguing standard practices for the creation of biomedical illustrations and visualizations of common concepts in physiology (Paper D).
7. A survey of expert and non-expert preferences for the visual communication of common concepts in physiology (Paper D).
8. A preliminary study of the color palettes used for molecular visualizations from two domain perspectives: (1) biomedical illustration and (2) application domains, e.g., structural biology, systems biology (Paper E).

Method:

9. A heatmap matrix display with nested glyphs to explore ratios of metabolite concentrations from magnetic resonance spectroscopy (MRS) data (Paper B).
10. The *DimLift* method for semi-automated creation and editing of semantically-related dimensional groupings, called *dimensional bundles*, to explore subtle patterns in high-dimensional data. This method semantically connects and tracks data transformations through visual mappings and interactions that allow on-the-fly recomposition of dimensional bundles (Paper C).

Application:

11. *SpectraMosaic*, a visual exploratory analysis pipeline and tool which provides an interface for linking of structural, spectral, and patient data. It includes group creation and uncertainty communication. This application integrates our heatmap matrix display with nested glyphs to explore metabolite ratios through different attributes. This approach is validated with a clinical case study and feedback from three MR spectroscopy research experts (Paper B).
12. An application interface for the *DimLift* method that allows experts to interact with complex, high-dimensional data. We validate this approach in a physiological context through a cerebral small vessel disease case study (Paper C).

Research replicability:

13. All software for the applications developed is open-source and certified replicable through the Graphics Replicability Stamp Initiative (Papers B and C). Empirical data packages and accompanying visualizations are also publicly available (Papers A, D, and E).

1.4 Thesis Structure

This thesis is divided into two main parts. **Part I** summarizes the motivations, challenges, and scientific results of this thesis. **Part II** contains the publications that contribute to this thesis, which remain unchanged apart from adjustments to fit the format

of this template. We also streamline the bibliographies of the respective publications, combining them to a single bibliography at the end of the thesis.

Part I is organized as follows: Chapter 1 (the current chapter) describes the problem space, provides necessary background material, and outlines the scope and contributions of this thesis. Chapter 2 provides an overview of related work in the following areas: visualization of physiology, discovery-centric visualization, and communication-centric visualization. Chapter 3 builds on related work to discuss the contributions of this thesis in the areas of theory, empirical findings, method, application, and research replicability. Chapter 4 concludes Part I and discusses opportunities for future work.

Part II includes five publications that provide further details on the contributions of this thesis.

- Paper A corresponds to contributions 1, 3, and 13.
- Paper B details contributions 4, 9, 11, and 13.
- Paper C provides further details for contributions 5, 10, 12, and 13.
- Paper D corresponds to contributions 6, 7, and 13.
- Paper E expands on contributions 2, 8, and 13.

Chapter 2

Research Context

Do you think miners stand around all day talking about how hard it is to mine for coal? They do not. They simply dig.

Cheryl Strayed

In this chapter we discuss works related to this thesis in the areas of (1) visualization of physiology, (2) discovery-centric visualization, and (3) communication-centric visualization.

2.1 Visualization of Physiology

Many of the theoretical frameworks guiding visualization design are constructed around data characteristics or user tasks. Munzner [354] and Tominski & Schumann [503] consider both elements in their framing of the purpose of visualization as the exploration, description, explanation, communication, and/or presentation of data. Munzner provides a nested model of visualization [353] that enables visualization researchers to validate their design choices according to data, user, and task at four levels, ranging from algorithm design up to domain characterization. A task-specific framework is that proposed by Brehmer & Munzner [66], who present a multilevel task typology on the what, why, and how of visualization tasks. Our work draws inspiration from these frameworks in our consideration of visualization tasks in the context of physiology.

Physiology has received extensive attention from the visualization community. However, this has been in a fragmented, unevenly distributed form across numerous subtopics, data sources, and visualization techniques. Few works extend their focus beyond one or two scales, e.g., only molecular-scale [267], molecular- and cellular-scale [169], or organ-scale [287]. Surveys on a particular data type, e.g., PC-MRI by Köhler et al. [260], are similarly restricted in scope. This paints a limited picture of physiology's true multiscale nature. Technique surveys, including by McGee et al. [331] on multilayer network visualization, Bach et al. [33] on space-time cubes, or Preim et al. [401] on medical animation, may cover multiple scales, but physiology is not the primary focus and is often just one application area of many under discussion.

Still others have surveyed the multiscale challenges in visualizing biomedical data from a high-level perspective [10, 79, 245, 330, 523], though again without a specific focus on physiology. The bioinformatics and systems biology domain has perhaps come the closest to capturing the multiscale visualization challenges for physiology from the molecular- to population-scale. However, these works are relatively brief and lack detailed discussion of visualization tasks due to their different domain focus [452], or retain a heavier focus on particular data types or scales, as in O'Donoghue et al.'s [376] review on the multiscale challenges of visualizing omics and imaging data. The fragmented nature of these various reports on the topic of physiology as a whole motivated our work for a broader-scope study that outlines mature and open opportunities in visualization research for physiology (Paper A).

Color in Visualization of Physiology. Color is a valuable tool in the development of any visualization, with numerous works making use of color to facilitate user tasks in visualizations of physiological phenomena. Surveys by Silva et al. [465] and Zhou et al. [579] provide an overview of color guidelines and use of color maps from a general visualization perspective, many of which are applicable to physiology. Particularly ubiquitous at the organ-scale is the use of rainbow maps to indicate physiological parameters, and Borland et al. [57] discuss the controversies and pitfalls of this particular color map in detail. Numerous tools enable the creation and selection of different color maps and palettes for use in a visualization, including Colourmap Hospital [128], ColorBrewer [185], Colorgorical [173], and Adobe Color [6]. Many of these tools are designed for basic charts, rather than complex 2D or 3D shapes or environments, and can require discretion when applied to specific topics in physiology. Furthermore, these tools can still offer a staggering array of options and may demand a degree of expertise in color theory from the content author.

Color in Molecular-Scale Visualization. Although valuable across all scales of physiology, color application is a particularly open and challenging research area at the molecular-scale. Illumination models, such as Hermosilla's diffuse color bleeding model [195], may boost colors in a visualization to provide structural cues on a molecular surface that are imperative for the identification of binding sites in molecular reactions. Szafer et al. [490] have also shown the value of ambient occlusion and directional lighting when interpreting molecular surfaces that are in shadow, finding that stylized rendering approaches hinder interpretability. Furthermore, since shadows diminish luminance range, this study found luminance-varying ramps more performant than isoluminant ramps. These works suggest the assignment of high luminance values to molecules of interest or in shadow to direct attention and readability of structures of interest. The application of high luminance values to structures of interest is common in biomedical illustration, and has carried to other visualization works such as Waldin et al.'s [535] technique for color application that spans atomic to whole-virus resolution. Using a systematically adjustable color scheme across different scales through primarily hue shift, their method enables visualization of structures of interest at a particular scale, e.g., atoms or the structural compartments of a virus. They use an analogous palette of saturated colors, and use luminance to draw attention to salient features, e.g., the quantity of a type of amino acid in a protein or a protein's secondary structures. Klein et al. [256] take a similar multiscale adaptive coloring approach in the context of

microtubule dynamics. Similar to Waldin et al., the initial color assignment is largely arbitrary, and focuses on structure rather than function. Based in part on the ideas introduced in these works, Paper E of this thesis explores considerations for the development of semantic coloring guidelines in molecular visualizations. The intent of such guidelines is to facilitate communication, as well as discovery-related exploration and analysis tasks.

2.2 Discovery-Centric Visualization

Visualizations of physiological phenomena for the purpose of discovery orient towards generating and verifying hypotheses [66] through exploration and analysis. The focus of such work is often for research or clinical experts. Visual exploration and analysis of multifaceted, high-dimensional data is a grand challenge in the visualization research community, with applications across numerous domains. Discussion of efforts in this general area are beyond the scope of this thesis, but are detailed with a survey of advances in recent years by Liu et al. [307]. This thesis explores strategies for visual exploration and analysis of multifaceted data on two levels. The first level is specific to a particular data type that is little-explored in the visualization community: magnetic resonance spectroscopy (MRS) data. This is the focus of Paper B. The second level looks more broadly at strategies to visualize complex data that include a variety of different data sources, as is often the case in clinical cohort data. This more general level of multifaceted data exploration is the focus of Paper C.

Visualization of MRS Data. Magnetic resonance spectroscopy (MRS) is an *in vivo*, non-invasive technique used to estimate the concentrations of metabolites in a tissue region. These concentrations provide insights to the activity of metabolic pathways in that area. Acquisition techniques include single-voxel spectroscopy (SVS) or multi-voxel spectroscopy (CSI), where single-voxel techniques cover a smaller area and are more suited for detailed, quantitative metabolite analysis in a smaller region [513]. Paired with structural imaging methods such as magnetic resonance imaging (MRI), MRS shows potential to improve clinical diagnosis and treatment monitoring of numerous conditions of the central nervous system [517]. A major challenge in visualizing MRS data is that each voxel acquisition is itself a multivariate dataset. Our work in Paper B draws inspiration from tools such as InSpectr [14], which employs comparative visualization techniques [163] and multiple linked views to provide insights into the composition of a multivariate, multimodal data source. We similarly combine acquisition techniques, but in a different domain context with a different analytical focus. Spectral Similarity Maps extend the InSpectr framework [141], which adopt isosurface similarity maps proposed by Bruckner and Möller [72] to show correlations between spectra. Rather than mapping energy correlations, we map metabolite ratios using a similar visual encoding.

Prior approaches to visualizing MRS data have generally been limited to a subset of metabolites at a time, and are summarized in a state-of-the-art report by Nunes et al. [372]. Highlights from this report include work by Feng et al. [136], who present scale-driven data spheres (SDDS), a technique using colored spheres to indicate metabolite abundance in a voxel. Further work expanded this approach to include par-

allel coordinates and scatter plots to compare metabolites concentrations, although this remains limited to a subset of metabolites. Marino and Kaufman [321] combine MRI with MRS and positron emission tomography (PET) data in a direct volume visualization to understand the dimensions and characteristics of a prostate tumor. However, their use of MRS is limited to a single metabolite ratio, and is constrained to a single individual at one time point. Nunes et al. [373] combine ComVis [328] and MITK [559] in a visual analysis framework for radiotherapy that uses metabolite values to identify a biological target volume. Retention of spectra is not the focus of this work, and it provides limited functionality for comparing metabolites. Jawad et al. [229] present a system for composition analysis of segmented brain tissue to identify the metabolic signatures of brain tumors, with a follow-on reproducibility study [230]. This tool is optimized for multivoxel data, and focuses on statistical outcome measurements. Further work by Jawad et al. [228] describes an approach for the comparative analysis of cohort single-voxel spectroscopy data. They employ violin and parallel coordinate plots to analyze spectral metabolite relationships. Our work in Paper B uses similar data inputs and processing tools in our *SpectraMosaic* application. While our work uses similar abstract 2D visualization techniques as these various approaches, we are not restricted to a particular analysis use case of MRS data, e.g., tumor composition identification. We furthermore allow for the simultaneous comparison of all metabolites in a sample with focus+context facilities to compare different regions, across time, between individuals, or in different brain activity states.

While our first iteration of *SpectraMosaic* focuses on exploratory analysis of single-peak metabolites directly from spectral graphs [150], our work in Paper B expands this application for a complete analysis of simple and complex metabolite signatures and integrates robust MR spectroscopy quantification tools. The practical usability of the tool is also improved, with facilities for assigning group membership and displaying the variation of the input data that align with several recent recommendations from Truong & Duncan [508] for visualizing MRS data.

Visualization of Multifaceted Mixed-Type Data. Expanding beyond a specific data type related to physiology, we now discuss work that broadly relates to the visualization of multifaceted, mixed-type data that are typical to more complex studies in physiology. Dimensionality reduction methods play a large role in visualization strategies for these data.

Sacha et al. [429] provide an overview and classification of dimension reduction methods as used in visual analysis. Our *DimLift* approach described in Paper C incorporates dimensionality reduction methods in a subset of the interaction scenarios they identified: data selection & emphasis, data manipulation, and feature selection & emphasis.

Similar work to our approach include Tatu et al. [493], who employ an interestingness-guided subspace search algorithm to identify subspaces for further visual analysis. However, their facilities for user-driven subspace composition are limited. Dowling et al. [110] present the SIRIUS system, which explores the interplay of dimension and item space in high-dimensional data with the help of multidimensional scaling (MDS), a nonlinear dimensionality reduction technique. Their system uses linked views between dimension and item space to visualize correlations. We similarly link between dimension and item space, but employ a linear dimensionality reduction method for

clearer data provenance. Guided workflows such as DimStiller by Ingram et al. [215] enable users to find a single global optimal composition through the dimensionality reduction process. In contrast, our approach does not consider a single optimum, but rather many possibilities that depend on the user's line of inquiry. Yuan et al. [576] combine a Dimension Projection Matrix, an extended scatterplot matrix, with a Dimension Projection Tree to explore data and dimension subspaces. While we similarly allow analysis of data subspaces at item and dimension level, our approach enables user exploration and comparison of subspaces for interesting correlations. Most inspiring to our approach in Paper C is the Dual Analysis approach by Turkay et al. [510], which enables the simultaneous exploration of dimensions and items for hypothesis generation. This work, along with a follow-on clinical application study [511], introduces a visual analysis workflow where users seamlessly move between the comparative analysis of dimensions and items to identify outliers and correlations of interest. Brushing and linking facilities provide clear visual feedback during the analysis process. Müller et al. [351] extend this approach to incorporate mixed data (continuous and categorical) with facilities for the visualization of missing data. Our *DimLift* method expands further on the reciprocity between dimension and item space by introducing *dimensional bundles* for the analysis of high-dimensional data.

The projection of relevant data features to low-dimensional space often yields results that are difficult to understand. Although principal component analysis (PCA) [394] is a well-known and broadly applicable method used in dimensionality reduction, it suffers from this interpretation issue. Müller et al. [357] discuss design solutions to visualize the connection between data inputs and results from a PCA. However, many of these solutions do not scale well with high-dimensional data. iPCA [235] is one solution that bridges PCA results to the original data. This approach uses multiple linked views to depict PCA results with interaction facilities for user adjustment of dimension contributions within any principal component. Any adjustments dynamically update the visualization of the PCA results. Our *DimLift* method similarly relies on visual elements and user interactions to connect original and derived dimensions, although we use this link to facilitate hypothesis formation rather than to explain the results of a PCA.

Visual Representations for Multifaceted Data. A number of visual representations enable the subsequent analysis and exploration of multifaceted data. This thesis focuses primarily on two strategies: heatmap matrices and parallel coordinates, both of which include nested elements, such as glyphs. We discuss work closely related to these strategies in the remainder of this section.

Numerous solutions have leveraged small related graphics series to visualize multifaceted data, as first introduced by Bertin [47]. We base our method and application from Paper B, *SpectraMosaic*, on this concept, and extend this by including a second layer of nested visual encodings. This is inspired by ATOM [390], a grammar for unit visualizations where individual data items are represented by unique visual marks (units) in a visual encoding system. PivotTable, which was extended by Polaris [482] and subsequently trademarked by Microsoft, enables exploration and analysis of multidimensional data with the flexibility to modify visual encodings, graphics, and the configuration of the table itself. Klemm et al. [257] build on this concept for linked vi-

sualization of image-centric heterogeneous cohort data. Our approach is related in that we allow on-the-fly reconfiguration of our matrix inputs.

Parallel coordinates are a popular means for representing multifaceted data [189]. Nested or hierarchical plots, adapted from the traditional flat parallel coordinates plot, can visualize and be used to evaluate structural relationships in the data. Numerous parallel coordinates solutions aggregate data according to item relatedness, in part as a means to reduce clutter and noise in the plot [24, 80, 143, 144, 418, 518, 533]. These methods focus on the hierarchical construction of sets of items, while we aim one level above this to hierarchically construct sets of dimensions. Several approaches use parallel coordinates to visualize dimension-level aggregation, which are created through algorithmic methods or pre-defined data hierarchies. These methods allow varying degrees of user interaction. For example, Dunica et al. [112] and Wang et al. [540] use parallel coordinates to visualize results of a single-run PCA, where each axis represents a different principal component. While we similarly incorporate principal components, our approach iteratively produces principal components that consist of subsets of related dimensions.

Approaches to parallel coordinate dimension hierarchies often incorporate other plots. These are either separate or directly integrated into the parallel coordinates plot. Huang et al. [207] present hierarchical clusters of dimensions in parallel coordinates through dendrograms that attach to each axis. DOFSA [571] and InterRing [572] are connected tools enabling interactive visual exploration and modification of hierarchical data. These modifications are made on InterRing and link to other plots, e.g., parallel coordinates. By contrast, our method does not split user attention over different graphical interfaces. Furthermore, the DOFSA hierarchy is flattened in the parallel coordinates representation, with its order informed by the hierarchy constructed in InterRing. We do not flatten the hierarchy in this manner. Weidele [548] present the conditional parallel coordinates method, which connects and reveals additional dimensions to the range of a given parent dimension when certain conditions are met. Perhaps most visually similar to our approach, Brodbeck & Girardin [70] and Andrews et al. [18] create aggregated dimension axes in parallel coordinates plots. These axes may be expanded to reveal the contained dimensions. In contrast to these methods, our *DimLift* approach described in Paper C does not use pre-defined hierarchies, and instead allows for flexible dimension regrouping as hypotheses evolve.

2.3 Communication-Centric Visualization

The use of visualization design principles to communicate science through illustrative and data-driven means is another important facet of this thesis. Several works [167, 168, 231, 232, 239, 551] emphasize visualization for communication with the aid of *illustrative* techniques which often come from a practice-based perspective. Sousa et al. include illustrative approaches in their illustrative scientific visualization framework to help non-artists approach and solve visualization tasks [471]. This framework parallels a traditional illustration pipeline that begins with receiving and recording information, producing and refining sketches, followed by rendering and annotation. Similar to these works, we take a broad view of visualization for communication that includes illustrative and data-driven techniques. We additionally draw on prior ideas of

abstraction spaces and build on the existing body of qualitative visualization research in developing communication-oriented visualizations for physiology.

Visual Abstraction. Abstraction is inherent to visualization. Viola & Isenberg provide a formalization of abstraction in visualization [531]. Their definitions and updated formalization [530] of visual abstraction serve as the basis for the abstraction spaces which we discuss further in Chapter 3 and in detail in Paper D. Rautek et al. [412] describe abstraction as a powerful tool in visual communication, which can lend additional insights to one's data. Andrews [17] takes a similar view of abstraction from a biomedical illustration perspective, discussing instances where illustration is an optimal medium to visualize certain concepts. This includes removing “visual garbage,” i.e., irrelevant or distracting structures from the core message, or superimposing structures that may otherwise be impossible to do without illustration. Tufte's concept of “chart junk” is a more extreme instance of visual abstraction whereby a visualization is reduced to only its essence, with no additional visual elements beyond the data to drive a message [509]. Abstraction, when fit appropriately to the task, user, and data, lays the foundation for a successful visualization.

Communication With Color. While we discussed color earlier in this chapter as a tool to facilitate user tasks in visualization for physiology in the broad sense, we now discuss color use as it relates to communication tasks.

While biomedical illustrators follow perceptual design principles in coloring a visualization, they often take artistic license in their selection of a color palette. This is particularly true for molecular visualization. David Goodsell's watercolor paintings of molecular machinery are foundational to the practice of molecular visualization, where he uses color strategically to describe the spatial organization of molecules [165, 169]. However, he notes that most of his color choices are “*completely arbitrary and are chosen solely for aesthetic appeal* [168].” Biomedical illustrators also commonly employ color as a narrative or attention-drawing device to key elements in a visualization [231]. For example, Jenkinson et al. [232] applied low contrast, desaturated colors to context molecules while assigning highly saturated, complementary colors to focus molecules (ligand and receptor) to all treatments in their perceptual study on molecular scene complexity relative to learning outcomes. In their guide to the visual analysis and communication of biomolecular structural data, Johnson & Hertig [239] suggest a similar color assignment approach. In a Points of View series article, Wong [561] states that small objects in a scene require increased hue, saturation, and/or brightness in order to stand out in a visualization. He furthermore suggests a simple trick of squinting at a visualization to evaluate color evenness and visibility. Although these approaches are valuable resources for using color to guide a narrative and improve the aesthetics of a visualization, none suggest color use to semantically highlight structural or functional features in a consistent manner.

Color can elicit different emotional and psychological associations and reactions [222, 316] that often vary across cultures [4]. These different cultural associations are well-summarized in the visualization “Colours in Culture”¹, created by David McCandless and AlwaysWithHonor.com. This graphic shows the color black associated

¹<https://informationisbeautiful.net/visualizations/colours-in-cultures/>

with, e.g., *intelligence* for Asian cultures versus *style* for Hindu and Japanese cultures. Color can also take on different meaning in different contexts within the same culture, suggesting a more subjective and nuanced interpretation of color. For Native American cultures in the “Colours in Culture” graphic, we see black associated with both *balance* and *death*. Similarly, Wexner’s [553] study of color-mood associations with 94 psychology students at Purdue University found strong associations for the color black to *despondent*, *dejected*, *unhappy*, and *melancholy* as well as *powerful*, *strong*, and *masterful*. A ground-breaking study by Adams & Osgood [4] analyzed the affective meanings of color across 23 different cultural groups with mixed agreement. Their analysis used bipolar adjectives grouped into three factors: *Evaluation (E)*, e.g., good ↔ bad; *Potency (P)*, e.g., strong ↔ weak; and *Activity (A)*, e.g., active ↔ passive. Filmmakers frequently take advantage of affective color associations to define the mood and genre of a film. To analyze the consistency of these choices, Wei et al. [547] collected and classified multidimensional feature vectors, including movie dynamics, pace, and dominant color ratio, to determine a film’s genre and mood (from a Western perspective [316]) with approximately 80% accuracy.

Color affect is also well-researched in the visualization community. Bartram et al. [37] studied the strength of affective associations with certain color palettes. Their findings include a strong association of *calm* with cool-hued colors that have high lightness and low saturation, an association that was again found in a different study investigating affective word clouds [275]. Not all associations are so strong, and Bartram et al. note a need for more nuanced analysis of color palettes and harmony patterns. While we previously discussed the different meaning of color between and within different cultures, color affect may also be highly individual [8]. This is particularly true when associating color with concepts or lesser-known objects [456], which need sufficient context for the correct meaning to be inferred [443]. This is one of the challenges in visualizing molecules, which are often unfamiliar to the public. However, colors with strong semantic associations, e.g., yellow ↔ banana, can be leveraged with the right audience to positively impact performance in comparison-related visualization tasks [303]. Although there are numerous reasons for subjective and variable interpretations of color semantics, e.g., culture or color blindness, we can use insights from such works to tell more consistent stories through color in the visual communication of physiology.

Empirical Visualization Studies. While empirical studies are increasingly core elements of visualization research to assess communication and other tasks [87], conducting a good empirical study can be challenging [549, 580]. For example, conducting expert reviews rather than a broad user study is strongly dependent on the type of visualization as well as the stage of that visualization’s development. Tory & Möller found expert reviews particularly valuable when evaluating early prototypes [505]. Our study on the practice and preferences for visualizing physiological phenomena, described in Paper D, targeted experts from diverse domains. We took this approach in order to ensure sufficient knowledge to provide high-quality feedback from our participants on all assets and scenarios that we presented. Empirical visualization research may attempt to understand the field of visualization as a whole, e.g., visualization research keywords [219], or specific aspects, such as memorability [55, 56, 301]. Our study design

and the use of keywords is inspired primarily by the approach presented by Isenberg et al. [219].

Empirical studies on biomedical visualization are often controlled and narrow in scope, e.g., to evaluate a specific technique. Such studies tend to focus on perceptual and cognitive aspects, e.g., Baca et al.'s [32] study assessing efficacy based on usability, aesthetics, and iterability for a visualization of combustion. Comparative studies may examine traditional illustration methods, e.g., pen and ink, relative to computational renderings mimicking the traditional illustration style [221]. Such approaches may also assess different computational techniques, e.g., stylization and color adjustments to improve surgical field imagery aesthetics [48], semi-transparent structures in volume rendering [122], or perceptual comparison of aneurysm anatomy with an embedded flow visualization [34]. Our study described in Paper D takes a more qualitative approach focused on the comparison of assets produced using different biomedical illustration or visualization techniques. As in Baer et al.'s work [34], we asked participants to indicate personal preferences in their selections to help us understand, e.g., possible background biases.

(Fig. 3.1B and C, Papers B and C), followed by select challenges in the visual communication of physiology (Fig. 3.1D and E, Papers D and E).

3.1 Visualization for Physiology

3

Human physiology is multiscale in that it integrates the individual functions of molecules, cells, tissues, and organs into a whole organism [180]. This multiscale nature of physiology allows us to understand how, e.g., reactions at the molecular scale affect organ-scale events such as the beating of our hearts. An understanding of these normal, multiscale processes in the body further allows us to identify abnormal processes, e.g., irregular heartbeat patterns, and their root causes. Rapid advances in hardware and software have enabled researchers to capture data related to physiology with more sophisticated experimental and imaging paradigms to the point where it is now possible to model physiology across multiple scales. Consequently, the time has come to discuss visualization for multiscale physiology. The state-of-the-art report in Paper A offers a broad overview of the challenges and opportunities in visualization for physiology from a multiscale perspective.

Modern clinical workflows include a range of imaging protocols and tests related to physiology used to guide therapy, monitor treatment response or disease progression, and identify new biomarkers for medical research. Advanced technology and hardware capture an unprecedented diversity and volume of data through various acquisition techniques, e.g., spatial transcriptomics, as well as through models and simulations, e.g., advanced numerical simulations of blood flow. Data are often multivariate, ranging from static to time-dependent, from 2D to 3D images, and from scalar to tensor fields. The processes these data capture range spatially from nanometers to full body length and temporally from femto-/nanoseconds up to hours, months, and, in some cases, even years, as previously described in Fig. 1.1. However, these data are often specific to a relatively narrow spatio-temporal range. Establishing links between data across scales has been identified for years as a grand challenge in systems biology [375, 378], visualization [159, 385], and in a multidisciplinary 2018 Dagstuhl Seminar [7]. Linking these data across space and time through analytical models and visualization approaches necessitates multidisciplinary teams, exemplified by initiatives such as the Physiome/Virtual Physiological Human and affiliated subprojects [30, 137, 211, 498, 527] and the National Institutes of Health's Human BioMolecular Atlas Program (HuBMAP) [208].

In spite of the abundance of acquired and simulated data for physiology, a data-driven visualization approach is not always optimal or even possible. Hand-crafted biomedical illustrations are an invaluable supplement or alternative to data-driven visualization for representing physiology [167, 168, 180, 239, 427, 434]. However, these illustrations are time- and labor-intensive to create. This often precludes their use in, e.g., highly personalized data visualizations. Throughout this thesis, we highlight select illustrative works to demonstrate opportunities where illustration can augment or inspire data-driven approaches.

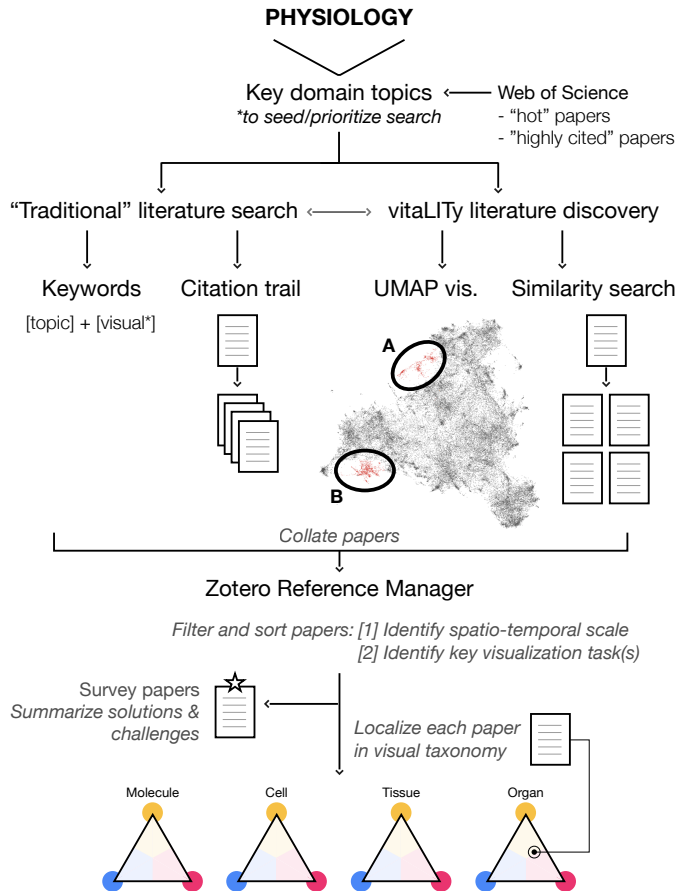


Figure 3.2: Our literature search process included both traditional search methodologies and vitaLITY [361]. VitaLITY’s UMAP visualization allowed us to identify two main groupings of physiology-related visualization literature: (A) contains molecular-scale visualization literature, while (B) contains cell, tissue, and organ-scale works (*from Paper A*).

3.1.1 Mixed Methods Literature Search and Classification

The report that opens the contributions of this thesis sketches out trends and opportunities in visualization for physiology across multiple scales, with an emphasis on human physiology. Fig. 3.2 provides an overview of our method. Our report is restricted to works discussing key domain topics that we identified from timely and highly-cited physiology research. These works furthermore are mainly application-oriented, and apply visualization techniques beyond a standard chart or figure for the topic domain. Our literature search focused both on core visualization venues as well as a selection of relevant works (approximately 46% of the total works collected) from the application domains.

Our literature search approach combined a traditional keyword-based search method with vitaLITY [361], a recent visual analysis tool for exploring academic literature.




Figure 3.3: Distribution of literature by spatio-temporal scale, excluding surveys. **Left:** Literature abundance according to scale. Each work is counted once, according to the scale to which it contributes most. **Right:** Works over spatio-temporal space, encoded by darkness. The x-axis shows temporal scale in units 10^n seconds, while the y-axis describes spatial scale in units 10^m meters (from Paper A).

vitaLITy’s UMAP visualization, shown in Fig. 3.2, allowed us to identify main groupings (A and B) of thematically-relevant papers. Works that we found using vitaLITy complemented those found with standard search methods. After collating works from our initial search we conducted a second pass that reviewed titles, abstracts, and figures to determine topical fit for our report. Our total literature set includes 305 individual works and 61 surveys or works that provide an outlook on some aspect of visualization for physiology.

Literature Taxonomy. We classified all works according to a two-level taxonomy to better identify challenges and opportunities for physiology. Our first level is inspired by the organization of physiology textbooks [180], and classifies literature into scales along a spatio-temporal axis that is roughly discretized according to biological complexity: **molecule**, **cell**, **tissue**, and **organ** (Fig. 3.3, left). Our approach bundles temporality along with size because increasingly large structures also tend to be involved in longer, more biologically-complex processes [106, 170, 452, 472]. This phenomenon is observable in the heatmap visualization in Fig. 3.3. In this figure, each work spans a segment of spatio-temporal space in increasing powers of ten, beginning with the smallest spatial and temporal scale that the input data resolve to and spanning to the size of the largest structure and time range of the process of interest discussed in the paper. Layering these ranges from each work produces a heatmap that enables rapid identification of common and less-commonly visualized processes.

A second level categorizes the literature that we collected by high-level visual task: **exploration**, **analysis**, and **communication**, which we draw from Brehmer & Munzner [66]. Depicted in Fig. 3.4, we use a triangle metaphor for each scale to array works according to their high-level task classification. Each triangle represents a single scale space, where the three triangle points represent the three respective visualization tasks. Circles indicate the position of each work according to its balance of exploration, analysis, and communication task(s). Circle darkness and size dually encode the number of works with a given task categorization that we collected in our report. This classification provides insights into the domain task requirements for each scale of




Figure 3.4: Distribution of literature according to scale and high-level task (task classification key on left). Many approaches support a combination of tasks. Darkness and size dually encode the number of works that are categorized with a given task combination within each triangle (from Paper A).

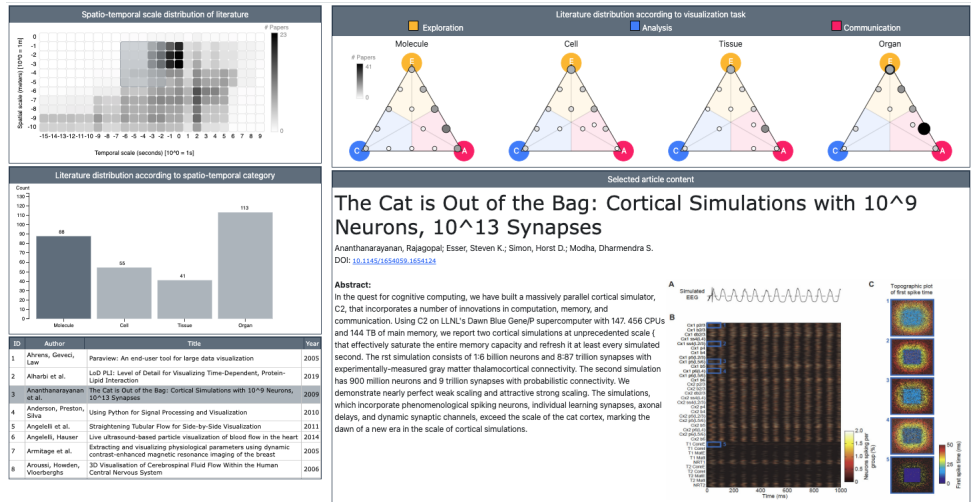


Figure 3.5: Screenshot of the STAR Literature Visual Explorer, our interface to visually navigate the literature collected and categorized for our report on visualization for physiology.

physiology and facilitates identification of ongoing visualization needs and challenges.

STAR Literature Visual Explorer. One of the artefacts produced in the course of this research is a visual interface for navigating the literature database that we created for our report on visualization for physiology, shown in Fig. 3.5. Its role is to provide a way for interested visualization and application domain researchers to easily browse the complete library of literature that we collected in our report. With this tool, we also begin to bridge the gap between exploration and communication, since this is intended in part as a means to present the various approaches in visualization for physiology in an interactive way to a diverse group of experts.

The tool consists of five panels. The top two panels directly relate to the two-level taxonomy discussed earlier: the first level categorizes literature according to organizational complexity and ranges from molecule to organ. A second level identifies any of three high-level visualization tasks within a given work: exploration, analysis, and communication. The visualization task categorization of a work is indicated by a circle glyph positioned at the appropriate location in the triangle. The circle increases in diameter and darkness with the increasing density of works with the same task allocation.



Hover facilities over a circle in a given triangle display the works that are contained in that categorization. The remaining two panels on the left side of the interface include a bar chart displaying the number of works categorized to each scale, discretized to molecule, cell, tissue, or organ. The bottom left panel displays a table that lists paper ID, authors, the title of the work, and the publication year. The large panel on the right provides details on a selected work, including an image preview and abstract. All views are linked.

By default, the database does not include the set of related surveys that the report found and rather shows only the individual methods, tools, etc. papers that we cataloged. However, the database input can be updated to include the surveys for those interested in visualizing this information.

The main entry point to using this interface is the table visualization at the lower left. Hovering or clicking on any of the rows in the table highlights that work's main spatio-temporal category (bar chart), its particular spatio-temporal distribution (heatmap), its task classification (triangle chart), and its abstract along with a representative image from the paper (article content panel). To filter works according to spatio-temporal category, the user may click on any one of the bars in the bar chart above to filter by, e.g., molecular-scale visualizations. Currently, only one category is selectable at a time.

3.1.2 Challenges and Opportunities

The following section summarizes the challenges and opportunities in visualization for physiology that we identified from our report in Paper A.

Visual Analysis in Literature Search. Our literature collection approach combined traditional search methods with new visual analysis tools from the community [361]. Although *vitaLITY* does not span the complete space of literature for our report, its visualization literature database is extensive and proved invaluable in identifying possible gaps in our literature search.

Spatio-Temporal Saturation and Opportunities. Our literature classification along a spatio-temporal axis uncovers some interesting patterns that we can identify in Fig. 3.3. The most salient grouping of works near middle top of the chart captures organ-scale processes that include brain, lung, and heart function as well as blood flow. Given that the data acquisition methods capturing these processes are generally well-established with strong clinical motivations, it is unsurprising to see an abundance of work at this scale.

A second, moderately-dark region to the mid-lower right of the chart occurs over minutes for large molecules and cellular substructures. This corresponds to works focused on visualizing data related to transcription and translation processes in gene expression. The experimental methods, e.g., proteomics and next-generation sequencing, capturing these processes are relatively coarse-grained [105] and often capture the result of these processes, which typically occur in the span of minutes in the case of a single gene [344]. The density of work we observe in this area corresponds to comparatively recent advances in experimental paradigms to measure gene expression, which is also reflected in the growth of visualization works in recent years (Fig. A.4 in Paper A) that we expect to continue in the coming years.

The light region between these two darker areas corresponding to a temporal resolution of a few seconds represents a few possibilities. The first relates to the previously mentioned coarse temporal resolution of experimental methods to capture gene expression [105]. The second relates to the timespan of human physiological processes themselves—while myriad processes bridge this time range, we found comparatively few that are confined to this time frame. Molecules may diffuse across a human cell over a few seconds [344], but this is more often of interest to researchers in the context of a larger molecular pathway or when studying whole-cell behaviors that encompass a larger temporal range.

Molecular-scale processes bridge an enormous time span, which presents both a challenge and an opportunity in visualization research to provide experts with the necessary tools to discover events of interest that are easily lost in temporal noise.

We found few cell- and tissue-scale visualization works relative to molecule- and organ-scale works, as clearly shown in the bar chart in Fig. 3.3 and in the abbreviated literature bibliography found at the end of Paper A in Sec. A.13. This is reflective of the history and trends in the availability of the source data. With increased computational power, data that truly capture living, dynamic processes at the cell- and tissue-scales have only become a reality in recent years. While visualizations of whole-cell models are now possible, this is still an extremely challenging area, and the difficulty is compounded when we increase biological complexity to modeling tissue-scale dynamics and interactions. Although multiscale models exist in some contexts, e.g., angiogenesis from the molecule- to organ-scale [411], corresponding visualization approaches are either non-existent or available only in unlinked forms with limited interaction facilities. This is an enormous opportunity and challenge for visualization research. Additionally, as part of their exploratory process many experts wish to visualize the raw data alongside any derived visual information. Research incorporating these data alongside multiple linked interactive views that furthermore incorporate tools for quantifying uncertainty is another valuable direction.

Communication Lacking. Although high-level task distribution differs slightly across all four scales in our report, communication-oriented works are consistently limited relative to exploratory or analytically-focused works, as shown in Fig. 3.4. The majority of works we came across are motivated by expert collaborators who work with specific data types and have concrete task requirements that are more often exploratory or analytical. These data are often cutting-edge, and domain experts first need support in understanding the data themselves before sharing or presenting these data. They reach out to visualization researchers to understand these data in a new or different light beyond their standard approaches. A second possibility for the lack of communication works is that the data are often inaccessible. This inaccessibility can be due to patient- and research subject-related protections. Alternatively, those datasets that are available, e.g., omics data packages, often require extensive processing efforts that hampers their usability. Finally, visual communication of complex, multifaceted physiology is simply hard, requiring a high level of understanding from the content author to distill this information into a clear message. This is an instance where drawing from other disciplines, such as biomedical illustration, can help bridge gaps in data accessibility or availability to show salient points at a higher level of abstraction. This finding

inspired our further work on visual communication that Papers D and E detail.

Application Domain Use. Many visualization approaches have yet to fully permeate the application domain. These approaches are often highly-specific techniques or algorithms, while those that are widely adopted in the application domains are usually more generally applicable, designed for ease of use, and have dedicated support teams to ensure stable production releases. Across all scales, visualization methods in the application domains are relatively restricted to direct visualization of time-lapse sequences of imaging data, as in light microscopy or medical imaging, e.g., fMRI. Depending on the data, approaches may extend to volume rendering and surface mesh creation with limited features for exploration and analysis [9, 131, 476]. Similarly across all scales, standard visualization approaches such as scatter or bar plots are common for analytical tasks such as identifying the frequency or distribution of features of interest. These are easily created in tools such as Microsoft Excel. Such visualizations often exhibit limited interactivity, e.g., brushing, and rarely link across multiple views. Ensuring greater adoption of advanced visualization techniques can occur in a variety of ways. Coordinated efforts with funding agencies can help in establishing initiatives for broader deployment and access to cutting-edge visualization techniques. Furthermore, visualization researchers may consider developing advanced approaches as plugins to established domain tools rather than stand-alone solutions. This could increase domain accessibility and mitigate the resources that would typically need to be expended for continued tool support.

3

3.2 Exploratory Analysis of Multifaceted Data

From the challenges and opportunities that we identified in our report on visualization for physiology, we now focus on two areas that stand to benefit from further visualization research. These areas fall within the context of expert users who wish to make discoveries about their data. The task combination that we target in these works is termed *exploratory analysis*, which integrates human interactions for organic exploration with computational analysis. This combination leverages the relative strengths of humans and computers that are particularly interesting in the context of visualization. Our contributions in this space go beyond a standard approach to the visualization of physiology and encompasses non-spatial visual analysis of different types of complex cohort data. The first area we examine is data-specific: magnetic resonance spectroscopy data. Our second area is broader and explores a method for exploratory analysis of multifaceted cohort data, such as data acquired from clinical routine.

3.2.1 Exploratory Analysis of MRS Data

The identification of biomarkers can lead to early diagnosis and assessment of numerous diseases and disorders of the central nervous system. For example, a brain tumor diagnosis is among the worst of all malignancies, and early detection is key to an improved prognosis. This in part has inspired clinical research into acquisition modalities such as magnetic resonance spectroscopy (MRS), an *in vivo* non-invasive technique used to estimate changes in concentrations of certain biomolecules, i.e., metabolites,

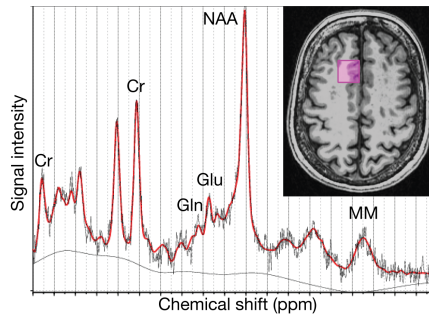


Figure 3.6: Plot of the typical visual output for spectral quantification by LCMoDel [408], with processed data in black and model fit in red. Structural localization and annotations of key metabolite peaks have been added (from Paper B).

in a tissue region [517]. However, while recent technological advances in MRS acquisition have enhanced data quality and resolution [525], the visualization of MRS data remains largely unexplored.

MRS produces a considerably different readout than magnetic resonance imaging (MRI). While MRI captures the spatial distribution of atomic nuclei at a high spatial resolution to produce a greyscale image of recognizable anatomical structures over many voxels, MRS trades spatial resolution for detailed chemical information using the same device. It acquires an abstract spectrum per single voxel, as shown in Fig. 3.6. This means that, while MRI may be used to roughly identify the extent of a tumor, MRS can support identifying the type of tumor [190]. However, translating MRS data into clinically useful biomarkers is an open challenge with limited attention from the visualization community.

MRS Data Characteristics & Task Abstraction. MRS produces a spectrum of signal intensity as a function of frequency for each measured voxel, as shown in Fig. 3.6. A chemical shift describes an intensity peak at a given resonance frequency. These chemical shifts, expressed in parts per million (ppm), arise from the different nuclear properties of the measured biochemical structures and denote **metabolites** in the acquired voxel [513]. A metabolite may consist of single or multiple intensity peaks. The most commonly measured signal comes from hydrogen atoms, i.e., proton MRS (^1H -MRS). This technique can detect metabolites in concentrations 50,000 times lower than fat or water as imaged in conventional MRI.

Most tools quantifying single-voxel spectral data produce only rudimentary visual output, such as the spectral graph in Fig. 3.6. Recognizing the metabolites that correspond to these graphs is challenging. Although it is critical to see the spectrum for quality assurance, metabolite concentrations are the most clinically relevant output from this method. These concentrations are typically depicted in a simple table in standard domain tools, which does little to advance interpretation, understanding, or comparison between acquisitions. MRS presents an opportunity to leverage visualization to identify metabolite patterns that could be indicative of early-stage diseases or disorders [489]. Ratios are the most clinically useful way to present these data, but existing approaches exclude ratios entirely or permit a only limited number for si-




Figure 3.7: Typical task flow for MRS data analysis. Users begin with data discovery (A) to review spatial voxel position, associated spectral graphs, and relevant acquisition parameters. (B) continues with data selection and filtering, where spectral voxels of interest are selected and divided into groups. Data production (C) calculates all possible ratios of selected metabolites. In (D) ratios are compared and summarized for different patients or different brain regions. Each of these steps may be revisited (*from Paper B*).

3

multaneous visualization and comparison. Our method described in Paper B crucially supports simultaneous exploration and comparison of all metabolite ratios in a study. These can be examined according to spatial location, individual, time point, and brain activity state. We developed our visual approach for exploratory analysis of MRS data over one year through weekly meetings with collaborators in radiology and medical physics.

We frame user tasks in the context of Brehmer and Munzner’s multi-level task typology [66]. These abstracted tasks form the workflow shown in Fig. 3.7. The first step, **data discovery**, supplies an overview of the data inputs and initiates a user’s line of inquiry. Users then **select** and **filter** the data to identify groups for further analysis. This step **produces** new artefacts, i.e., the set of all selected metabolite concentrations as ratios. These new artefacts are then **summarized** and their values can be **compared** to learn how metabolite concentrations vary between or within study groups or identify new sources of variation, e.g., a particular brain region.

Metabolite Visualization. The core of our method in Paper B is our novel visualization of metabolite concentrations. We do this at three levels of detail: (1) metabolite concentrations selected for a study session, (2) concentration ratios of pairwise metabolites, and (3) acquisition metadata.

On the first level, we depict the **concentrations of those metabolites** selected for an analysis session. These concentrations are the output of a pre-processing step. An analysis session may include up to two groups at a time, where each group contains either all metabolites or only a subset of metabolites that the user has selected. Since each MRS spectrum is essentially a multivariate set, where each metabolite is a variable, each metabolite in the spectrum has its own set of unique statistical information. In each group, we calculate the mean for each metabolite. We calculate median, minimum, and maximum when a group contains two or more spectra, and, for five or more spectra, we also calculate the interquartile range.

For each group, we plot metabolite concentrations along a common axis [90] in encodings that vary according to the number of spectra within that group. We use bars to encode metabolite concentration when metabolites from only one spectra are added to a




Figure 3.8: Box and bar plots encode metabolite concentrations (tier 1 visual encoding). (A) employs bars where height encodes the concentration of each metabolite for a single spectral input. (B) is used for two to four spectral inputs on an axis, where height encodes the median value and whiskers encode the minimum and maximum metabolite concentration values, respectively. A box plot (C) is employed for five or more spectral inputs on a given axis (*from Paper B*).

group, shown in Fig. 3.8A. In a group with two or more spectra, the median concentration maps to bar height, and whiskers encode the minimum and maximum metabolite concentrations, as in Fig. 3.8B. Statistics for groups with five or more spectra map to a box plot [556] as inspired by Blumenschein et al. [52], shown in Fig. 3.8C.

These metabolite concentrations are the inputs to the second visualization level, which is the set of all combinations of **pairwise metabolite concentrations, presented as a ratio**. These ratios are shown in aggregate and filtered according to key attributes of interest. We calculate aggregate ratios by dividing a metabolite's mean concentration in one group into the mean concentration of each metabolite in the opposite axis group. We show these ratios using a heatmap metaphor calculated from the metabolite groups along the input x- and y-axes, as shown in Fig. 3.9. Each cell shows the aggregate ratio of the metabolite on the x-axis position to the corresponding y-axis metabolite. The ratio value maps to a diverging red-blue colormap [185]. This strategy is intended for similarity assessment, where subtle differences in ratios are important to track, as well as for large difference assessments such as those seen in tumor studies, or in research cases, where echo time is varied. Red indicates a higher x-axis metabolite input while blue indicates a higher y-axis metabolite input. Equivalent inputs map to white. If an input value is 0, i.e., the metabolite was not detected in the acquisition, we map the cell color to dark grey. We initially thought to exclude such values from the visualization, but on further discussion with our collaborators felt these were useful to preserve context.

Metabolite ratios can be filtered according to **four key data attributes**: (1) spatial region, (2) individual, (3) time point, and (4) brain activity state. The spatial region indicates the voxel sample location. Individual refers to a given subject in the study. Time denotes the multiple recorded samples within a single session, as in a time-resolved MRS study, or the number of acquisitions performed on an individual over a study period, as in a longitudinal study. Finally, brain activity state denotes if the subject was in an *active* (task-explicit) state or *resting* (task-negative) state during the acquisition. For each pairwise metabolite ratio we determine ratios for each of these key attributes, when available, in a nested fashion that mirrors the preferred order of user analysis: the ratio for each spatial region (using the average of all individuals for this region), each individual (using the average of all states for the given individual in a given region),

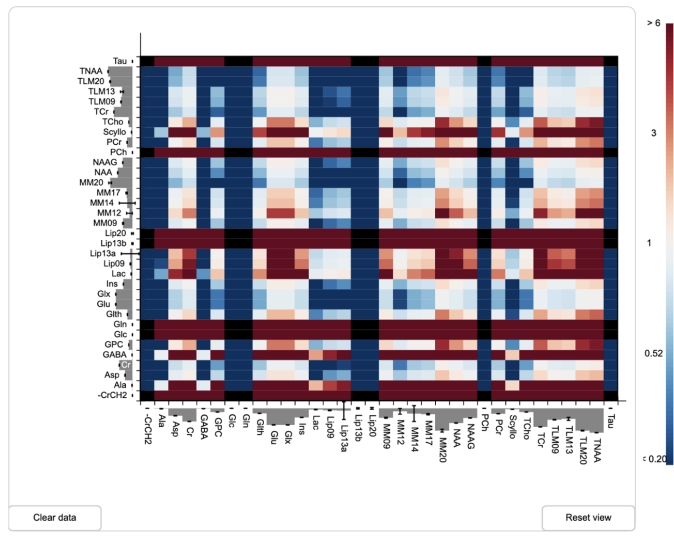


Figure 3.9: A heatmap matrix visualizes the ratio of metabolites in the x-axis voxel group against the metabolites in the y-axis voxel group. Ratios with higher x-axis metabolites are red, while higher y-axis metabolites are blue.

each state (using the average of all time points for a given state of an individual from a given region) and each time point, as shown in Fig. 3.10.

Because a majority of the studies conducted by our collaborators are small, proof-of-concept studies, we use glyphs inspired by the ATOM grammar [390] to represent each of the four attributes with limited concern for issues with display and perceptual scalability. Inspired by Ward et al. [542, 543], we designed a system of nested glyphs that could satisfy the 16 possible use cases that our method must cover. This nesting approach maintains context during analysis and mirrors our collaborators' preferred order of analysis. The left columns of Fig. 3.11 and 3.12 depict each of these use cases, which are created from a synthetic dataset containing up to 3 spatial regions, 20 individuals, 3 time points, and both active and resting brain states (360 voxels in total). Our glyph design is summarized as follows. We map spatial region to a rectangular glyph with rounded corners, which distinguishes the attribute from the sharp-cornered heatmap cell. We create one rectangular glyph for each spatial region in a study, e.g., three regions for one individual in Fig. 3.11, case 5. Individuals are depicted as circles which nest inside their relevant spatial glyph, as shown in Fig. 3.11, case 9. When time series data are present, these circles expand to rounded squares to provide more room for the time points that we visualize as spark lines [341]. An example of this is shown in Fig. 3.11, case 3 for a single individual with three time acquisitions. In studies with a single individual over multiple spatial regions, the spark line nests instead inside the rectangular spatial glyph, as in Fig. 3.11, case 7. In studies where different brain activity states are measured, we divide the glyph for the individual in half horizontally, as in Fig. 3.11, case 2 for a single individual sampled from one region. If there is a single individual sampled in multiple regions, we instead divide the rectangular glyph, as in Fig. 3.11, case 6.




Figure 3.10: Nested metabolite ratio calculations for spatial region, individual, brain state, and time (from Paper B).

We conducted a preliminary evaluation to validate our glyph design with three experts for all 16 basic use cases. For each case, we assigned a basic task and asked experts to rank the interpretability of the design on a scale of 1–10 [558] for their ability to complete the stated task. The right columns of Fig. 3.11 and 3.12 describe each of these tasks and the experts' rankings. While brief, this feedback ensured that our approach is usable for the study sizes we expect to use in our method and application.

At the third level, we present **metabolite metadata** important for context and selection that are unnecessary to encode explicitly as glyphs in the visualization: gender, age, and acquisition settings, e.g., echo time, can have varying impacts on the resulting concentrations and ratios of metabolites [475, 568]. We depict this information in a table below the heatmap (Fig. 3.13G).

SpectraMosaic Application. The result of our design study (Paper B) is *SpectraMosaic*, a novel application for the visual exploration of magnetic resonance spectroscopy data in ratio form. *SpectraMosaic* consists of two main views, as shown in Fig. 3.13: the left supports spectral inspection and customized group creation while the right serves as the spectral analysis interface in the form of a heatmap matrix. The general workflow of this application is as follows: after an offline processing step, data are loaded into the web tool (Fig. 3.13A). Data of interest for analysis can be explored, selected, and added (Fig. 3.13B–D) to a spectral ratio heatmap for deeper inquiry and hypothesis verification (Fig. 3.13E). A legend provides information on the encodings used in the tool (Fig. 3.13F). A table below the heatmap summarizes salient acquisition information (Fig. 3.13G).

SpectraMosaic Case Study. The following describes an exploratory workflow with *SpectraMosaic* for a *Giardia* single-voxel study, where we walk through an expert user's exploration of possible neuroinflammation in three different brain locations compared across two different patients.

Following data preprocessing and quantification in an offline routine, the expert loads study data into the application using the drag and drop window feature.

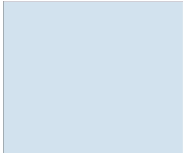
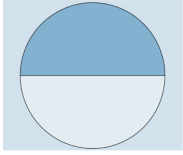
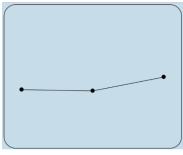
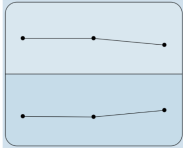
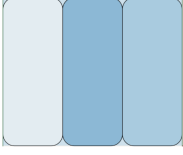
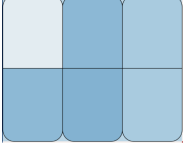
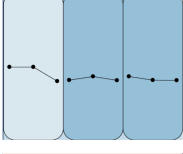

Case No.	Example	No. Voxels	Task	Interpretability
1		1	What is the ratio of Ins/TCho for patient X in the precuneus region?	A: 10 B: 10 C: 10
2		2	How does the ratio of Ins/TCho differ for patient X when acquired on the same date in a resting state versus in an active, i.e., task-positive, state in the precuneus region?	A: 10 B: 10 C: 8
3		3	How does the ratio of Ins/TCho compare for patient X acquired in the precuneus region, acquired pre-op, post-op, and in a long-term follow-up?	A: 10 B: 10 C: 9
4		6	How does the ratio of Ins/TCho compare for patient X acquired in the precuneus region, acquired pre-op, post-op, and in a long-term follow-up? How does this ratio differ between active and resting brain state acquisitions?	A: 9 B: 9 C: 8
5		3	How does the ratio of Ins/TCho compare for patient X between the precuneus, prefrontal, and hippocampal regions?	A: 10 B: 10 C: 10
6		6	In patient X, how does the ratio of Ins/TCho compare between active and resting states for each of 3 regions: precuneus, prefrontal, and hippocampal regions?	A: 8 B: 10 C: 8
7		9	What time acquisitions contributed the most to the difference in Ins/TCho seen in the leftmost spatial region (hippocampal) relative to the other two regions (prefrontal and precuneus) in patient X?	A: 9 B: 10 C: 9
8		18	How does Ins/TCho vary between active and resting brain states between spatial regions (hippocampal, precuneus, prefrontal) in patient X? Are there outliers in the time acquisition that contribute to these differences?	A: 8 B: 8 C: 8

Figure 3.11: SpectraMosaic Detail Case Scenarios 1–8

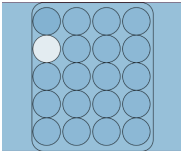
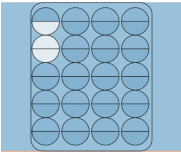
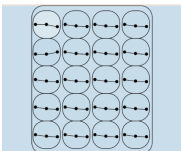
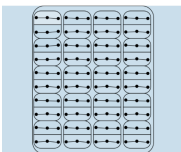
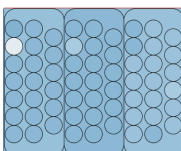
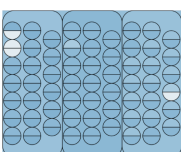
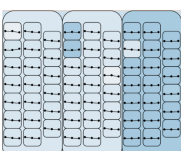
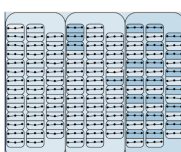
Case No.	Example	No. Voxels	Task	Interpretability
9		20	How does Ins/TCho vary between 20 patients in the precuneus region?	A: 10 B: 10 C: 10
10		40	What patients show a marked difference between active and resting state Ins/TCho in the precuneus region?	A: 10 B: 10 C: 10
11		60	Do patients show variation in time acquisitions in the precuneus spatial region?	A: 8 B: 6 C: 8
12		120	Do any patients show a difference in their active and resting states over the three different time acquisitions in the precuneus region?	A: 10 B: 10 C: 9
13		60	How does Ins/TCho vary between 20 patients in each of the three spatial regions acquired (hippocampal, prefrontal, precuneus)?	A: 10 B: 10 C: 10
14		120	What patients show a marked difference between active and resting state Ins/TCho in the three spatial regions? Which state seems to vary the most between regions?	A: 10 B: 10 C: 10
15		180	In which brain region(s) is there the most variation in Ins/TCho over time? In higher variation patients is it just one time point, or all, that contribute to the outlying ratio?	A: 10 B: 10 C: 9
16		360	In which spatial region do we see patients with the most variation between active and resting Ins/TCho values over time?	A: 10 B: 10 C: 10

Figure 3.12: SpectraMosaic Detail Case Scenarios 9–16

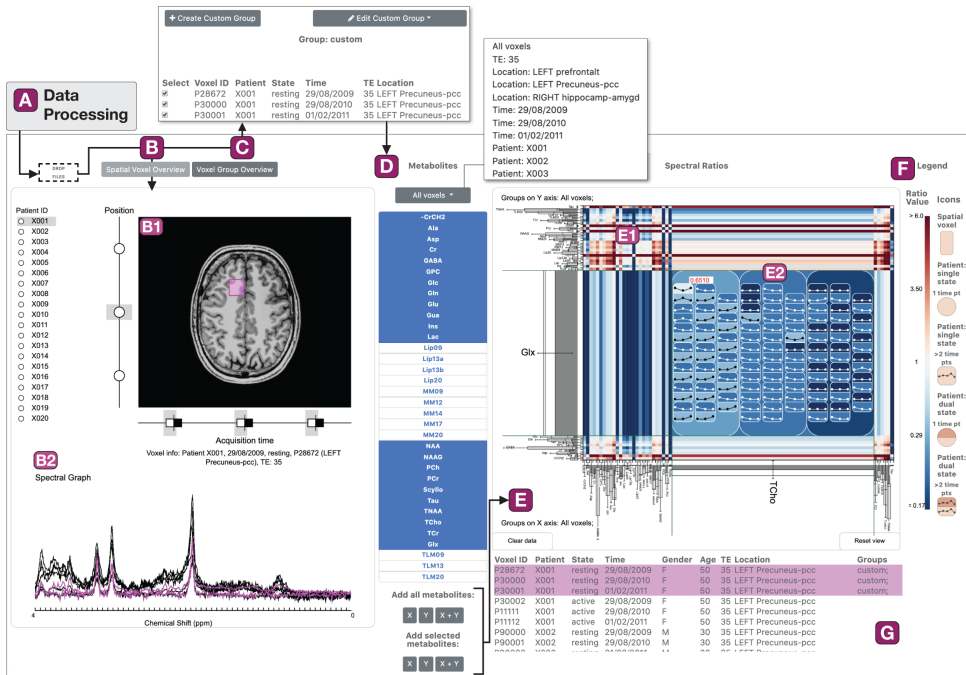


Figure 3.13: SpectraMosaic application overview (from Paper B).

These data remain linked semantically in the application to ensure continuity between voxel location, its spectral output, and patient-specific information.

The expert next can inspect the input data. In the patient selector panel on the left side of the interface, radial buttons allow the display of one patient with associated acquisition data at a time. The vertical axis to the left of the structural data image shows the set of voxel regions acquired for a selected patient. The horizontal axis below the image indicates the number of different time acquisitions performed for this voxel. In this study the expert is interested in comparing all spatial voxels of each patient for only one time acquisition. Hovering over or selecting each voxel node calls the spectral graph for the voxel below the image (Fig. 3.13B2), as well as additional information such as voxel ID, and echo time used in the acquisition protocol (Fig. 3.13G).

To compare metabolites, the expert can either select an automatically-created group immediately from the metabolite selector panel, (Fig. 3.13D) or can first create a more customized group for analysis using the voxel group overview panel (Fig. 3.13C). This view enables creation of a custom group whose membership can be updated as desired throughout an analysis session. Group membership is updated automatically in the bottom right table (Fig. 3.13G).

Once a group of metabolites have been added to each axis, a heatmap matrix is drawn onscreen (Fig. 3.13E1). The matrix size is determined by the number of metabolites added as inputs. Users can select a subset of metabolites to review, or all metabolites at once (Fig. 3.13E). The visualization updates as the inputs are changed. The height of the metabolite bar encodes, in the case of these two groups, the median concentration, with whiskers extending to indicate the input value range (Fig. 3.14).

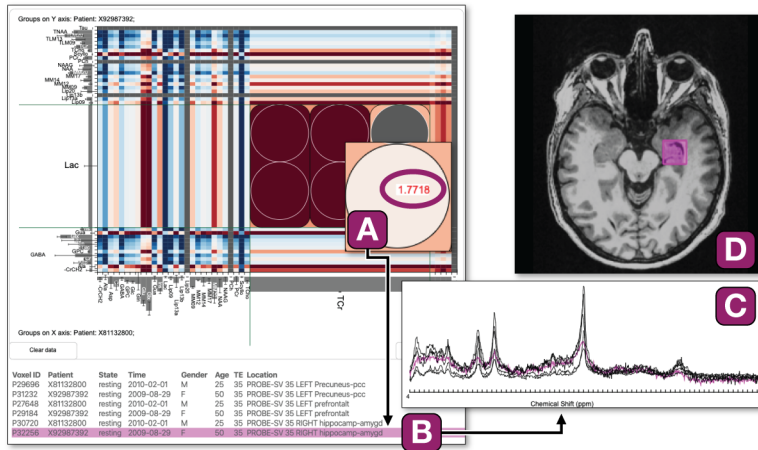


Figure 3.14: Following overview of the metabolite ratios between the individual patients, the expert inspects Lactate (Lac)/Total creatine (TCr) ratio between two patients for all three regions at TE 35 ms. They note a high lactate measurement for the female patient relative to the other measured regions (A). Subsequent inspection of spectral metadata (B), the spectral graph (C), and the brain region in which this measurement was acquired (D) help establish reasons for this difference (*from Paper B*).

The expert is then interested in seeing detailed metabolite ratio information for a given cell. They inspect individual ratio variation for *Lactate (Lac)* against *Total Creatine (TCr)*, since *Lac* is often elevated in giardia infection (Fig. 3.14A). On expansion of a cell, hovering over glyph elements highlights linked patient spectra, voxel locations, and voxel metadata in the other views to provide additional information on the metabolite ratio in question (Fig. 3.14B-D).

Through the similar hues of the color map, both patients are quite similar between the prefrontal and precuneus voxels, but the hippocampal region shows a large difference, indicating that *Lac* and *TCr* have strongly differing levels in this area. This arises from the female patient, since the male patient had no recorded *Lac* as indicated by the grey color of the circle glyph. The whiskers on the bars also confirm that the large input range for *Lac* causes this variation. However, the red tooltip indicates that the model fit for this metabolite to the raw spectrum may not be reliable, and needs further investigation. Subtleties like this would be difficult to understand through visualization of the spectral graph or through statistical measures alone. The expert can reset the heatmap matrix view to its initial state at the bottom right of the interface, or clear it for a new line of inquiry with new group inputs by clicking the button at the bottom left of the view. This enables an iterative, exploratory analysis workflow for MRS data.

From *SpectraMosaic* we move beyond the focus on a particular type of data to look more broadly at ways to enable users to explore and visualize interesting relationships in multifaceted data.

Algorithm 1: *Dimensional bundle* creation for two or more dimensions

```

1 initialize pool = all dimensions in dataset
2 do
3   mark all dims in pool as possibly contributing
4   initialize new bundle
5   perform FAMD on pool
6   for all dimensions in pool
7     if PC1 loading  $\geq$  contribution threshold
8       move dimension from pool to new bundle
9     else
10      mark dimension as non-contributing
11 while pool contains dimensions marked as non-contributing
12 for all bundles
13   perform FAMD on bundle
14   store PC1 and PC2 for bundle

```

3

3.2.2 Exploratory Analysis of High Dimensional Clinical Cohort Data

Dimensionality reduction techniques can reduce the complexity of high-dimensional data through projection to a lower-dimensional space. However, these techniques can emphasize strong, uninteresting patterns in the data and hide variations when used alone and monolithically.

Clinical efforts investigating biomarkers for early disease diagnosis require extensive analysis of multifaceted clinical cohort data. Analysis of these data can be overwhelming. Many items are missing, and the data are both heterogeneous and high-dimensional. While visual analysis approaches that incorporate dimensionality reduction methods can make analysis more manageable, the risk of such methods is that subtle and interesting patterns may be lost in the process. The *DimLift* approach that we summarize here and in detail in Paper C uniquely handles complex, multifaceted data in a way that preserves these subtleties in the data for subsequent exploration. We do so by combining an iterative dimensionality reduction method with user interactions to edit and inspect groupings of dimensions, so-called *dimensional bundles*, which contribute similarly to the overall variance of the dataset. This avoids a monolithic treatment and instead produces hierarchical bundles of dimensions that retain the expressivity of the original dataset. Demonstrating with the simple synthetic dataset in Fig. 3.15, smoking (b3) is an important clinical indicator of cardiac risk (o2). However, a standard dimensionality reduction process may not easily show this relationship. It instead buries these dimensions in all five principal components (Fig. 3.15, left). In contrast, our approach (Fig. 3.15, right) extracts subsets of similarly-contributing dimensions. For example, our approach places the lifestyle-related dimensions education level (b1), workout frequency (b2), and smoking (b3), together with the similarly-contributing variable cardiac risk (o2). We furthermore visually retain data provenance to allow clinical researchers to identify original dimensions for further analysis. Our approach is intended as an early exploratory analysis step prior to a more detailed investigation in SPSS or similar.

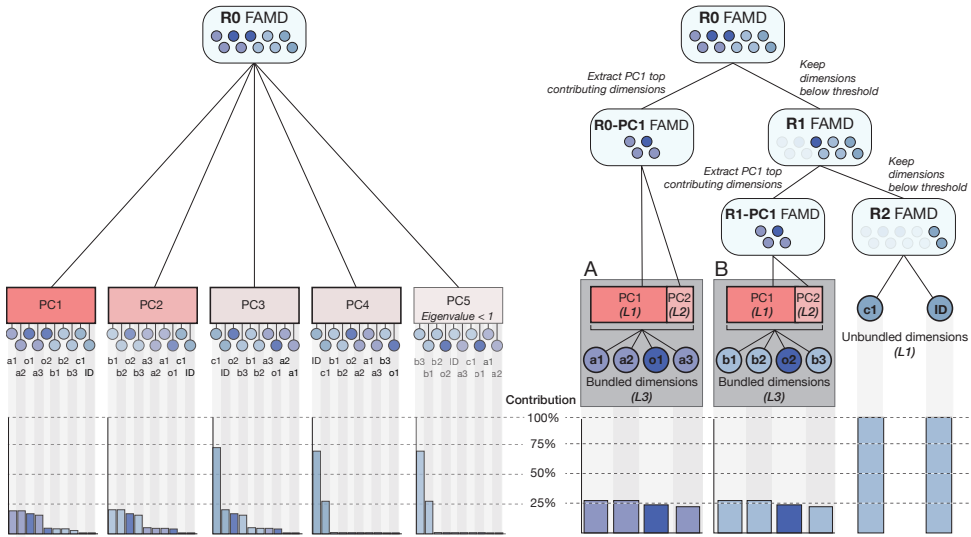


Figure 3.15: We contrast our iterative algorithmic approach (right) with a standard approach (left) using a synthetic ten dimensional health and lifestyle dataset comprised of four quantitative [height (a1), weight (a2), waist circumference (a3), BMI (o1)] and six qualitative [education level (b1), workout frequency (b2), smoking (b3), gender (c1), eye color (c2), and cardiac risk (o2)] dimensions. A standard approach contains all ten dimensions in each principal component (PC), e.g., cardiac risk (o2) is present in all PCs. In contrast, our approach produces a pair of *dimensional bundles* (**A**: body measurement, **B**: lifestyle) containing only dimensions with similar variance contributions, where cardiac risk is bundled into **B**. Dissimilarly contributing dimensions, i.e., c1 and c2, remain unbundled (*from Paper C*).

Create Dimensional Bundles. In the following, we briefly describe our approach to creating dimensional bundles, which is also described in pseudocode in Algorithm 1.

To create *dimensional bundles*, we use factor analysis of mixed data (FAMD) [289], a dimensionality reduction technique that is applicable to complex, mixed-type data, to extract and bundle dimensions that meet a contribution threshold in the first principal component [394] of each algorithm run from a main pool of dimensions (line 7 of Algorithm 1; Fig. 3.15, right). These extracted dimensions are the inputs to a *dimensional bundle*. We run the FAMD algorithm again on these extracted dimensions to create the bundle components (line 12). In the bundle components, we retain the first and second principal components (PC1 and PC2) to describe the overall variance and to give context to the variation that the bundle captures. We additionally include the original contributing dimensions in the bundle. We continue iterating and extracting contributing dimensions from the main pool until the contributing variance for the possible bundle falls below the variance contributed by any of the dimensions alone [575]. These dimensions remain unbundled.

Our method can furthermore handle data with missing items and mixed types, e.g., continuous and categorical, as are common in clinical routine data. We impute missing data through one of four options (mean, hot-deck, multiple imputation of chained equations, or principal components), with MICE [554] as the default after discussions with

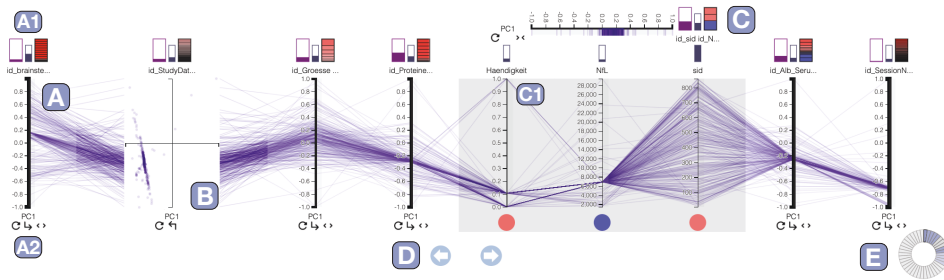


Figure 3.16: Parallel coordinate axes (A) map to the first or second principal component (PC1, PC2) of a *dimensional bundle*. Glyphs (A1) provide feedback on variance contribution, missingness, and composition. View interactions (A2) allow users to pan (D) through the dataset, swap axes between PC1 and PC2, drill-down into a PC1 vs. PC2 score plot (B), or drill-down further to the *dimensional bundle* component dimensions (C) and their relationships (C1). A chart at the bottom right (E) provides an overview of all *dimensional bundles* and unbundled dimensions, a subset of which are visualized as plot axes (from Paper C).

3

our collaborators and a short study on the results of the different imputation approaches. We briefly discuss the effects of the different imputation methods in the context of a cerebral small vessel disease dataset at the end of this section.

Explore *Dimensional Bundles*. Projecting to a lower-dimensional subspace in dimensionality reduction often creates a disconnect from the source data [235]. We solve this issue by preserving and mapping the semantics of the *dimensional bundles* directly to visual elements. We echo the natural hierarchy produced through data aggregation in a visual aggregation step with a parallel coordinates plot as the basic metaphor, with facilities for dynamic user interaction to explore and recompose *dimensional bundles* to pursue a line of inquiry.

Parallel coordinates are a common technique to visualize relationships and correlations in multifaceted datasets [189]. Our modified parallel coordinates plot supports three layers of visual analysis within and between each *dimensional bundle*, as displayed in Fig. 3.16. Each parallel coordinate axis represents a *dimensional bundle* (Fig. 3.16A) or an unbundled dimension (Fig. 3.16C1). The stroke widths of the axes are set according to the number of dimensions they contain [18]. By default, a bundle axis represents the first principal component of a *dimensional bundle*. Items are plotted along the axes according to their scores. A second layer expands a secondary axis from the main axis to show a scatterplot of PC1 versus PC2 item scores (Fig. 3.16B), which is inspired by earlier similar approaches [187, 207, 226]. An innermost third level consists of all dimensions contributing to the *dimensional bundle* (Fig. 3.16C) and plots the original item values. This representation is inspired by Andrews et al. [18]. Rectangular filled glyphs, positioned above each axis, provide details on bundle composition and variance contribution (Fig. 3.16A1), including the eigenvalue and explained variance, amount of non-imputed, data, and the contributing dimensions with their respective loadings. In the lower right, a donut glyph (Fig. 3.16E) provides an overview of the total number of dimensions and bundles that have been created. It shows the bundles in purple and unbundled dimensions in grey.




Figure 3.17: *DimLift* view interactions allow iterative exploration of *dimensional bundles* (from *Paper C*).

Navigation facilities enable experts to explore the contents of *dimensional bundles*, as illustrated in Fig. 3.17. These controls are mainly accessible through the options below each bundle (Fig. 3.16A2). **Panning** through the parallel coordinates plot allows the user to explore correlations between all bundles. **Brushing** over a bundle axis selects a **subset** of items, which is adjustable in an adjacent panel. **Drilling down** to a plot of PC1 vs. PC2 item scores shows item spread or clustering within a bundle. **Drilling** further to the contributing bundle dimensions allows assessment of correlations within a bundle. **Swapping** the bundle axis from the first (PC1) to the second (PC2) principal component provides a different perspective on the variance captured in the bundle.

Modify Dimensional Bundles. Automatically-created *dimensional bundles* may not always be conducive to a user's analysis goals, or the user may find that the bundles still bury features that they find interesting. Structural interactions allow for recomposition of *dimensional bundles* on-the-fly. Bundles can be created from scratch based on inter-




Figure 3.18: *DimLift* structural interactions allow for the creation or modification of *dimensional bundles*. Using our synthetic health dataset we **create** a new bundle combining smoking (b3) with gender (c1); a resulting eigenvalue above 1 shows a fair grouping with equal dimension contributions. The left diagram provides a conceptual overview of this process (from Paper C).

3




Figure 3.19: (a) In the analysis of a synthetic health dataset we may suspect gender (c1), to have interesting correlations with the lifestyle-related bundle, i.e., education level (b1), workout frequency (b2), cardiac risk (o2), and smoking (b3). We **add** gender (c1) to this bundle and observe that gender shows no contribution (loading = 0) to the bundle variance. (b) We inspect a *dimensional bundle* comprised of lifestyle dimensions, e.g., education (b1), workout frequency (b2), smoking (b3), cardiac risk (o2), and gender (c1). We suspect a correlation between cardiac risk (o2) and gender (c1), so then *lift* these dimensions to better target and test our hypothesis by removing all other dimensions from this bundle. With an eigenvalue above 1 and changes in contributions/loadings indicated by hue at the bottom of the axes, we note a **subtle correlation** that was previously undetectable. This is conceptually illustrated on the left (from Paper C).

esting groupings by descriptive statistics from an accompanying dimension overview scatter plot [351], as shown in Fig. 3.18, or by searching directly for dimensions of interest to combine into or remove from a bundle, as in Fig. 3.19.

During the analysis a user may wish to visualize the degree that a group of dimensions that relate conceptually, e.g., all lifestyle input variables in our synthetic health dataset, are correlated. Similarly, dimensions that do not seem to be conceptually-related may exhibit similar descriptive statistics, e.g., similar mode or diversity measures, that would be interesting to apply dimensionality reduction to for detailed correlation assessment. Fig. 3.18 demonstrates the workflow for creating a new *dimensional bundle* based on similar descriptive statistics. Rather than creating a new bundle, the user may also modify an existing bundle and either add (Fig. 3.19A) or remove dimensions in place (Fig. 3.19B). After running the FAMD on the selected set, the bundle is created or updated in the parallel coordinates plot for subsequent exploration and

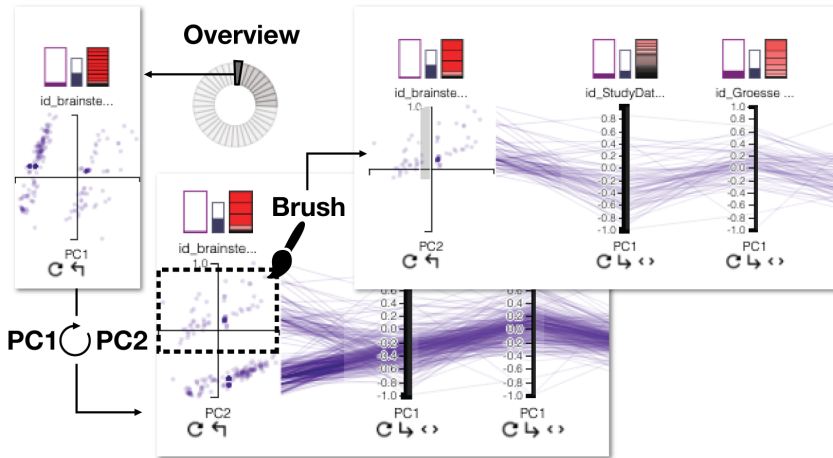


Figure 3.20: In exploring a clinical cohort dataset for cerebral small vessel disease (CSVD) in the *DimLift* application, experts select a bundle of primarily imaging data for closer examination and drill-down to observe two distinct clusters. Swapping the axis to PC2 allows subset creation of the top cluster; this corresponds to selection of non-imputed items within the bundle. Addition of APOE-related dimensions to the bundle allows for correlation assessment of these interesting dimensions within a single bundle (*from Paper C*).

possible further modification. These structural modification tools empower the user to flexibly reconstruct the dimensional hierarchy to suit their discovery goals. With feedback on their suitability in the parallel coordinates plot, users may quickly gain insights into their data by lifting interesting dimensions from their original bundles.

***DimLift* Application & Case Study.** We next briefly demonstrate our *Dimlift* application workflow for the exploratory analysis of a clinical cohort dataset for cerebral small vessel disease (CSVD), which was part of a two-expert case study including our clinical collaborators. Responsible for one in five strokes worldwide, cerebral small vessel disease is the most common cause of vascular cognitive impairment in the elderly, and the identification of new biomarkers for this disease is of key interest to neuroscience researchers. The study consists of data from 307 patients collected from clinical routine. The data are mixed, consisting of 168 dimensions containing laboratory, education, demographic, and lifestyle information. Twenty-four additional dimensions describe the volume of 24 brain structures, e.g., caudate and hippocampus, as derived from T1-weighted magnetic resonance imaging data. As is common with this type of data, approximately 76% of entries are missing because of, e.g., missed appointments or the fact that not all patients require the same tests.

On an initial overview, experts discover two notably interesting bundles generated. One bundle includes lacunes and microbleeds together, while another combines nicotine, intracerebral hemorrhage, and enlarged perivascular spaces. Undetectable in their standard analysis approach, this generates a new hypothesis regarding the correlation of these dimensions that requires further analysis.

For closer inspection of a bundle containing primarily imaging data, experts can swap the axis and then drill down into this bundle to note two distinctive clusters in

the principal component score plot, as displayed in Fig 3.20. Experts further note in PC2, which is now the dominant axis in the parallel coordinates plot, that the Boston/STRIVE criteria dimension is shown as one of the more influential dimensions in this bundle. This is an interesting find, since Boston/STRIVE criteria taken together are known biomarkers for CSVD.

Experts add APOE, another known biomarker for CSVD, to this bundle to explore possible correlations between these biomarkers. Swapping the axes allows for rapid subset selection of the top group. Subsequent exploration of correlations within this subset shows non-imputed items for Boston/STRIVE criteria, sex, and group, a consistent range for white matter and CSF volume values, and consistent APOE values that merit further investigation.

3

User-Adjusted Imputation in Cerebral Small Vessel Disease Case Study. To better understand the effects of different imputation methods on our *DimLift* approach, we additionally conducted a small-scale study of four different imputation methods in the context of the cerebral small vessel disease study just presented. Our goal for this small study was not to identify the most accurate or best method of imputation. This determination is not possible without a ground truth, and for the clinical study we tested this was unavailable. Our goal instead was to establish the degree to which our bundling approach is preserved in spite of different imputation methods, and to explore the semantic relevance of differences that did occur in the bundling process.

We narrowed down to the following four imputation methods after reviewing the state-of-the-art in data imputation approaches [108, 438] and discussions with our clinical collaborators. Each of these imputation methods are implemented and available for use in our application:

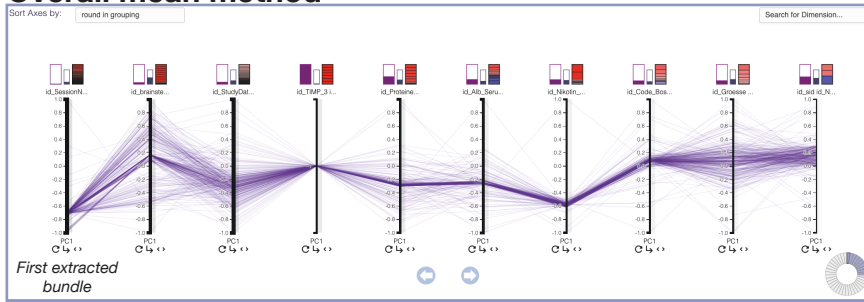
- Overall mean value for quantitative data/“not defined” for qualitative data [387, 438]
- Hot-deck imputation [19]
- Multiple imputation of chained equations (MICE) [554]
- Principal components imputation [28]

Although we initially investigated cold-deck imputation [19] as well, this was not possible with our clinical study since there was no suitable dataset to which we had access to for comparison.


We show the top-level bundle results from each imputation method in Fig. 3.21. Bundles are sorted by their order of extraction from the main pool of dimensions. We can see that each method produces 10-11 bundles. The hot deck and MICE methods produced ten bundles, while overall mean and principal components methods produced 11 bundles. The bundles generally maintain the broad semantic themes of tests, lifestyle, and measurements. The main difference between the four imputation methods is in the granularity of their bundling.

For example, the first extracted bundle is generally consistent between imputation methods. It consists mostly of lacune and microbleed measurement dimensions from different regions of the brain. This bundle shows only 17% complete data, demonstrating a solid use case for the robustness of our bundling with different imputation methods. As an example of the difference in bundling outcomes with the different imputation methods, we observed that the dimension classifying smoker/nonsmoker is

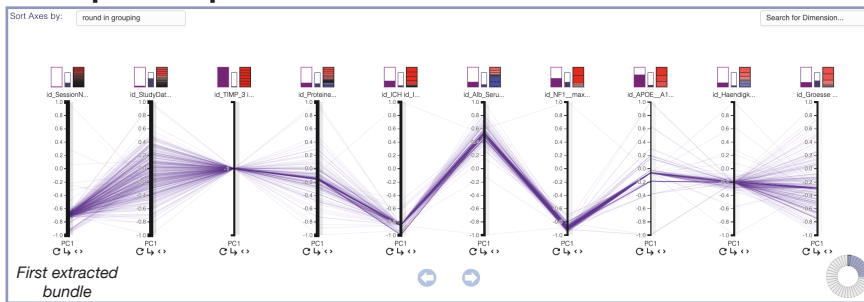
Overall mean method



Hot deck method



Principal components method



MICE method




Figure 3.21: Visual output of top-level *dimensional bundles* produced with each of the four imputation methods tested: overall mean, hot deck, principal components, and MICE.

bundled differently in each method. However, after discussions with our collaborators about this, each of these bundles still make semantic sense. This dimension has myriad effects on other dimensions, and in each bundle its relations to its companion dimensions present interesting lines of further inquiry. For example, the MICE and principal components methods bundle smoking with alcohol. However, hot deck imputation leads to a bundling of smoking and alcohol along with education level, gender, and blood pressure or cholesterol measurements. In overall mean imputation, smoking is bundled with enlarged perivascular spaces (EPVS), which presents an intriguing connection that our clinical collaborators noted for further inquiry.

This brief qualitative assessment of the effects of different imputation methods found that our approach generally preserves bundle patterns. Nuanced differences are apparent in each imputation method. However, these still make semantic sense on inspection. No imputation method is ever entirely ideal. Its utility is highly dependent both on the specifics of the data and the goals of the analyst. However, our findings indicate that our *DimLift* approach may flexibly accommodate the results of different imputation methods with a degree of robustness that presents an interesting area for further research.

3

3.3 Visual Communication of Physiology

While expert exploration and analysis of physiological data is crucial to establish understanding, communication of these data both between experts and to non-experts, e.g., the general public, is imperative to the scientific process. However, visualization research efforts tend to focus more on exploratory and analytical tasks, and communication tasks are often fulfilled by other disciplines such as biomedical illustration. Motivated by the lack of communication-oriented works in our report in Paper A, we conducted two studies to understand the interplay of visualization and biomedical illustration techniques in visual communication for physiology. The first was a two-part exploratory study that we detail in Paper D. The first component of this study examined the practice of creating visualizations to communicate physiology. The second component explored the criteria that audiences use when evaluating a physiological process visualization targeted for communication. This study introduced several interesting questions, one of which we followed up on in a short study and call to action on the value of establishing guidelines for semantically-meaningful color palettes in molecular visualizations. This is described in detail in Paper E.

3.3.1 Practice & Preferences for Experts and Non-Experts

Designed over a series of expert interviews and focus groups, our interdisciplinary study detailed in Paper D focused on common communication scenarios of five well-known physiological processes and their standard visual representations. Feeding into these scenarios was a set of visual assets we developed that captured the range of standard practices by biomedical illustrators and visualization researchers when crafting a visualization of a particular topic, e.g., blood flow. We framed these scenarios in a survey with participant expertise spanning from minimal to expert knowledge of a given topic. We summarize our process in Fig. 3.22. Our goal in this study was to gain

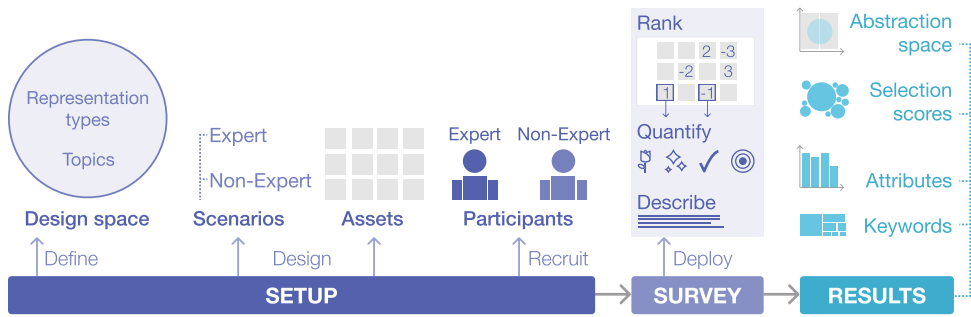


Figure 3.22: Three-phase study pipeline. **Setup:** define the design space, create audience scenarios and visual assets, and recruit survey participants, **Survey:** deploy survey asking participants to rank, quantify, and describe their top and bottom asset selections for each scenario, and **Results:** review survey results for patterns in selection abstraction space, scores, attribute rankings, and frequent keywords (from Paper D).

insights into how visualization and biomedical illustration techniques are used and assessed by differing audiences for visual communication, and to identify opportunities for further growth and convergence of techniques.

Setup. The first part of our study setup involved defining and constraining the design space according to (1) representation types and (2) physiological topics. We relied on discussions with practicing visualization researchers, biomedical illustrators, reviewed online biomedical illustration galleries and resources from the Association of Medical Illustrators¹ to establish the constraints in these respective areas. We limited the space of representation types to only color media that are either still or animated, and that are typical representations for a given topic. We excluded interaction or annotations, with the exception of arrows in a few instances, to maintain a reasonable scope. Our five selected topics span the micro- to macroscale and include patho- and physiological processes that are frequently depicted in visualization and biomedical illustration. These serve as a proxy for the large space of representations of physiological processes. Our microscale topics include (1) signal transduction, a normal process whereby a signal is relayed along a chain of molecules in the body, and (2) constitutive activation, a process where one or more molecules in a signal chain is always switched “on” to create a looping signal relay. At the mesoscale, we selected (3) normal blood flow and (4) an aneurysm. Our macroscale topic is (5) tumor metastasis. This focuses on the movement of tumors from their origin site to other organs. This is in essence a multiscale topic as well, since metastasis can occur as a result of constitutive activation at the microscale and tumors are carried through the bloodstream.

We designed communication scenarios and assets for each topic with input from expert focus groups, following in part the framework for creative visualization-opportunities workshops described by Kerzner et al. [249]. This helps control the design space and provides important in-depth insights into discipline-dependent visualization practices. Each group consisted of three to four biomedical illustrators and/or visualization scientists. Providing a scenario is an effective means to capture specific

¹<https://ami.org/>

goals and research questions in a domain [278]. These scenarios drive comparison and selection of assets in our survey. For each topic, we created an expert and a non-expert scenario, which are detailed in Tables 3.1 and 3.2. Our asset production pipeline included the Adobe Suite (Illustrator, Photoshop, AfterEffects) [5], Blender [51], ParaView [9], and 3D Slicer [250]. We produced animated assets as short, looping GIFs. Our full asset library is available at https://github.com/lauragarrison87/Biomedical_Process_Vis/tree/main/2-assets. The same set of assets are used for both expert and non-expert scenarios for a given topic.

Table 3.1: **Expert Audience Survey Scenarios**

Topic	Scenario
Signal Transduction	An immunology researcher is publishing in an immunological venue on the newly-discovered pivotal role that a ligand plays in a signaling pathway. Their goal is to communicate the specificity of the activation pathway and its location in the cell with a visual supplement to their publication.
Constitutive Activation	An oncology researcher would like a visual supplement that demonstrates to the readership of an immunology journal the mechanism of disease in which a key molecule in the signal transduction chain is constitutively activated, which produces an unregulated positive feedback loop.
Blood Flow	A researcher studying vascular flow would like a visual to supplement their publication that explains the variation of laminar flow (i.e. smooth movement of fluid with no swirls), in normal hemodynamics (i.e., blood flow behavior).
Aneurysm	A researcher publishing in a medical venue would like to include a supplementary image or animation to describe the final shape of an aneurysm, resulting from abnormal hemodynamic forces (i.e., blood flow in helical or swirling patterns) and morphological properties of the vessel wall.
Metastasis	A radiation oncology researcher publishing in an oncology journal is focused on describing the metabolism and movement of metastatic tumors as the basis of validation for their novel radiation therapy approach.

Our target survey participants included biomedical illustrators, visualization researchers, clinicians, and domain scientists familiar with the five topics. We were also interested in responses from non-experts, i.e., participants unfamiliar with the topic and lacking formal training in this area.

Survey. We followed Tory’s [504] principles for a comparative survey design, and organized topics so that a healthy/normal physiological topic precedes a corresponding pathological topic. This format provides the necessary pathological context.

We collected basic personal information and asked participants to self-rate their expertise on each topic. We used this self-rated expertise to divide participants into expert (4 or higher) and non-expert groups (3 or lower).

For each topic scenario, participants were asked to rank their top three and bottom three choices. The bottom choices were useful to capture, as this encouraged participants to explore negative aspects of a visualization that can be illuminating. For the top- and bottom-most choices we asked participants to quantify their selections according to four attributes: aesthetics, scientific accuracy, visual clarity, and communication success. Our use of these attributes was inspired by the work of Abdul-Ramen et al. [2]

Table 3.2: **Non-Expert** Audience Survey Scenarios

Topic	Scenario
Signal Transduction	An introductory biology student is studying for an upcoming exam. Their goal is to understand how a “message” is relayed through a series of messengers inside a cell.
Constitutive Activation	The same introductory biology student is tasked with identifying where in the signaling pathway a molecule is constantly activated when it should not be. This causes the entire signaling pathway to be always switched “on.”
Blood Flow	A person with little/no prior knowledge on the topic is interested in learning more about their body. They visit a popular health and well-being website, e.g. WebMD, to understand how blood moves and delivers nutrients throughout the body.
Aneurysm	A person has recently been diagnosed with a cerebral aneurysm. Their doctor shows them a visual to communicate what aneurysms are and why they must be closely observed.
Metastasis	A patient recently diagnosed with cancer has been told by their doctor that their cancer may metastasize, meaning that the cancer may spread to a different part of the body from where it began. To help them understand this concept, their doctor shows them a visual.

and by the Association of Medical Illustrators (AMI) juried salon judging criteria². We also asked participants to select keywords describing the relative strengths and weaknesses of their chosen assets. We drew these keywords from the previously mentioned AMI salon judging criteria. New keywords could also be entered manually. Finally, we included an option to add freeform comments. For further details on this survey design, we refer to Appendix D in this thesis.

We conducted a pilot study with five participants to test our survey design, and administered the final result via Typeform [349].

Results. To create a common foundation for comparing audience preferences within and between the five topics, we applied two **abstraction constructs**, model and visual abstraction, to every asset as depicted in Fig. 3.23. These draw from the terminology and definitions of abstraction by Viola et al. [530]. Model abstraction describes the relative knowledge precision, i.e., the creator’s mental model, of the input data and its temporality. Visual abstraction describes the relative visual simplification of the model. The abstraction spaces and asset scores (described below) for each topic are shown in Figs. 3.24, 3.25 and 3.26.

Asset scores are weighted such that $final\ score = 3s_1 + 2s_2 + s_3$, where s_1 , s_2 , and s_3 indicate the summed counts for an asset selected as the 1st, 2nd, or 3rd choice for a given scenario. We show these scores in the corners of each asset in Figs. 3.24, 3.25 and 3.26, and demarcate assets within the top 20th percentile for expert top (dark blue), expert bottom (dark red), non-expert top (light blue), and non-expert bottom (pink) choice selections.

A meaningful visual abstraction eases visual processing and reduces cognitive load [530]. Our findings indicate that both audiences preferred a middle space of visual and model abstraction. They dislike either extreme realism or extreme abstraction. An-

²https://github.com/lauragarrison87/Biomedical_Process_Vis/blob/main/4-Judging-Form-General.pdf

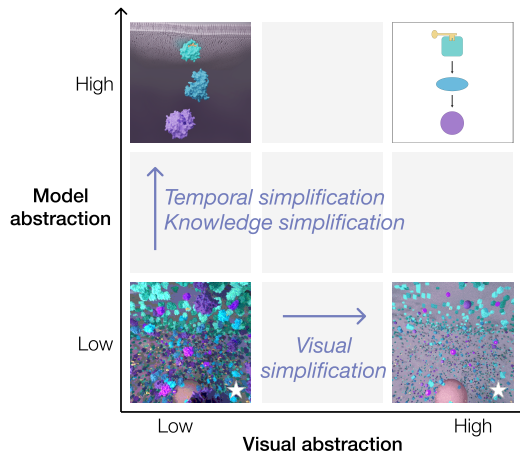


Figure 3.23: Conceptual abstraction space. Model abstraction spans the relative knowledge precision, i.e., the creator’s mental model, of the source data and its temporality, while visual abstraction encompasses the relative visual simplification of the model (stars denote animated assets) (from Paper D).

imated assets were often preferred to static representations, although in alignment with prior research we can see that the added value of dynamic visualizations is variable and highly dependent on audience and communication goals [232, 402]. We furthermore note that participant backgrounds likely unavoidably influence their selections, as people may be more likely to select what feels most familiar to them. This is especially noticeable for the expert top choices in the blood flow topic, shown in Fig. 3.25. Here, the top left group of illustration selections were predominately made by clinicians or biomedical illustrators, and the lower right grouping of streamtubes and streamribbons visualizations were more often chosen by visualization or domain scientists.

Fig. 3.27 shows the average **attribute rankings**, i.e., aesthetics, scientific accuracy, visual clarity, and communication success, for expert (upper row) and non-expert scenario (lower row) top and bottom choices for all five topics. Attribute rankings are similar for the top- and bottom-most choices across all topics, and are relatively similar between audiences. Scientific accuracy is the natural exception to this point, as non-experts cannot easily judge the accuracy of a topic that they do not know. An interesting exception, however, is the expert ranking for the aneurysm scenario. This has similar rankings for both top (Fig. 3.25, C12) and bottom (Fig. 3.25, C4) selections. This supports the point that communication success holds more weight than aesthetics in the selection of a visualization.

Participant **keyword** choices, as shown in Fig. 3.28, indicate different selection criteria for bottom choices, but similar selection criteria for top choices. This aligns with our observation that top choices overlapped more extensively than the bottom choices between the two audience groups. This suggests that participants may place an equally high priority on positive visual clarity and communication-related factors, i.e., *informative, easy to read, clear*. However, their criteria for a poor visualization differ, as does their concept of what constitutes *confusing*. Experts consider *simplification* confusing, while a non-expert audience finds overly *distracting* or *excessive* visualizations

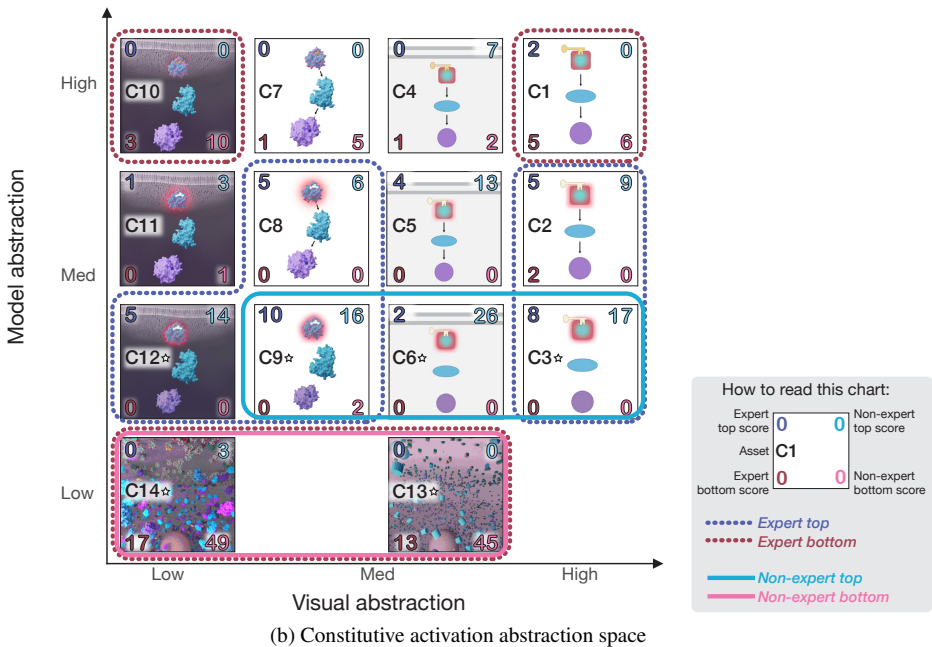
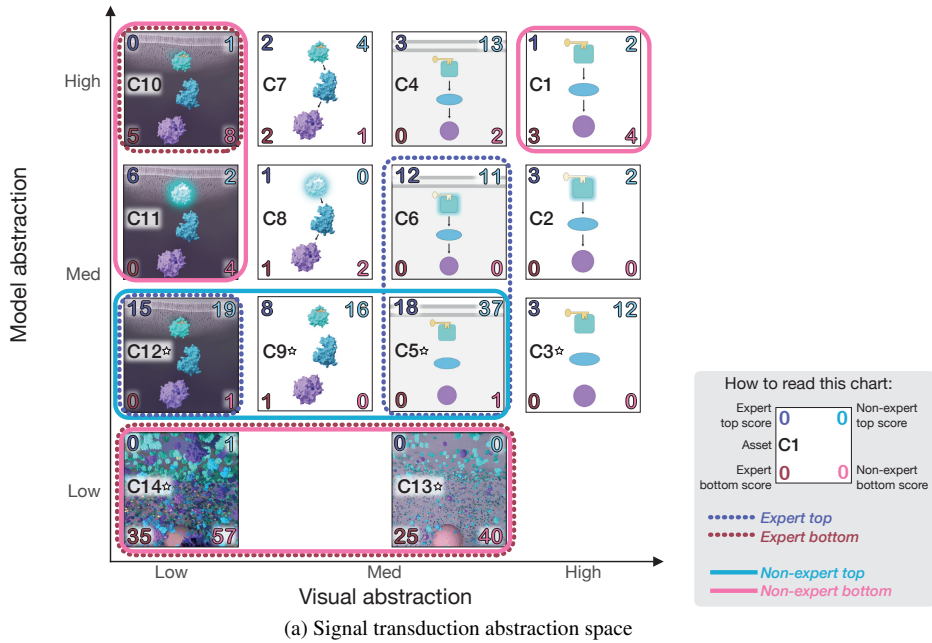

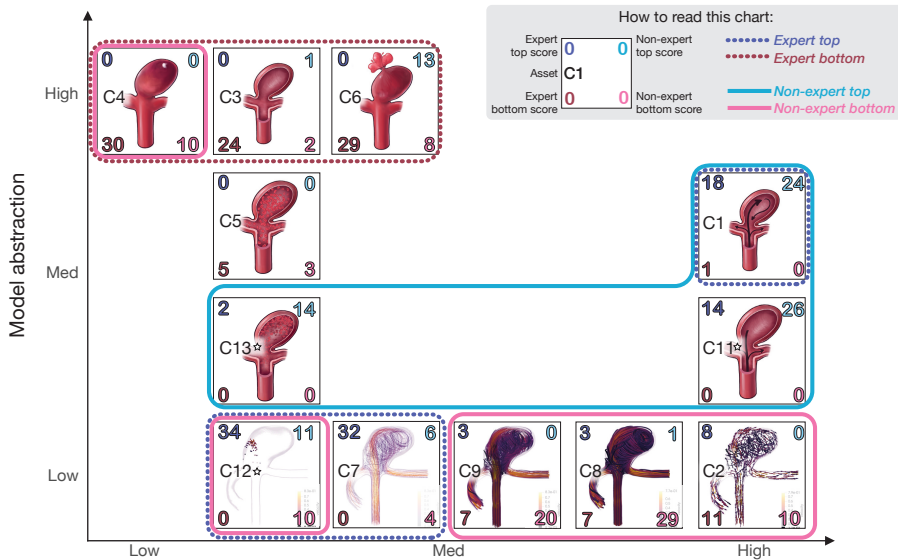


Figure 3.24: Assets are arrayed in the space by degree of model (y-axis) and visual abstraction (x-axis). Animated assets are denoted with a star glyph to the right of the asset name. Values in the four corners of each asset represent a weighted score for its selection frequency as the first, second, or third choice for an expert or a non-expert audience scenario. Encircled regions indicate assets with scores in the 20th percentile of each scenario (from Paper D).

3



(a) Blood flow abstraction space



(b) Aneurysm abstraction space

Figure 3.25: Assets are arrayed in the space by degree of model (y-axis) and visual abstraction (x-axis). Animated assets are denoted with a star glyph to the right of the asset name. Values in the four corners of each asset represent a weighted score for its selection frequency as the first, second, or third choice for an expert or a non-expert audience scenario. Encircled regions indicate assets with scores in the 20th percentile of each scenario (from Paper D).

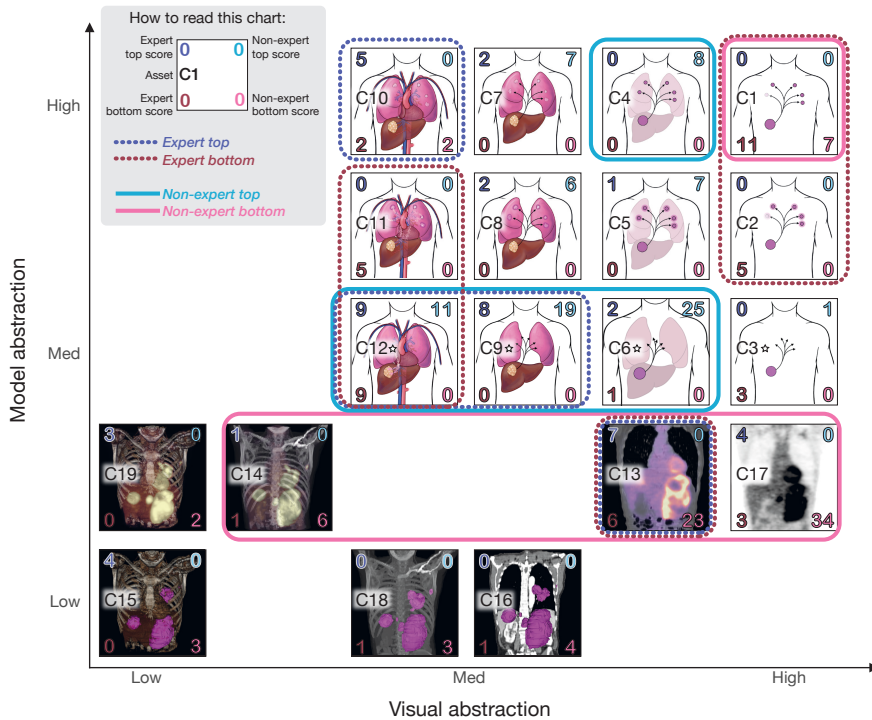


Figure 3.26: Metastasis abstraction space. Assets are arrayed in the space by degree of model (y-axis) and visual abstraction (x-axis). Animated assets are denoted with a star glyph to the right of the asset name. Values in the four corners of each asset represent a weighted score for its selection frequency as the first, second, or third choice for an expert or a non-expert audience scenario. Encircled regions indicate assets with scores in the 20th percentile of each scenario (from Paper D).

confusing. This makes sense anecdotally, as without sufficient background or explanation, an audience can become easily lost in a visualization.

Detailed Preferences Per-Topic. Finally, while not part of the core contribution of this work, we additionally created finer-grained visualizations of participant preferences on a per-topic basis using Tableau Public that are intended for communicating and easily sharing our results to a broader audience. Selected screenshots of these visualizations are shown in the lower right quadrant of Fig. 3.31.

Study Summary. Our findings from this study show frequent overlap in abstraction preferences between expert and non-expert audiences, with similar prioritization of visual clarity and the ability of an asset to meet a given communication objective. Aesthetics is not always a desirable attribute for experts, especially if the asset fails to meet the stated communication objective. We also found that some illustrative conventions are unclear, e.g., glows have broadly ambiguous meaning, while other approaches were unexpectedly preferred in some scenarios, e.g., biomedical illustrations in place of data-driven visualizations. Our findings suggest numerous opportunities for the con-




Figure 3.27: Expert and non-expert attribute rankings for top and bottom choices for all five topics (from Paper D).

3



Figure 3.28: Word cloud of keywords chosen to describe top and bottom choices for expert and non-expert scenarios for all topics (from Paper D).

tinued convergence of visualization and biomedical illustration techniques for targeted visualization design. Semantically-meaningful use of color in guidelines developed with stakeholders in both fields and in other relevant application domains is one such direction, which we summarize in the next section and detail in Paper E.

3.3.2 Color Semantics in Molecular Visualization

Visualization and biomedical illustration techniques can provide a window into complex molecular worlds that are difficult to capture through experimental means alone. Biomedical illustration techniques frequently employ color to tell a molecular story, e.g., to identify a key molecule in a signaling pathway. Visualization techniques can apply color in illumination models to cue structural features on molecular surfaces [195, 490] or to draw attention to structures of interest across multiscale molecular visualizations [256, 535]. However, color assignment to molecules on the whole is largely arbitrary, and often chosen based on cultural factors, the client, or the content author's personal taste.

Molecular biology is still a relatively young field, with some stakeholders arguing that establishing color guidelines would throttle the field's growth. Instead, content authors are free to choose an aesthetic supporting the type of story they wish to tell. However, this creative freedom comes at a price. Color design is challenging, particularly for those without training in color theory. Varied color palettes for the same structures dilutes their semantic meaning, which can negatively impact a visualization's interpretability and effectiveness. For example, COVID-19 spike proteins have been illustrated in a variety of colors. However, it is unclear if a blue spike protein is still immediately recognizable relative to the red used in the well-known version produced by Alissa Eckert and Dan Higgins for the CDC³. If such structures are not easily recognizable, how does this impact the success of, e.g., public health communications?

The CPK coloring convention [92] is a well-known and widely used standard for coloring atoms. The cellular-scale also exhibits fairly established coloring conventions, where red for red blood cells is perhaps the most obvious example. Immune cells are also often depicted in cool tones that echo the soothing blue color popular in medicine. Molecules can similarly be classified according to their various structures and functions. With an established standard for coloring atoms, and an informal semantic coloring practice for cells, it is reasonable to consider guidelines for semantic coloring of molecules.

In this final contribution of the thesis, we first discuss some factors that contribute to this array of color palettes. Next, we provide a brief sample of color palettes employed in the industry and research sectors. Finally, we suggest considerations for developing guidelines for color palettes applied to molecular visualization. A detailed discussion of these points can be found in Paper E.

Drivers of Palette Choice. Several factors influence a molecular color palette. The pharmaceutical industry is a major driver of the development of molecular visualizations. These companies often request palettes that adhere to corporate branding or convey a particular mood. Visualizations generated for academic or educational sec-

³<https://phil.cdc.gov/Details.aspx?pid=23311>

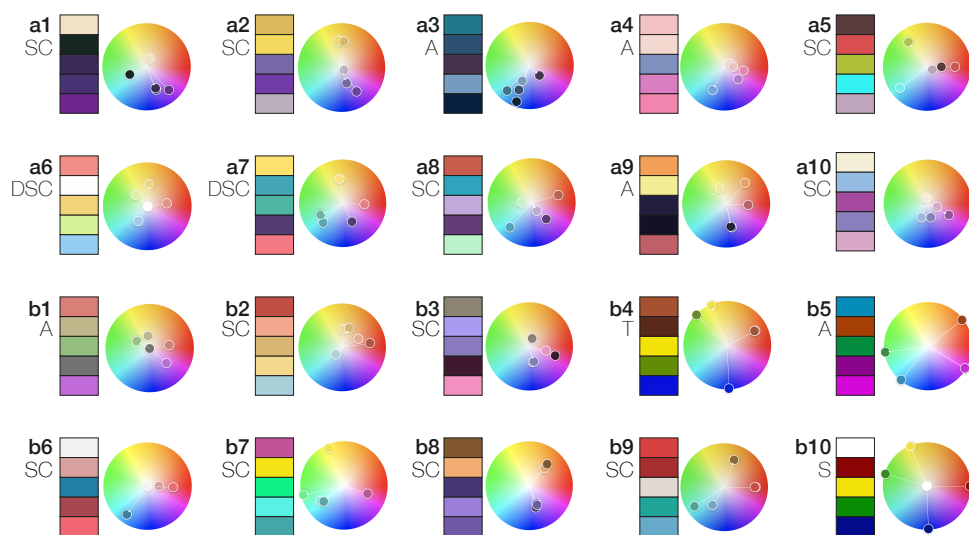


Figure 3.29: 20 sampled molecular visualization color palettes. (a) 2021 AMI Online Salon color palettes. (b) 2021 VIZBI Protein poster color palettes. *Rule abbr*: SC: split complementary, DSC: double split complementary, A: analogous, T: triad, S: square (from Paper E).

tors may use color to engage viewers and use colorblind-safe colors to reach a broader audience. Culture is also a significant driver of color selection, where different cultures may have different color-mood associations. Many molecular visualizations incorporate blue, or a similar color, into their palette for the calm, comforting emotions that are often associated with this color [4, 37, 96]. North American biomedical illustrators often use red to indicate abnormal molecular activity, e.g., constitutive activation. This choice is not necessarily obvious to other cultures. However, the most likely factor for color choice in molecular visualizations is the content author's personal tastes and aesthetic preferences.

Color Palettes in Practice. We conducted a small study to demonstrate the broad use of color in contemporary molecular visualizations. We extracted the color palettes from 20 molecular visualizations that were produced in the last year to ensure that our sampling captures recent color design trends.

Since such works may be created by bioinformaticians, biomedical illustrators, structural biologists, and visualization researchers [148], we sampled ten palettes from each of two venues that attract these professions. The Association of Medical Illustrators (AMI) is a global society of biomedical illustrators which hosts an annual juried salon. The ten palettes we sampled from this salon included works where molecules are the main focus, either in static images or animation stills. Visualizing Biological Data (VIZBI) is an annual meeting bringing together diverse professions to discuss research advances in the visualization of biological data. Each meeting includes a poster session that is divided into categories. We sampled ten palettes from the 2021 Proteins poster category. While VIZBI is advertised as an international venue, we note a North

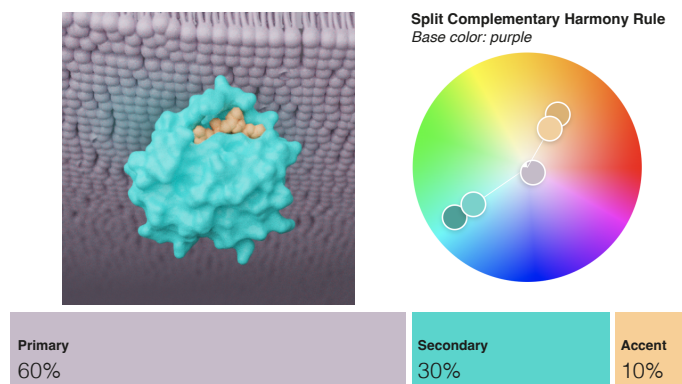


Figure 3.30: Example of 60-30-10 rule used in a biomedical illustration. Explorable color palette at https://color.adobe.com/color-name_LG-color-theme-19646985/ (from Paper E).

American bias in our sampling of biomedical illustration work from the AMI, as this organization is primarily oriented to North America.

We used Adobe Color [6] to generate color palettes from each molecular visualization. This tool enables semi-automatic extraction of a color palette from an image, which is then stored as a “custom” color harmony rule. Our intervention was necessary to manually match this custom palette by eye to the closest harmony rule. The results are shown in Fig. 3.29. This figure shows *split complementary* as the most commonly used rule, employed in 11 palettes across both groups. Two additional *double split complementary* palettes, which is closely related to the *split complementary* rule, are used in the AMI group. Palettes using an *analogous* rule are the most abundant (five), and occur more frequently in the AMI group.

In both groups, molecular coloring is consistent for structural elements within each visualization, e.g., molecules comprising a cell membrane are all colored the same or *analogously*. However, the semantics of the color choices are often unclear, except in the case of the red COVID-19 spike protein in b2, b6, and b9. Focusing attention or creating a visual hierarchy to drive a narrative appear to be the primary motivations in color application. Across all samples, the coloring of particular types of molecules, e.g., ligands, is only sporadically consistent. However, a stronger trend may be revealed in a larger sample.

Considerations for Best Practices. **Aesthetics** can crucially draw and guide attention within a molecular visualization. The *60-30-10* rule of thumb, borrowed from interior design, is useful to guide the balance of color in a molecular visualization. Exemplified in Fig. 3.30, 60% of a visualization should use a primary color, 30% a secondary color, and 10% an accent color. Molecular visualizations following basic perceptual principles and defined color harmony rules can be more aesthetically-pleasing, and furthermore easier for audiences with different degrees of expertise to interpret [231, 239, 561].

Color is an important element in a molecular visualization’s **interpretability**, or readability. It can be used as a device to focus audience attention to salient features of

the visualization, which may ultimately lead to a more effective visualization. Although different cultures or contexts may associate different semantic meanings, saturated, light value, or warm colors tend to draw attention, while desaturated, dark value, or cool colors often recede, especially when used with, e.g., warm colors [222]. This perceptive feature can facilitate the interpretability of a molecular visualization by using color to draw attention to the principal elements of a visualization. Importance functions may be useful in developing rule-based methods that could assign an appropriate hue, saturation, and lightness values to elements in a molecular visualization.

Finally, a molecular visualization is **effective** if it is correctly read by the intended audience. Just as a general audience may not know precisely, for example, the structure or function of a red blood cell in the body, through its color assignment they can recognize some of its basic properties and relate it to the flow of blood in our bodies at a larger scale. We argue that this strategy can be extended to the molecular-scale by coloring hemoglobin red. Molecules may also be colored according to the class of pathway that they are involved in, e.g., a signal transduction versus a gene expression pathway. This coloring may expand to include the use of glow or fresnel effects to show abnormal pathway activity. Alternatively, molecules may be assigned color families and harmony rules based on their structure, e.g., the secondary or tertiary structures of a protein. Establishing guidelines that retain the flexibility for creative expression and innovation is challenging and requires further research and stakeholder discussions. However, with the increasing prevalence of molecular visualization in mainstream culture, the implementation of semantically-meaningful coloring guidelines can improve public understanding of molecules.

3

3.4 Research Replicability

The final and often-overlooked step of any research project is making data and code publicly available for other researchers to attempt to replicate experimental results, and to enable other works to build from these data and code. In some instances, data confidentiality makes this step impossible. However, when possible in this thesis we made a concerted effort to make our code publicly available, to document the steps necessary to set up and run our code locally, and/or to make available the source data used for our case studies in validating our approaches (Fig. 3.31). Ensuring that visualization research is both available and replicable for domain experts to more easily use is one way to bridge the data accessibility gap that we often face in visualization for physiology.

The **Graphics Replicability Stamp Initiative** (GRSI)⁴ is an organization that certifies the replicability of code and results from data for sharing in and beyond the visualization community as a non-commercial resource. Our source code for *SpectraMosaic*, available at <https://github.com/mmiv-center/spectramosaic-public>, and *DimLift*, available at <https://github.com/lauragarrison87/DimLift>, have both been certified through this process. In open-sourcing and certifying the replicability of both of these applications, we hope to enable scientists from the application domains that these works target to be able to use these tools for visual exploration and analysis of their own data.

⁴<http://www.replicabilitystamp.org/>

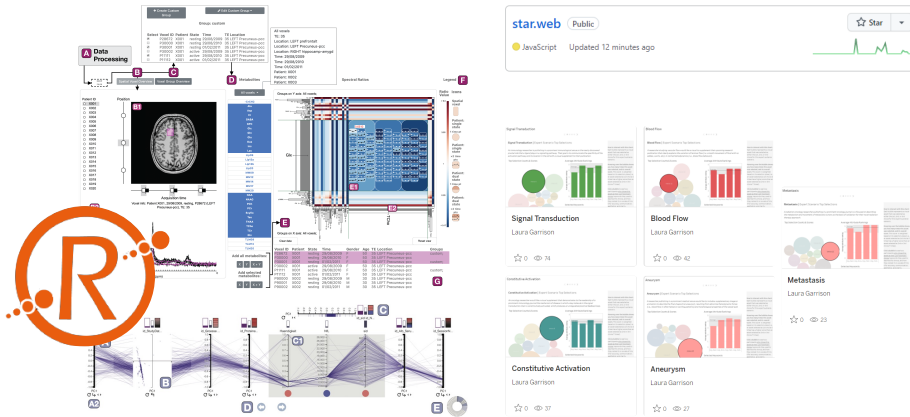


Figure 3.31: Our two applications contributions, *SpectraMosaic* and *DimLift* are certified through the GRIS. Our STAR literature visual explorer and database (upper right) is also open-sourced. We used Tableau Public to create easily adjustable and shareable visualizations of the results from our survey on the practice and preferences of visual communication of physiological processes (lower right).

We also assembled a data package comprising the results of our exploratory study on the practices and preferences of domain experts and non-experts for the visual communication of physiological processes. Our main database stores expert and non-expert preferences and feedback for five different topics that are commonly visualized in physiology: signal transduction, constitutive activation, blood flow, aneurysm, and metastasis in expert and non-expert scenarios in each topic. The data package furthermore contains all visual assets created for the study (75 in total). Their use is permitted through a Creative Commons CC-BY-NC 3.0 License⁵, meaning that they can be freely shared and adapted for non-commercial purposes. We also include screenshots of each page of the survey for other researchers interested in reproducing our survey design in their work. The complete data package is available at https://github.com/lauragarrison87/Biomedical_Process_Vis. Part of this work includes a set of publicly-available visualizations of the data using Tableau Public (<https://public.tableau.com/app/profile/biomedsurvey2021>). These visualizations are freely editable and customizable for other users, and can be embedded onto other pages to share the results of this research.

Lastly, we have open-sourced the code and database for our STAR Literature Visual Explorer tool, available at <https://github.com/lauragarrison87/star.web>. Interested visualization and application domain researchers can follow the documentation to set up and run this tool locally, or run from the github server directly. The database can optionally be updated with new literature with a pull request, giving this tool the possibility to be a living repository rather than a snapshot of the state-of-the-art at the time the literature was collected.

⁵<https://creativecommons.org/licenses/by-nc/3.0/>

Chapter 4

Discussion & Outlook

The future is already here — it's just not evenly distributed.

William Gibson

4

The myriad interconnected biochemical processes that contribute to physiology are complex, yet necessary, to understanding life. Data-driven research has introduced a new paradigm where researchers first acquire physiological data in ever-increasing resolution and volume and then try to ask questions of the data. Increasingly multi-disciplinary teams are necessary to solve challenges in exploring, analyzing, and communicating these data, and visualization plays an integral role in identifying and solving these challenges.

This thesis brings a cross-disciplinary perspective to visualization for physiology for expert and non-expert user profiles and their according high-level tasks over multiple spatio-temporal scales. Our work contributes to the areas of theory, empirical findings, method, application, and research replicability. We began with a broad discussion of the landscape of visualization for physiology spanning the visualization as well as application domains, which we classified in a two-level system according to task and spatio-temporal scale. Using this classification we identified two important task areas for further investigation: exploratory analysis and communication. In the area of exploratory analysis we started narrow with a specific under-visualized data type, magnetic resonance spectroscopy. We conducted a design study that extensively reviewed data characteristics, user tasks, and contributed a novel method and application for exploratory analysis of these data. Our work then broadened from a specific data type to multivariate, mixed clinical cohort data and developed a method and corresponding application enabling discovery of subtle patterns in these data, and briefly explored the effect of imputation on missing data in this context. Moving from expert-oriented tasks we pivoted to visual communication for experts and non-experts. Curious about the overlap in techniques between biomedical illustration and visualization, we conducted a survey of the practice and preferences of communicating a subset of processes in physiology to expert and non-expert audiences. Part of this work included the development of a visual asset library that capture standard approaches to visually communicate these processes from both disciplines. The findings of this work led to our proposal for the development of best practices for inclusive and semantically-meaningful color

palettes in molecular visualizations. Finally, all application code and study results are open-sourced and available for reuse by the scientific community and public.

The work presented in this thesis is not comprehensive in its coverage and discussion of visualization for physiology. To do so within a manageable scope would have been impossible. Instead, our work turns to representative topic areas in an effort to capture the visualization opportunities and challenges this field represents. The studies conducted in our work are also predominately qualitative, with relatively small sample sizes. While the results are sufficient to draw preliminary conclusions, particularly in narrow fields such as magnetic resonance spectroscopy research, a clear opportunity for further work is building on these studies with larger and more demographically-varied cohorts. This is particularly true in the case for developing guidelines for semantically meaningful palettes in molecular visualizations. Beyond conducting larger studies to further validate the ideas we propose, we see a number of exciting opportunities for future research directions from our work, some of which we highlight in the remainder of this chapter.

4

Core to what makes a visualization “effective” is its ability to put information in front of people so that they can understand and make decisions from this information¹. This is true whether the high-level goal is to explore, analyze, and/or communicate facts about the data being visualized. The use of visualization technique is thus driven by the user’s main goal for the data, and is further impacted by their expertise or background knowledge. The methods that we developed for multifaceted exploratory analysis of physiology data are developed for experts in a research setting. However, our ultimate goal, which is shared by many who develop visual methods for medical data, is to see our work realized in clinical routine where clinicians and patients can see an immediate and positive impact on health management. While our heatmap matrix display with nested glyphs for metabolite exploration is useful for spectroscopy researchers who already possess a degree of experience in spectral data interpretation, this method inside *SpectraMosaic* is too time-consuming and unwieldy for clinicians to use effectively in their daily work. Given the immense value of MRS in assisting diagnostics, continued work on visual solutions for clinical interpretation of MRS is an exciting research direction. The same is true for our *DimLift* method—while designed as an upstream approach for e.g., clinical researchers to identify new biomarkers for disease, continued work that explores adapting this method for clinicians to use opens a new line of inquiry that can be of enormous help to clinicians who are buried in diagnostic data.


Framed in the context of the user’s tasks and goals, we have learned that the scale of the process in question also has an impact on the visualization technique. Processes occurring over broad scales, whether spatial, temporal, or both, present extra visualization challenges that require the visualization designer to carefully balance insights with complexity. This is particularly challenging and necessary for molecular dynamics simulations, or any data that capture molecular activity *in situ* in cells, tissues, or even organs. These data are the future of understanding the body, and visualization plays a crucial role in facilitating our understanding of these data.

Dimensionality reduction strategies are common with a number of physiology data types, and the ongoing visualization challenge is ensuring that the user can relate the results back to the original data. Our state-of-the-art report on visualization for physi-

¹Ben Fry, VIZBI 2022 Keynote

ology found that many researchers in the application domains often find it necessary to review the raw data as part of their analysis, e.g., time-lapse imaging of cellular dynamics. Fully automatic solutions are not that desirable, as they can crucially inhibit user trust and understanding of the results of any data derivations. Semi-automated systems that allow the user the flexibility to steer their own inquiries if they find something interesting in the data are beneficial to continue exploring, building on the concepts we presented with, e.g., our *DimLift* method. Such approaches may be effective for experts with mainly analytical goals, but are often too abstract for non-expert users. This introduces a delay in the sharing of information that might be critical to, e.g., public health. As a community we need to consider how our bespoke approaches for a certain type of expert can be adapted or shared with other communities, as well as to the general public. Additional studies building on the findings in our survey on the practice and preferences for communicating physiological processes may help identify strategies to adapt complex visualization approaches for experts in a target domain to experts in a different domain, or to the general public.

A broad range of methods and modalities capture physiological data. The majority of methods and modalities are limited to a particular spatio-temporal range in physiology, e.g., metabolite concentration detection or changes in vascular diameter through imaging. This limited range correspondingly limits the availability of visualization techniques that can still accommodate a given visualization task for a certain level of expertise. Furthermore, some methods and modalities have received little attention from the visualization community, many of which we identified in our state-of-the-art report, e.g., hierarchical phase-contrast tomography (HiP-CT) [536] and nascent chain tracking (NCT) [105]. This will continue to be the case as hardware and software advances continue to yield new paradigms for data acquisition. The development of visual approaches that facilitate not only initial exploration and analysis of these data, but ultimately communication of these data to various stakeholders, is critical for advancing public health and health literacy in our global society. The research presented in this thesis provides a valuable foundation for developing visual methods for exploratory analysis and communication of multifaceted physiological data.



*“Time is the substance I am made of.” —Jorge Luis Borges, **Labyrinths***

Part II

Scientific Results

Paper A

Trends & Opportunities in Visualization for Physiology: A Multiscale Overview

Laura A. Garrison^{1,2}, Ivan Kolesar³, Ivan Viola^{4,5}, Helwig Hauser^{1,2}, and Stefan Bruckner^{1,2}

¹Univ. of Bergen, Norway, ²Mohn Medical Imaging and Visualization Centre, Haukeland Univ. Hospital, Norway, ³Rainfall AS, Bergen, Norway ⁴King Abdullah University of Science and Technology, Saudi Arabia ⁵TU Wien, Vienna, Austria

Abstract

Combining elements of biology, chemistry, physics, and medicine, the science of human physiology is complex and multifaceted. In this report, we offer a broad and multiscale perspective on key developments and challenges in visualization for physiology. Our literature search process combined standard methods with a state-of-the-art visual analysis search tool to identify surveys and representative individual approaches for physiology. Our resulting taxonomy sorts literature on two levels. The first level categorizes literature according to organizational complexity and ranges from molecule to organ. A second level identifies any of three high-level visualization tasks within a given work: exploration, analysis, and communication. The findings of this report may be used by visualization researchers to understand the overarching trends, challenges, and opportunities in visualization for physiology and to provide a foundation for discussion and future research directions in this area.

A

A.1 Introduction

Human physiology describes the functions and mechanisms of the human body that make it a living being. Forming the link between the basic sciences (biology, chemistry, and physics) and medicine, human physiology is multiscale in that it integrates the individual functions of molecules, cells, tissues, and organs into a whole organism [180]. Physiology is an important aspect of systems biology, which has been characterized as an approach to understanding multiscale interactions in a biological system [259]. While systems biology tends toward data-driven and quantitative methods, an integrative physiology approach emphasizes concepts through experiments and observation across multiple scales [164]. The multiscale nature of physiology allows us to, for example, link how signaling events at a molecular level lead to the normal, i.e., healthy, contraction of cardiac muscle in a normal heartbeat. An understanding of the normal processes and functions of the body allows us to recognize those that are abnormal, such as in atrial fibrillation, a heart problem where the upper chambers of the heart do not follow a regular beating pattern. With recent advances in hardware and software, as well as in experimental and imaging modalities, it is now possible to model many of these processes across several scales. Consequently, it is time for a discussion of visualization techniques for multiscale physiology. This survey provides a broad overview of common approaches and highlights research opportunities in visualization for physiology across multiple scales.

Modern clinical workflows involve a battery of tests and imaging protocols related to physiology. These are used to guide therapy, monitor disease progression or treatment response, and identify new biomarkers for medical research. Improved technology and hardware capture an unprecedented volume and diversity of data through models and simulations, e.g., advanced numerical simulations of blood flow, as well as through various acquisition techniques, e.g., *fluorescence lifetime imaging microscopy* (FLIM). Data range from 2D to 3D images, from static to time-dependent, from scalar to vector to tensor fields, and are often multivariate. The visualized physiological processes range spatially from nanometers to full body length and temporally from femto-/nanoseconds up to hours, months, and, in some cases, even years, as shown in Fig. A.1. However, these data are often specific to a particular and relatively narrow spatio-temporal scale, and establishing links between these multimodal data types from the nano- to macroscale has been described as a grand challenge for many years from the perspective of systems biology [375, 378], visualization [159, 385], and in a multidisciplinary 2018 Dagstuhl Seminar [7]. Linking these data requires multidisciplinary teams to develop analytical models and visualization approaches that can bridge the range of spatial and temporal scales. The Physiome/Virtual Physiological Human and affiliated subprojects [30, 137, 211, 498, 527] have aimed to model processes that range from the molecular to organ scales, and beyond, to understand the multiscale interplay of physiology. The National Institutes of Health's Human BioMolecular Atlas Program (HuBMAP) aims to comprehensively map the human body at single-cell resolution from both a structural and functional perspective [208]. Numerous works have blossomed from these initiatives, such as OpenCMISS-Zinc [63], a library for building multiscale models and visualizations of physiological processes.

Despite the wealth of collected and simulated data for physiology, not all of this information can be, or is optimally, visualized through data-driven means. Hand-crafted




Figure A.1: The general spatial and temporal ranges of human physiology, partitioned according to scale. Bold text indicates areas of focus in this survey, with example processes labeled in each scale.

medical illustrations are an alternative or supplement to data-driven visualization for representing physiology. Illustration remains ubiquitous when communicating information to a broad audience where simplification and abstraction of concepts are essential [168, 180, 239, 427, 434]. More generally, illustrations are invaluable in communicating abstract concepts, theories, and models. In this sense, illustration can provide a source of inspiration for abstraction in data-driven visualization. Paired with computer-supported solutions, an illustration can be brought to life through interactivity and adaptability to different scenarios. However, the time and labor cost for creating such illustrations prevents their use for, e.g., patient-specific data visualization. Throughout this report, we highlight select illustrative works to demonstrate opportunities where illustration can inspire or augment data-driven approaches.

Physiology has received extensive attention from the visualization community but in a fragmented, unevenly distributed form across subtopics, data sources, and visualization techniques. Few of these works extend their focus beyond one or two scales, e.g., only molecular [267], molecular and cellular scales [169], or organ [287]. This paints a limited picture of the true multiscale nature of physiology. Similarly restricted in scope are surveys on a particular data type, e.g., PC-MRI by Köhler et al. [260]. Technique surveys, such as by Bach et al. on space-time cubes [33], McGee et al. on multilayer network visualization [331], or Preim et al. [401] on medical animation, may span multiple scales, but physiology is often only one application area of many under discussion and is not the primary focus. Still others have explored the multiscale challenges in visualizing biomedical data from a high-level perspective [10, 79, 245, 330, 523], though without a specific focus on physiology. While our survey does not go into the level of detail that these surveys visit for their respective areas, the novelty of our work is in presenting a unified overview across multiple scales. We do so by discussing the coverage

of these surveys alongside representative individual works that contribute to the same scale. We intend this report as an introductory resource for the space of challenges and opportunities for visualization research applied to physiology. Our framing for the work we survey additionally provides a different perspective than related work. We embed each article in a spatio-temporal context and draw from Brehmer & Munzner [66] to discuss its contribution according to the high-level user task(s) that it addresses.

Perhaps closest to our survey in terms of the scales of biological organization covered, Secrier & Schneider [452] discuss general visualization techniques from the bioinformatics domain for physiology from the molecular to population scale, but this review is brief and high-level. O'Donoghue et al. [376] review the use of omics and imaging data in biomedical research from molecule to population level for the primary purpose of exploration. However, their discussion is from a systems biology and bioinformatics perspective and mainly focuses on the visualization of molecular data to understand multiscale physiology. Our work covers a broader set of data types and a wider range of physiological processes.

To our knowledge, this work is the first of its kind to broadly overview the space of visualization for physiology that covers a scope similar to Lipşa et al.'s survey of visualization for the physical sciences domain (astronomy, chemistry, etc.) [305]. Our main contributions include:

- This is the first literature survey paper of its kind that provides a view into mature and open opportunities in visualization research for physiology. Our work surveys both within and beyond the core visualization venues.
- We focus the content of our survey on physiology topics that are highly-researched and cited both within the visualization community and in related physiology domains.
- We introduce a novel taxonomy that addresses these different topic areas and their respective opportunities by embedding works within a spatio-temporal context according to the high-level visualization task that they address.

In the following, we provide a brief background on physiology in Sec. A.2, followed by a discussion of our survey methodology (Sec. A.3) and classification structure (Sec. A.4). Sections A.5–A.8 are each dedicated to a spatio-temporal scale in our taxonomy, in order of increasing biological complexity. In each section, we first introduce the necessary background information and relevance of the physiological processes discussed at the given scale. We then provide an overview of visualization conventions and trends that we observe in the related literature according to task. A brief discussion of the mature and open challenges in visualization concludes each scale before transitioning to the next section. Fig. A.2 provides an overview of our organizational approach for these sections. Readers interested in a specific topic, e.g., molecular pathway visualization for exploration, may easily navigate to the part of the paper that is most relevant to their interests and needs. Section A.9 provides an overview of true multiscale visualizations uncovered in our search and leads into a discussion of the challenges and research outlook on visualization opportunities for physiology (Sec. A.10–A.12).




Figure A.2: Our approach for Sections A.5–A.8 organizes physiology topic areas in order of increasing spatio-temporal scale, with detailed per-scale and per-process points for each topic area (when relevant.)

A.2 Physiology Background

Normal human physiology requires a careful balancing act, known as homeostasis, of numerous processes that occur over a broad span of time and space, as shown in Fig. A.1. The smallest entity in the human body with the functional characteristics to sustain life independently is the **cell**. The cell itself contains **molecules**, such as water and ions, and organic molecules, such as proteins, that participate in processes necessary for its survival. Genes are the basic unit of heredity in cells that are made up of DNA and which encode the synthesis of RNA, which directs protein synthesis. Genes, proteins, and other molecules interact in sequences of reactions and interactions that are described as pathways. These pathways form networks and contribute to specific cellular functions. Molecular structures fall generally in the range of nanometers, and **molecule-scale processes** over a broad temporal range from femtoseconds, e.g., bond vibration between atoms in a molecule, to seconds, e.g., global motions or reaction sequences in a molecular pathway [193], to minutes and hours in the case of pathways involved in metabolism, gene expression, and signal transduction [31].

Human and other eukaryotic **cells** contain specialized cellular structures called organelles that participate in and facilitate the molecular pathways that keep the cell functioning. The mitochondrion is one type of organelle known as the energy powerhouse of the cell, while the nucleus is another organelle that provides the housing for our genes. A cell has the ability to communicate and exchange nutrients with its environment through its semipermeable membrane. This membrane contains specialized molecules, known as receptors, as well as channels and other structures that facilitate communication and exchange. **Cell-scale processes** relate to cells that average in the range of tens of microns in size, and with a temporal range of milliseconds, e.g., ac-

tion potential generation, to minutes, e.g., mitosis, to a day for a complete cell cycle in humans.

Human cells differentiate into specialized cells with shared properties that group together to form **tissue**, which is classified into four different types: muscle, epithelial, connective, and nerve tissue. The physiological properties of these different tissues reflect the function that they serve. Cardiac muscle tissue, for instance, is responsible for the periodic contraction of the heart. Skeletal muscle tissue moves our limbs, while smooth muscle tissue moves food through the digestive system. Abnormal tissues occur where the comprising cells take on different characteristics than normal tissue. For example, cancerous tumors have completely different tissue features than the surrounding tissue in which they occur. **Tissue-scale processes** span hundreds of microns to millimeters and temporally range from milliseconds, e.g., signal propagation, to weeks or even months, in the case of tissue growth and development.

Organs are composed of different types of tissue and perform major physiological processes according to their location, form, and composition. For example, one of the functions of the heart is to pump blood that contains life-giving oxygen and nutrients to the body's cells. **Organ-scale processes** that we consider in this report tend to fall within a narrow spatio-temporal window: organs like the heart measure in the range of centimeters, and the temporal range of a complete heartbeat or a complete breath cycle is in the range of seconds. Organs with similar functions are grouped into systems. For example, the cardiovascular system consists of the heart and blood vessels that are responsible, among other tasks, for carrying oxygen and nutrients to the body's cells. Our organ systems are interdependent. The cardiovascular system cannot function without the respiratory system, which includes the lungs, because the lungs handle blood re-oxygenation. The healthy functioning of an organism is dependent upon the systems of the body working in concert.

A

A.3 Scope and Methodology

This survey sketches out trends and opportunities in visualization for physiology across multiple scales, with an emphasis on human physiology. Fig. A.3 provides an overview of our methodology.

Thematic Topics in Physiology. We restrict our survey to timely, highly-cited thematic areas in human physiology to ensure that our survey presents a relevant research agenda. For this, we used Web of Science's "hot papers" and "highly cited" filters with the keyword "physiology." A "hot paper" is any paper published in the past two years that has received enough citations to rank in the top 0.1% of papers in its field. A "highly cited" paper ranks in the top 1% of cited papers for its field and publication year. To get a sense of the diversity of topics, we took the top 20 papers from each of these filters and excluded works that did not relate to humans or other mammals. We keyed these papers to topical area of physiology, e.g., molecular pathways or heart function, following standard medical physiology textbooks [180]. For a complete list of these papers and their topical areas, we refer the reader to Tables 1 and 2 in Appendix A.

Search Criteria. Our survey focuses on visualization research works for understanding physiology. We excluded pure methods works, meaning that the visualization lit-

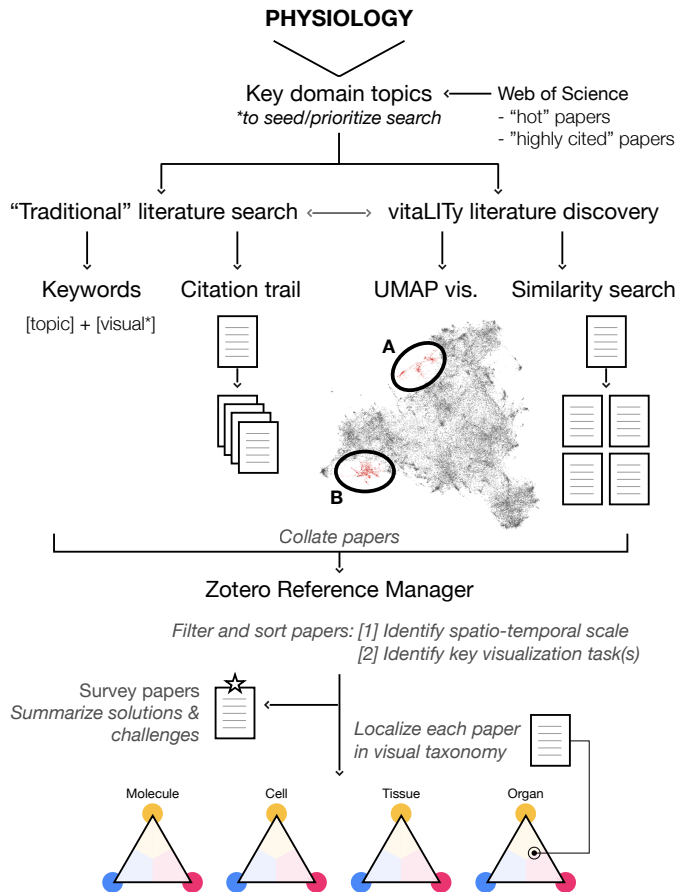


Figure A.3: Our literature search process included both traditional search methodologies and vitaLITY [361]. VitaLITY’s UMAP visualization allowed us to identify two main groupings of physiology-related visualization literature: (A) contains molecular-scale visualization literature, while (B) contains cell, tissue, and organ-scale works.

erature we included must have a clear discussion of the domain science as a possible application for the proposed method. We excluded works where the main visualization goal is to understand structure, although we included limited examples of instances where a physiological process is used to visualize a structure, e.g., 4D PC-MRI data to describe vessel boundaries [40]. We focused primarily on input data that is either itself dynamic, or is being used to capture snapshots of a dynamic process. We excluded purely longitudinal studies. We limited our search and discussion of works in areas that have already been well-covered in visualization and looked more comprehensively in less well-covered areas. In summary, we included application-oriented papers that center around a key topic area we identified from timely and highly-cited physiology research and that apply visualization in a novel way for the topic domain.

We focused our literature search on core visualization publication venues: IEEE TVCG, CGF, C&G, BioVis, VIZBI, and VCBM. The domain sciences may adapt vi-




Figure A.4: Distribution of works (excluding surveys) according to publication year.

sualization techniques in a novel way to interpret their data. We did not extensively review work published in domain-specific venues, but included a selection of works relevant to our included physiology subtopics to show visualization’s use from this perspective. These domain-specific works contributed to approximately 46% of the total works collected.

Search Process. We conducted our initial search using a combination of Google Scholar, PubMed, and IEEE Xplore based on keyword search [*physiology topic*] AND *visual**. The literature search was divided between two coauthors.

We used *vitaLITY* [361] to complement our search, a recent visual analysis tool that allows for serendipitous discovery of academic literature. The *vitaLITY* database at the time of this writing consists of 59,000 literature items from 38 computer science venues that include our core venues listed above. These are searchable in a standard table that includes paper title, abstract, keywords, and authors, as well as a similarity search and a 2D UMAP visualization of the embedding space for the entire collection. For details on these tool features, we refer to Narechania et al. [361]. In the UMAP visualization, we identified two main groupings of literature, shown in Fig. A.3, that helped focus our search: **(A)** groups works for visualizing molecule-scale processes: molecular dynamics, interactions, and pathways, while **(B)** includes works for visualizing cell, tissue, and organ-level processes. Within each of these groupings, we searched for existing surveys and state-of-the-art reports to identify saturated topics. For example, since a number of reports have been written on visualizing different topics at the molecular scale, we devote less space to discussing this scale in our work and focus more comprehensively on scales and physiology topics with less coverage. UMAP exploration also helped us to identify relevant individual papers. We used works found in *vitaLITY* to seed our more traditional search approach and vice versa. This allowed us to perform a more complete literature search that accounted for terminology differences between domains.

Refining Process. In a second detailed pass of our collected works, we reviewed titles, abstracts, and figures to determine topical fit for our survey. At this stage, we used the publication year as a secondary check for our search coverage. If necessary, we revisited *vitaLITY* for topic areas that had a publication gap and resampled papers from this publication time frame.




Figure A.5: Distribution of literature by spatio-temporal scale, excluding surveys. **Left:** Literature are sorted by molecule, cell, tissue, and organ scales. This chart counts literature only once, according to the scale to which they contribute most. **Right:** Works are visualized in terms of density over spatio-temporal space, encoded by darkness. The x-axis shows temporal scale in units 10^n seconds, while the y-axis describes spatial scale in units 10^m meters. The dark region at the upper center indicates an abundance of works to visualize organ-scale physiology over the range of seconds, e.g., one heartbeat. The dark region at the right corresponds to works visualizing gene expression data.

Collection Summary. Our complete literature set includes 366 works, 61 of which survey or provide an outlook on an aspect of physiological data visualization. Approximately $\frac{1}{3}$ of these works have been published in the last five years, with $\frac{2}{3}$ of the total set published in the last ten years. Fig. A.4 shows the distribution of works by publication year. The peak in publications in 2010 is a point we discuss in Sec. A.10. Following literature collection, we classified all papers according to a two-level taxonomy to help identify challenges and opportunities in this domain. Due to limitations in space, we discuss a subset of these works in this paper, with the full library available in supplementary material and at https://lauragarrison87.github.io/star.web/vis_tool.

A

A.4 Taxonomy and Overview

Physiology spans the basic sciences and medicine, requires diverse domain knowledge, uses myriad data types, and employs a wide range of visualization techniques. Classification by domain, e.g., biology, chemistry, physics, or medicine, may seem the most obvious approach. However, these sciences are tied into each process and are difficult to classify separately, especially at the molecular scale. Molecular reactions are dictated by biology, chemistry, and physics and are core to disease diagnosis in medicine. In addition, different domains often adopt slightly different terminologies and classification systems. For example, biology works from gene, protein, tissue, organ, system, body, while neuroscience follows a neurochemical, neuronal, region/network, and brain classification scale. This creates more confusion when organizing literature according to domain.

Although classification by data source may feel most natural in the context of the visualization pipeline, this does not provide a unified axis for the scales we survey. Simulations and models may span the molecule to organ scale but tend to be heavily




Figure A.6: **Left:** Distribution of literature according to scale and high-level task, the latter of which is adapted from Brehmer & Munzner [66]. **Right:** Many visualization approaches support a combination of exploratory, analysis, and communication task(s). Darkness and size dually encode the number of works that are categorized with a given task combination within each triangle.

focused on particular topics, e.g., heart function [29] or lung function [251]. While new imaging technologies, such as hierarchical phase-contrast tomography that maps organ to cell level, have come closer to realizing this possibility [536], no unifying technology yet bridges from the molecule up to the organ scale from a general physiological perspective.

A.4.1 Spatio-Temporal Organization

To minimize semantic collisions or confusion, we classify literature into scales along a spatio-temporal axis that is roughly discretized according to biological complexity: **molecule**, **cell**, **tissue**, and **organ**. This discretization is inspired by the organization of physiology textbooks [180]. Fig. A.5 (left) shows the distribution of non-survey works we collected that are categorized according to this scale. Works that span multiple scales are counted once for each scale, e.g., a work that we classify as both molecule and cell scale is counted in both the molecule and cell groupings.

We bundle temporality into this scale discretization based on the fact that, as structures increase in physical size, they tend to be involved in more biologically-complex processes that take more time to complete. This relationship between increased structural size/complexity and time has been discussed elsewhere in different domain contexts [106, 170, 452, 472]. We can observe this phenomenon in Fig. A.5 (right), which represents collected works classified in a range according to the scale that the input data spatially and temporally resolve to and up to the spatial and temporal scale of the structure and process of interest, e.g., the whole brain. For example, while EEG measures neural activity, we cannot visualize individual neurons with this modality, and that is not the intent of conducting these types of studies—visualizations of EEG data fall in the organ scale. The relationship between space and time is not perfectly linear, as reflected by the dark groupings in the upper center and right regions of the chart. The upper corresponds to organ-scale processes that occur over the range of seconds, such as a heartbeat or a full breath cycle, while the right corresponds to the time for the expression of a single gene [344].

Our classification system does not formally extend beyond the organ scale for a few reasons. First, restricting the scales we examine keeps the scope of this survey manageable. In addition, our preliminary searches found limited visualization works that exist purely at the system- or organism-level, and the tasks and visualization techniques implemented are similar to those observed at the organ scale. We briefly review examples

of works beyond the organ scale, as well as selected works that use true multiscale approaches, where the visualization aims to facilitate a task at three or more scales, in Sec. A.9.

Although an enormous range of physiological processes occur at all scales of the body, we focus our survey on a few categories of processes at each scale that are timely for the physiology domain. Processes occurring at the **molecule scale** include molecular dynamics, e.g., the motion of atoms and molecules, reactions between molecules where electrons and/or atoms are exchanged, and molecular pathways, which describe a chain of molecular reactions. **Cell-scale** functions that we highlight include cell dynamics and interactions, such as how the cell develops and communicates with its environment. We include the dynamics of the cell's organelles at this scale, such as mitochondrial activity. **Tissue-scale** functions consider the behavior of aggregates of cells of the same type and include tissue dynamics, such as growth, and tissue interactions, such as signal propagation in neural tissue. At the **organ scale** we consider processes related to blood flow as well as the functioning of the heart, brain, and lungs. The body of visualization literature at this scale is large in correspondence to the maturity of data acquisition techniques available and the ease at which these processes may be captured or simulated.

Non-human studies present an issue in this measurement-based classification system. For example, a visualization of the neural pathways in a fruit fly brain exists in micrometers, while a human brain measures in centimeters. While the focus of this survey is on human physiology, there is immense value in considering model organism physiology. These experiments tend to be more innovative, with correspondingly greater likelihood of exciting visualization opportunities. In cases where we include model organism physiology visualizations, we map the organism's scale up to the human. Following this logic, we classify, e.g., a visualization of fruit fly brain activity, at the organ scale.

A.4.2 High-Level Visualization Task

A subsequent layer categorizes the literature according to any of three high-level visual tasks: **exploration**, **analysis**, and **communication**, as illustrated in Fig. A.6. These tasks are drawn from Brehmer & Munzner's typology of abstract visualization tasks [66]. We chose high-level, rather than low-level, tasks to provide a clear picture of the broad needs and challenges users face in visualizing physiology and how this compares across scales. We first considered categorizing works according to visualization technique, e.g., direct visualization. However, since task ultimately drives the chosen visualization technique, we feel that this is a more meaningful classification mechanism that furthermore has been the basis of classification in other surveys.

Exploration tasks often arise when the user is unsure of what the data contain. In the context of the data visualization pipeline, the user typically wishes to minimally abstract the data and produce a visual mapping that is as close as possible to reality. They do this to explore what the data actually contain. This is often a preliminary step in a larger analytical process. **Analysis** tasks occur when the user may be more sure of the intrinsic characteristics of the data, but now want to extract meaning from these data. Analysis often relates closely with exploration, where a user may begin with an exploratory approach to generate a hypothesis, then perform low-level analytical tasks

alongside statistical methods to follow up on their hypothesis. In the visualization pipeline, analysis involves production of new artefacts through data transformation, derivation, and abstraction [10]. Although important for any task, audience is a key part of a **communication** task, where a visualization is created to underline key concepts of data for presentation, education, or enjoyment to a particular group, whether to peers or to a broader audience. Visualizations developed for this task are often further abstracted from the data than in analysis- or exploration-oriented tasks, and can incorporate cinematic or storytelling elements to convey the author’s interpretation of the data. While nearly all publications include figures to communicate scientific results, for this survey, we identify uses of visualization for communication beyond what is achievable with standard, out-of-the-box tools.

Many visualizations cannot be defined through strictly one of these tasks and rather are often generated to meet a combination of tasks. A work created to explore the data may also specify a visual analysis task. For example, ZigCell3D [101] visually explores simulations of cellular functions while also providing tools for the visual analysis of the underlying simulation. The same data may be visualized for a communication task if the data are more visually abstracted or if annotations or glyph overlays are added to tell a story about the underlying information. This may also summarize key findings from, e.g., a visual analysis session for a broader audience. We apply weighted categorizations to each work, excluding surveys that cover many works. This produces vector of three values between 0 and 1 for exploration, analysis, and communication task, respectively. We then use this vector to position a work in a barycentric coordinate space, which allows us to compare and contrast between similar works within and between different spatio-temporal scales.

A

We array works graphically within a triangle where each of the three points correspond to a high-level task, as shown in the left of Fig. A.6. Exploration resides at the top of the triangle to reflect that, when exploring data, we are in a position of knowing the least about what we are looking for and/or the data are in their least abstracted form. Moving clockwise to the right corner is analysis, where we usually know something about the data and what we are looking for. Communication resides at the left corner, where data are highly abstracted and summarized in order to present, communicate, or serve data for enjoyment. The set of four triangular glyphs in Fig. A.6 summarizes our visual taxonomy. Each triangle represents a single scale space, where the three triangle points represent the three respective visualization tasks. Circles indicate the position of each work as encoded by its balance of exploration, analysis, and communication tasks. Circle darkness and size dually encode the number of works with a given task categorization that we collected in our survey.

Literature Overview. The scale and task categorizations for each literature item collected for our survey can be browsed in the References section. Scale is labeled with a grayscale color tile, and related surveys are labeled with a black tile. Individual categorized works include a miniature bar graph that indicates the task(s) addressed, i.e., exploration (yellow), analysis (magenta), and communication (blue), from a range of 0 to 1. An interactive overview of the complete literature collection is available at https://lauragarrison87.github.io/star.web/vis_tool.

A.5 Molecular Function

All physiological processes depend on events that occur at the molecular scale. Molecules are the smallest units of a chemical compound and are themselves made up of atoms. Molecules in living organisms are known as biomolecules. Large biomolecules, known as macromolecules, include DNA, RNA, proteins, and lipids, while small biomolecules include metabolites [267]. Molecules are dynamic, flexible structures that interact and react with nearby molecules or ions. These individual reactions link into pathways with cascading effects at larger spatio-temporal scales.

Data. A number of data types can be used to characterize molecular function. Omics data, which is an umbrella term that includes genomics, proteomics, metabolomics, and transcriptomics data, are used experimentally to characterize and quantify molecular patterns and behaviors that scale up to the behaviors of cells, tissues, and entire organisms [154, 176, 447]. Some of these data may be utilized as structural data sources. These include nuclear magnetic resonance (NMR), x-ray crystallography, cryo-electron microscopy [267], and mass spectroscopy [325]. High resolution microscopy techniques, such as fluorescence lifetime imaging microscopy (FLIM), may also be used to visualize dynamic signaling events between proteins and their specific locations in living cells [485]. Molecular dynamics simulations commonly pair with these structural data to describe conformational changes and reactions between molecules [461].

Related Surveys. Various aspects of molecular function have received considerable attention from the visualization community with a strong focus on visual exploration and analysis tasks. The main challenge with these large, multifaceted datasets is to balance insight with complexity. The BioVisExplorer tool by Kerren et al. [248] is a useful starting point to explore the space of methods for molecular data visualization according to data type, data properties, and data tasks. Alharbi et al. [12] contribute a brief survey of surveys of molecular visualization of computational biology data, where the main focus of many of the surveys included is on either structural aspects of molecules, or on visualizing molecular dynamics and interactions from simulations and structural data [199, 267, 377]. Visual analysis tasks for molecular interactions related to molecular cavity structure and dynamics have received considerable attention [272, 466]. More recent surveys that include discussions of methods to visualize molecular dynamics and interactions of structural data include those by Schatz et al. [440] and Martinez et al. [326]. Johnson & Hertig provide a communication-oriented guide for the visualization of molecular structural data with a short discussion on visualizing molecular dynamics [239].

Surveys centered around visualizing processes from omics data from a systems biology and bioinformatics perspective similarly emphasize the challenges balancing insight with complexity for visualizing these data, often with a focus on analysis of the data with a secondary focus on exploration [154, 393, 433, 447, 484]. These works provide an overview of data types, visualization tools, and methods for large-scale omics data. Their focus is on tools and methods for molecular interactions and pathways with the goal of understanding and interpretation, generally by experts. Approaches for using multilayer network graphs to visualize omics data are explored in McGee et al. [331]. From the pharmacology domain, Csermely et al. provide a comprehensive review of analytical tools for molecular interactions, pathways, and networks for the

purpose of drug discovery [95]. Visualization approaches highlighted are limited to node-link diagrams with abstract glyph representations of molecular entities.

A number of works target the visualization of genomic data, where understanding patterns of gene expression is an important facet [370]. Nusrat et al. survey the tasks, techniques, and challenges for visualizing genomics data, of which gene expression and interactions are an aspect [374]. They emphasize the need for tools that allow for exploration for hypothesis generation and follow-up analysis. Works by Goodstadt [170], Yardimici et al. [574], and Ing-Simmons & Vaquerizas [214] highlight several visualization methods that incorporate the 3D nature of gene organization in chromatin and chromosomes into the visual analysis of gene interactions and expression, which is of particular interest to experts in recent years.

In the following subsections, we review a selection of visualizations for three categories of processes that themselves increase in temporal and spatial scale: **molecular dynamics**, **molecular interactions**, and **molecular pathways**.

A.5.1 Molecular Dynamics

Molecules are flexible and dynamic structures that frequently transition between conformational states. These structural dynamics are due to interactions between a molecule's atoms, with nearby atoms from their environment, and environmental conditions like temperature and pressure [419]. Molecular dynamics are characterized by the time scale of their conformation fluctuations (kinetics) and the amplitude and directionality of the fluctuations (structure). These fluctuations form a multidimensional energy landscape. Local fluctuations typically occur over nanoseconds, while global fluctuations can span microseconds to seconds. These global fluctuations are big conformational changes that signify protein-protein interactions, or reactions that initiate a molecular pathway, e.g., signal transduction [193]. Domain researchers are particularly interested in this energy landscape as it applies to understanding mechanisms of disease and for drug design.

Visualization approaches that target the flexibility of molecular structures often use nonphotorealistic visualization techniques that show molecular surfaces at atomic resolution. Ball-and-stick and ribbon visualization representations are also commonly used. Color is often assigned to highlight differently-flexible regions. In addition, many visual analysis methods incorporate simple graphical elements, such as glyphs or pathline visualization techniques.

Approaches that are mainly exploratory in nature are intended to allow researchers to browse and familiarize themselves with the results of a molecular dynamics simulation. These approaches tend to use minimal abstraction and encodings that are familiar to the domain [177, 304]. This also includes tools like VMD [210] or PyMol [450], which are widely adopted in the application domain. Visual elements may be used to draw out features within the data for exploration, such as pathlines to indicate atomic paths that drive changes in overall molecular shape [97].

As researchers become more familiar with the data, they may switch from exploration in an overview to analysis of a particular region of a molecule or molecular complex. Visual abstraction methods that exploit the hierarchical structure of molecules are useful to facilitate toggling between exploration and more focused identification and comparison tasks [279].




Figure A.7: MoleCollar [78] enables the visual analysis of protein tunnel dynamics and biochemistry in large ensembles of molecular dynamics simulations. Reproduced with permission.

Analytical approaches tend to incorporate interactive techniques and/or statistical methods. These allow researchers to identify and compare specific information about how parts of the molecule are moving in relation to one another. These approaches are of particular interest for researchers in drug development and protein engineering. A structural visualization of the molecule is usually important alongside 2D plots showing, e.g., trajectory [172]. For example, Fioravante et al. [138] use principle component analysis to cluster molecules that have correlated motions, while Schmidt et al. [444] derive mean shape conformations from the data to allow researchers to identify and compare metastable conformations. Other methods incorporate additional visualizations, such as time curve plots and heatmaps, to help researchers identify particular shape changes or constraints of interest [98, 500]. Interactive filtering techniques also help researchers identify particular movements of interest [209].

Conformation changes of a molecule affect not only its outer shape but also the shape of cavities or tunnels in the molecule. The shape of these tunnels affects the ability of a ligand, i.e., a signaling molecule, to travel to its binding site within a molecular cavity or tunnel [180]. These approaches usually include a mix of direct 3D visualization methods alongside heavily abstracted methods to accommodate a specific goal, e.g., to understand how the shape of a tunnel changes over time. Visual analysis methods include aggregating a molecular dynamics time sequence to a single contour plot [76]. Heatmaps to show variation in tunnel properties, such as tunnel centerline length, amino acid composition, and bottleneck size, can also be paired with direct visualization of a molecule [78, 178], as shown in Fig. A.7. More extensive visual abstraction from the original molecule shape can be used to understand dynamic structural changes and energy landscapes without occlusion [265, 271, 295].

Limited visualization research is dedicated to the communication of molecular dynamics, as much of this work comes from collaborations with domain experts with specific exploratory or analytical goals. Communication-oriented works use graphical elements, such as arrow glyphs, to illustrate molecular flexibility [73]. Tools geared towards medical illustrators, such as Molecular Maya (mMaya) [89], allow artists to animate molecular motions.

Summary. Most of the works visualizing molecular dynamics are targeted at domain experts for a combination of exploratory and analytical tasks. The time scale over which molecules change shape ranges over at least nine orders of magnitude. This presents a yet-unsolved visualization challenge to provide exploratory and analytical tools to experts to review and identify movements of interest in a vast temporal space.

A.5.2 Molecular Interactions

Molecular interactions can lead to an (ir)reversible reaction between two molecules [116]. This can change the properties of the input molecule(s), synthesize a new molecule, e.g., polymerization, or destroy a molecule. Enzymes speed up the rate of a specific chemical reaction within a cell. Ligands form a complex with another molecule, often a protein, at a binding site [180]. This binding initiates a series of reactions. The time scale of molecular interactions is large, ranging from nano- to seconds, which presents a similar visualization challenge as we discussed in Sec. A.5.1.

Similar to molecular dynamics, visual approaches to molecular interactions are strongly spatial and tend to focus on exploratory and analytical tasks for domain experts. Experts are often interested in exploring a simulation of interactions between molecules and in analyzing those interactions, e.g., protein-ligand interactions, that could lead to binding events that trigger a molecular reaction. Multi-view visualization approaches are ubiquitous, where at least one view typically uses surface models and nonphotorealistic rendering techniques to visualize the molecule(s) of interest at atomic resolution. Coloring of the molecule(s) is often according to biochemical properties or measures of uncertainty. Standard information visualization techniques, e.g., line, bar, and scatter plots accompany the spatial view to describe interaction energies and other important simulation parameters.

Key research questions relate to positional relationships between the protein and ligand, which influence the likelihood of binding. In some instances, the researcher wishes to observe such interactions in living cells, as Kerppola's work demonstrates [246]. Detailed position and interaction information from structural and simulation data can be shown on a per-atom basis through direct visualization of structural and simulation data. Glyphs and color-coding on molecular isosurfaces often enrich the visualization with additional information [499, 512].

Works that target the identification of important interaction and binding events incorporate multiple data sources, e.g., simulation and mass spectrometry [325], often in interactive views with some level of guidance. Aggregation of trajectory data to aid the analysis process is common. In addition, navigational techniques help experts locate features of interest in these often large and highly complex datasets. Such techniques can allow users to reveal different levels of detail on-demand [13], incorporate focus+context techniques, as in the CLISD view of protein-ligand interactions by Schatz et al. [439], allow filtering of subsets of trajectories [242], and allow users to jump to different parts of the simulation timeline [114]. Visual abstraction of the hierarchical structure of molecules may also be exploited for the analysis of different configurations of protein complexes [147]. Many works visualize 3D molecular structures alongside interaction energies and other important molecular parameters. These parameters may be represented by glyphs as by Hermosilla et al. [194], or using scatter- and box-plots as by Furmanova et al. [146].

Works that emphasize the broad scale of space and time over which such interaction events can occur use adjustable aggregation measures to manage spatial and temporal complexity [77, 388, 528]. Other works eschew 3D structural information entirely in favor of abstracted graphics to visualize pairwise interactions of interest [445, 524, 578].




Figure A.8: The Compressed Ligand Interaction Sequence Diagram (CLISD) [439] provides an overview of protein-ligand interactions over the course of a simulation, with unimportant residues visually de-emphasized and important residues and related parameter values given greater visual emphasis through size and color cues. Reproduced under Creative Commons CC BY license.

Tools and techniques from medical illustration and animation can be used to explore and share possible hypotheses in modeling environments, such as Autodesk Maya [363], or to communicate molecular interactions between experts or to other stakeholders. Approaches that give users tools to create molecular interactions through rule-based frameworks enable exploration and sharing of the resulting simulation data [179, 368]. Guided, interactive exploration through a rule-based simulation to track interactions in a molecular environment allows users to see the direct output of the simulation results or understand the spatial context of reaction events between molecular structures. Some of these methods employ illustrative techniques, such as focus+context, that are approachable for education and outreach [355], and incorporate multiple temporal scales into the visualization [266]. Visual complexity of molecular interactions scenes is an ongoing challenge, but research has shown that oversimplifying the crowded environments in which such interactions take place can be counterproductive to learning [232].

Summary. Exploring and analyzing molecular interactions is valuable for experts to understand and identify features and behaviors that can be used for pharmacological research. The main challenge to visualization is to continue researching effective methods that allow experts to understand the massive simulation datasets that are generated. This can be achieved via interactive tools to enable the identification of reaction events that occur very briefly within a temporal space that spans several orders of magnitude.

A.5.3 Molecular Pathways

Reactions between molecules create small changes in their immediate environment that trigger other reactions. This chain of reactions describes a molecular pathway [180]. Metabolism, signal transmission, and gene regulation and expression pathways are essential to life. Metabolic pathways describe the sequence of chemical reactions that occur in our bodies, such as the process for a cell to break down food into energy, or a pathway that builds a new molecule. Signal transduction pathways move a signal from the exterior to the interior of a cell with the help of proteins embedded in the cell surface known as receptors. Gene regulatory pathways turn genes on or off. When a gene is turned on, this allows the process of gene expression to occur, which transcribes and translates DNA instructions to create, e.g., a specific protein [217]. Pathways do not exist in isolation and interact together in larger networks.




Figure A.9: Selected visualization approaches from Gosling, a grammar-based toolkit for scalable and interactive genomics data visualization [311]. Reproduced with author permission.

Understanding the participants, sequence, and timing of molecular pathways is key to understanding physiology at larger scales. Given the complexity of the input data, visual methods to address these goals tend to target expert user exploration and analysis tasks. These incorporate varying degrees of abstraction and interactivity. 2D information visualization techniques dominate, with networks being the most common technique, to show a sequence of steps in a pathway. Heatmaps, line plots, chord diagrams, and histograms are common for the visualization of gene expression.

The most straightforward visual methods allow experts to explore and identify the sequence of actors that participate in a given pathway(s) use node-link diagrams. Perhaps one of the most well-known pathway exploration tools, Cytoscape [460], uses node-link diagrams to visualize complex pathways and networks for users to explore and query. Such diagrams show entities in highly abstracted glyphs, often indicate reaction direction, and can indicate the location where the pathway takes place [243, 290]. Brushing and linking [176], filtering [299], comparison [448], and focus+context [237] techniques for detailed analysis are often supported. Numerous works have explored different layout algorithms to reduce crossover and clutter of these complex and often crowded visualizations [62], and implement graphical representations using, e.g., a subway map metaphor [297], that are approachable to broader audiences [81, 254]. Visualizations of pathway simulations can be abstracted in 2D as line charts or heatmaps [458] to help experts to better understand the timing of pathways. An entirely different pathway simulation approach by Le Muzic et al. [356] employs an agent-based approach with 3D molecular structures to tell a multi-temporal scale story that provides insights to both experts and broader audiences alike.

Identifying and comparing levels of gene expression can provide valuable information to researchers on the activity of a given pathway, while studying gene co-expression can provide understanding of patterns and similarity of certain expression pathways. Gene expression and co-expression data are most commonly displayed in heatmaps, parallel coordinates, and chord diagrams, as shown in Fig. A.9 [311]. Tools like Caleydo [300] enable the exploration and analysis of large-scale pathway data alongside gene expression data, using node-link diagrams and heatmaps in a 2.5D layout, while OmicsTide [183] uses clustering with profile plots in a Sankey diagram to compare trends from gene expression and proteomic data. Some tools capture the multiscale nature of gene expression in visualizations that span the scale of individual nucleotides to entire chromosomes [340]. Gene expression is a dynamic and fluctuating process. Other tools allow for exploration and analysis of temporal patterns of

these fluctuations [341], and in some cases use clustering methods to facilitate pattern identification [94].

Researchers are similarly interested in identifying and comparing concentrations of metabolites in specific locations of the body. This information provides another perspective on the activity of certain pathways. Tools for visual exploration and analysis of metabolite concentrations are useful to understand metabolic profiles of diseases at a molecular level [372]. Such approaches can use basic statistical methods alongside heatmaps [149], violin [228], or star charts [229].

Strong communication-oriented approaches to visualizing molecular pathways often draw inspiration from medical illustration and use cinematic elements to convey pathway information, such as Berry's animations showing the process of DNA transcription in real-time [45]. In this way, the molecular dynamics and reactions between molecules at key steps in the pathway can be visualized in a larger context. Large charts showing pathway elements, when mainly used for communication, usually rely on abstraction of visual elements to create a scene that balances accuracy with readability [74, 148].

Summary. Similar to molecular dynamics and interactions, the majority of visualization research works focus on expert-centered exploration and analysis tasks. The extraordinary complexity and volume of these data often necessitate guidance in interactive methods, and many approaches use statistical methods to reduce the analysis space alongside minimalist graphical elements. Further research into methods that facilitate a greater degree of exploration for hypothesis generation of these data, while managing the volume of information present, is an ongoing visualization challenge and opportunity for all molecular processes. Visual communication research for pathways is also important to develop further. Giving the public better tools to understand how diseases work, such as in COVID-19, can improve adherence and trust in public health protocols. Understanding physiology at this scale is essential, as these molecular dynamics, interactions, and pathways work in concert to trigger behavioral and physical responses that form the foundation of cell physiology.

A.6 Cellular Function

The cell is the structural and functional unit of life in humans and many other organisms. Cells are self-contained, bounded by an outer membrane holding several substructures (organelles) that perform specific functions and facilitate molecular pathways that keep the cell alive and within balance [491]. We acknowledge that the distinction between cell scale processes and molecular scale pathways can be blurry, particularly in the case of large-scale molecular networks that themselves define cell physiology. We categorized each work according to the scale that is most relevant to the user's interest. In cases where interest is primarily in understanding whole-cell behaviors, we categorized the work in this scale, while if user interest is primarily in the various molecules that form a pathway or network, we categorized corresponding works in the molecular section.

Data. Input data to visualize cellular function can be acquired experimentally, often through different time-lapse optical microscopy methods on living cells and most commonly through fluorescence microscopy. This technique allows researchers to tag cells

with specific proteins that fluoresce under the microscope, allowing for visualization of specific cellular structures and behaviors. For a complete overview of live cell microscopy methods, we refer to Jensen et al. [234]. Electron microscopy, which kills the cell, is often used to supplement live microscopy methods to visualize ultrastructure details inside the cell [160]. Biomechanical methods to experimentally determine the effects of different forces on cells and their organelles include atomic force microscopy and tweezing [38, 200]. Omics data, e.g., single-cell RNA sequencing (scRNA-seq) data, which provides the molecular expression profiles of live individual cells, can also provide detailed information on cell function and behavior [358]. Because a cell is a self-sufficient entity, it is often a natural starting point for physiological models of cell behavior [472]. The CellML repository, maintained by the Human Physiome Project, is a rich repository for cell behavioral models [309]. Stochastic simulations are also useful to simulate complex biological pathways and networks within the crowded and dynamic environment of the cell and its surroundings [532].

Related Surveys. Surveys covering the visualization of cell dynamics and interactions are sparse relative to the molecular scale. Pretorius et al. [406] identify six classes of visualization techniques: spatial embedding, space-time cubes, temporal plots, aggregate plots, dimension reduction, and lineage diagrams in their survey of visualization for live cell imaging. These techniques remain common in our report at this scale. Goodsell et al. [169] provide a review of visualization methods that combine experimental data from microscopy, structural biology, and bioinformatics to build structural models of entire cells, mainly through nonphotorealistic visualization techniques. These models include details of molecular behaviors and interactions that contribute to cell dynamics. Feig & Sugita review models for visualizing whole-cell dynamics at the resolution of the myriad molecular interactions that occur within a cellular environment [132, 133]. Their work highlights the use of surface, ribbon, and ball-and-stick molecular models at atomic resolution.

In the following, we discuss visualization trends and challenges for **cellular dynamics**, which essentially are processes that affect the cell itself, and **cellular interactions**, which are processes that involve a cell interacting with its neighbors.

A.6.1 Cellular Dynamics

The dynamics of a cell are dictated by molecular pathways and by behaviors of its organelles, which are themselves modulated by molecular pathways. These pathways drive the dynamics of a cell's organelles, the ability of a cell to move in its environment, the suite of internal mechanisms that dictate a cell's growth, and that lead to cell division and death, to name a few processes. We also discuss visualizations for whole cell models.

Organelles participate in and facilitate the network of pathways that drive the overall behavior of a cell. Visualization tasks related to organelles are often exploratory in nature, e.g., to observe the effects of an experimental condition under microscopy. Visualization methods from the domain often show time-lapse imaging data unmodified or surface renderings. This can help users to understand the shape changes a cell nucleus undergoes in response to experimental conditions [449], the movement of cellular vesicles [414], or the compaction of chromatin in the nucleus over different phases of




Figure A.10: Visualization of vesicle formation from molecular dynamic simulation data [398]. Reproduced under Creative Commons CC BY license.

the cell cycle [386]. Analysis-oriented approaches often color-code regions of interest on volume-rendered segmented data or raw data slices to identify and compare features that indicate the functioning of underlying pathways [160, 236]. Glyphs are used to annotate features of interest on imaging data as a process occurs [552], and heatmaps can quantify interactions between organelles over the course of an experiment [515].

Simulations using 3D surface models can help answer questions about how organelle structures move and behave. For example, Waltemate et al. [537] visualize membrane dynamics at molecular resolution in small “patches” of the cell membrane. More recent works visualize microtubule dynamics [256] or dynamics of mitochondria and cell transport vesicles [398], the latter of which is shown in Fig. A.10. Such visualizations are adaptable for use in education environments, with systems like LifeBrush designed to explore mesoscale environments, e.g., the mitochondrial membrane, at molecular resolution in VR [99]. Even further toward communication are hand-crafted animations, such as the ground-breaking *Inner Life of the Cell* [54], which shows the interplay between various organelles and molecules within the cell using cinematic techniques and visual abstraction to focus the narrative.

The individual dynamics and interactions of organelles influence and facilitate the cell’s response to input from its environment and internal mechanisms that push the cell through its life cycle of growth, division, and death. Visualization of cell movements in 2D or 3D to discover and understand behaviors under experimental conditions is common in the domain, e.g., to understand cell cycle progression [196, 431].

Researchers may be interested in localizing subcellular structures and interactions through direct observation of imaging data [26] or may supplement imaging output with histograms, time plots, and similar aggregate visualizations to quantify features of interest [569]. Approaches may also use such aggregate visualizations to visualize results of classifications based on live microscopy alongside gene and protein expression data [35]. Dimensionality reduction methods, such as tSNE and HSNE, are useful methods to characterize and compare different cell behaviors and types from high-

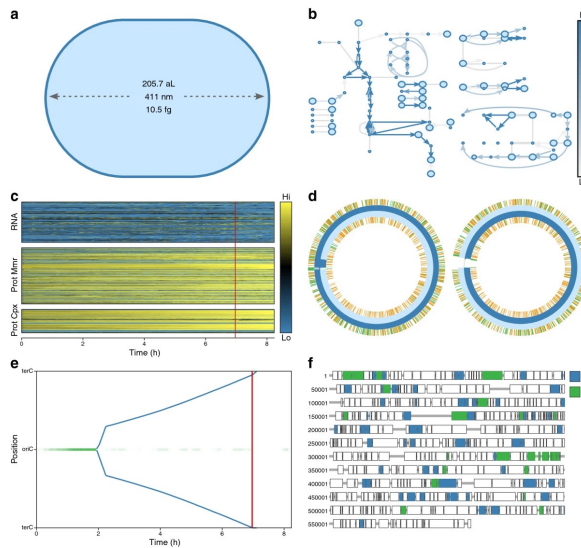


Figure A.11: WholeCellViz [292] modeling framework to explore and analyze simulations of cellular dynamics. Reproduced under Creative Commons CC BY license.

volume experimental data. Visualizations can map the results of these methods to a scatter plot, with color encoded to cell type according to an expression marker [203], and include guidance methods to more easily explore and analyze the data [204]. Others deal with these high-volume data by identifying exemplar cells, either automatically or through user specification, to track specific derived attributes, such as growth, in response to drug treatments for cancer research in a multi-view visualization [281]. Simulations allow researchers to manipulate cellular parameters for exploration and analysis. Simulations of parts of the cell cycle can be queried and explored for hypothesis generation in tools like CellCycleBrowser [58], which uses an interactive multi-view visualization containing heatmaps, line, and scatter plots to show the results of parameter changes to the system.

Whole-cell physiology is naturally multiscale, with limited works addressing visualization and specific user tasks at both molecular and cellular scales. These works enable experts to better understand intracellular functional and structural relationships. Highly abstracted approaches, similar to node-link network diagrams used to represent molecular pathways, can be used to visualize the multiscale interactions that occur within the cell [410] that can be explored or queried. Visualizations of whole-cell simulations in 3D are useful to put molecular pathways into context, such as the effects of signal transduction on the cell's function [125, 126] or the conditions and events that lead to cell death [124, 127, 446]. WholeCellViz [292] and ZigCell3D [101], the former of which is shown in Fig. A.11, are whole-cell modeling frameworks. These frameworks allow researchers to explore and analyze cellular simulations in a biological context, from the molecular scale to the entire cell. They include pathway information as maps, as well as animation. ZigCell3D also incorporates imaging data and 3D models. A recent structural model of a whole *Mycoplasma* cell [322] provides an

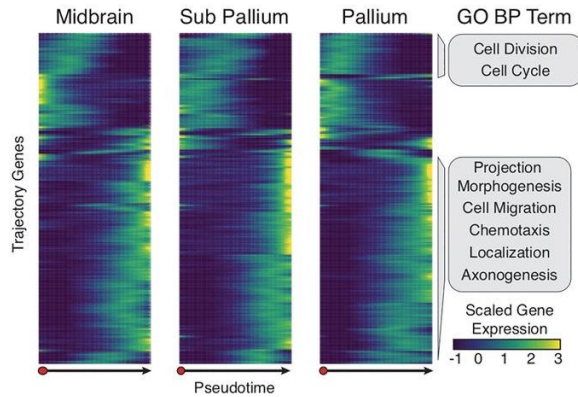


Figure A.12: sci-Space measures patterns of gene expression over time to understand patterns of cell differentiation and migration in neural tissue [474] Reproduced with permission.

unprecedented means for researchers and the general public to explore and understand the structural and functional relationships of entities within the cell.

Summary. The visualization of cellular dynamics puts molecular pathway information into a cellular context and enables understanding of overall cell behavior. Experts are often interested in exploring and quantifying these data directly from imaging methods, with analysis of key features in aggregated plots. Further research into more interactive methods to facilitate analysis that allow experts to move away from simple rendering of microscopy data is a possible direction to explore. Very few works, especially from the visualization community, support expert study of organelle dynamics and behavior. This is an open space for visualization research. Multiscale visualization becomes truly meaningful at this scale to connect molecules with cellular behaviors. While numerous methods allow for exploration of whole-cell physiology, analysis of such models remains relatively limited, and this is another future research opportunity. Finally, in some contexts, communication-oriented approaches can serve both experts and a broader audience equally, as cells are less conceptually-abstract entities than molecules. Research into such approaches, particularly with regards to public health and in facilitating conversations on the mechanisms of disease, is an exciting challenge.

A.6.2 Cellular Interactions

In reality, cells do not exist in isolation. Their physiology is strongly influenced by their interactions with their environment and neighboring cells. In this section, we discuss works that focus on the behavior and fate of individual cells, where understanding the environment and neighboring cell interactions are key to the user task.

As in previous topics, many visualization works are born out of collaborations with domain and mainly address exploratory and analytical tasks. These methods allow researchers to browse experimental, imaging, or simulation data to understand cell communication, lineages, and migratory patterns.

Direct visualization of live cell imaging data provides an overview of cell division, adhesion, signaling, and movement patterns within different environments, e.g., tumor

microenvironments [118]. Clustering methods can facilitate user exploration of inter-cellular communication networks from single-cell transcriptional data. COMUNET classifies and clusters cell types according to ligand-receptor pairs and visualizes these communication patterns in node-link diagrams [467]. Migration patterns of differentiating cells can be understood by visualizing, e.g., spatial transcriptomics data, as in sci-Space [474] (Fig. A.12), where heatmaps of gene expression patterns are measured over pseudotime to capture cell differentiation and migration patterns in a tissue context. Clustering methods can classify cells according to their migratory and other behaviors from microscopy data, which is valuable for comparative analysis [130]. Simulations with simplified 3D spherical models to represent individual cells provide information and control on per-cell properties of division, adhesion, and other environmental variables in an *in vitro* environment [156, 389].

Cell lineages contain valuable information on patterns of cell division, growth, differentiation, and death over generations of cells. This is particularly important with stem cells, which have unique regenerative abilities that have massive implications in cancer and other areas of medical research. Statistical methods to make sense of these patterns, in combination with the branching tree structures of lineage diagrams, help researchers identify and compare factors that influence cellular genealogies [162, 405]. These approaches can include visualization of cells within a spatial context, with navigational tools to observe how cells divide and where they migrate, e.g., to neighboring blood vessels [432, 534]. Uncertainty due to segmentation of microscopy data when tracking cell aggregation is a challenge. Tools like Uncertainty Footprint [539] attempt to visualize and quantify these uncertainties for domain experts.

A growing interest in public science education has led to the development of tools like Bioty [544], a real-time programming environment that visualizes cell interactions for non-experts. As for other topics, hand-crafted medical illustrations are used to educate audiences on cellular interaction processes, such as the communication between a neuron and muscle cell [166].

Summary. Visualizing cellular interactions adds a degree of complexity to cellular function visualizations, as these leave the self-contained environment of the cell to include external parameters that increase the complexity of the system. Collaborations with experts provide a means to explore and analyze data acquired either experimentally or through simulation, where gaining an understanding of the data through exploration is equally important to more targeted analysis tasks. Developing methods to facilitate exploration and analysis of cellular microscopy and lineage information through visual abstraction while retaining expert trust is one research opportunity. We found few works visualizing cellular migration and adhesion, particularly from within the visualization community. Given the importance of these behaviors in normal development and disease, this is yet another research opportunity.

A.7 Tissue Function

At the tissue scale, we see groups of cells of the same type that perform a specific function. These form tissue, which allows for coordinated behaviors to accomplish tasks impossible for single cells to perform. A tissue region also includes a container, known as an extracellular matrix, that holds the cells together and provides structural stabil-

ity [472]. Each of the four main tissue types in the human body serves a specific role: (1) epithelial: covers and lines body surface and cavities; (2) connective: protects and supports body structures, i.e., organs; (3) muscle: coordinates movement; and (4) nervous: facilitates communication of nerve cells through electrical signaling. The visualization works we discuss in this section aid user tasks where the goal is to understand the overall behavior or dynamics of cells in aggregate, rather than individual cells. One process that we highlight includes tissue growth, also known as morphogenesis. This process drives the development of, e.g., blood vessels or tumors. We also discuss methods for visualizing perfusion, the delivery of nutrients to tissue via small blood vessels called capillaries, and the propagation of electrical signals through tissues.

Data. Spatially-resolved gene expression data can characterize the overarching physiology and behavior of tissue [358]. These methods may pair with imaging methods, as in seqFISH+. For a comprehensive discussion of specific experimental methodologies, we refer to Waylen et al. [545]. Imaging methods for perfusion are well-established. These include positron emission tomography (PET), single-photon emission computed tomography (SPECT), computed tomography (CT), Doppler ultrasound, dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI), dynamic susceptibility contrast MRI (DSC-MRI), phase-contrast MRI (PC-MRI), ASL (arterial spin labeling), and optical methods such as widefield or fluorescence microscopy. Conventional widefield microscopy is common for visualizing tissue histology, although this requires fixing cells to a slide that not only kills the cells but can damage their spatial organization. Simulations and models often describe signal propagation within nervous or cardiac muscle tissue. Imaging data at the resolution to visualize individual neurons involved in signal propagation typically come from microscopy. The state-of-the-art techniques for imaging brain tissue are confocal laser point-scanning (CLSM) and spectral precision distance microscopy (SPDM) for their high resolution, improved signal-to-noise ratio, and removal of out-of-field fluorescence. For further reading on these techniques, we refer to Tröger et al. [507].

Related Surveys. Visualization works that cover tissue-scale physiology are limited and generally motivated by the needs of experts in the medical domain. Preim et al. [403] survey methods for the visual exploration and analysis of perfusion data. The authors highlight cine-movies, subtraction images, and color-coded parameter maps on a single slice as basic visualization techniques. Advanced visualization techniques covered include multiparameter visualizations, e.g., colored height fields, combining structural information with dynamic perfusion data, or extracted features, e.g., temporal curves. Schlachter et al. [441] survey visual computing methods for radiotherapy planning, which include detailed visual analysis of the metabolic profiles of tumors that can be acquired through perfusion data. Volume visualization techniques that fuse multiple data sources into a single image through overlays and color-coding are common in this area, particularly within the application domain. More advanced visualization techniques enable exploration and analysis of uncertainties in segmentation, or analysis of perfusion parameters in parallel coordinates, scatter, and star plots, among others. Qutub et al. [411] review modeling efforts for angiogenesis from an application domain perspective, some of which include molecular and cellular-level processes in the resulting visualization that we discuss further in Sec. A.9. Visualization techniques used to illustrate these models include node-link diagrams, line plots, and histograms to




Figure A.13: Compartment diagram and parameters for tumor growth [313]. Reproduced under Creative Commons CC BY license.

describe and analyze model parameters. Spatial visualizations include algorithmically-generated surface models and stacked image slices. Color maps to model parameters, e.g., tissue oxygenation or upregulation of a particular pathway.

In the following, we provide an overview of visualization approaches and challenges for **tissue dynamics**, or the behavior of aggregates of cells in a given tissue type. We then discuss **tissue interactions**, including the delivery of nutrients and the passage of electrical signals through specialized tissue.

A

A.7.1 Tissue Dynamics

Tissue dynamics refers to the behavior of aggregates of cells in a given tissue type. Visualizing tissue dynamics simulations allows researchers to explore the general mechanisms of tissue growth and development under changing environmental conditions. Understanding spatial relationships at this scale is particularly important, and many visualizations represent simulation data as 3D surface models to capture the development of blood vessels (angiogenesis) [411], liver tissue regeneration [201], embryonic limb [88] and organ development [86, 115], and in a non-human case, wing development [65]. Exploring changes in skin tissue as a response to aging in 3D is also of interest [212], with the ability to change parameters to simulate the impact of disease or dehydration. In addition to visualizing normal processes, simulations are valuable sources to visualize pathologies related to tumor growth under changing environmental influences [501, 502].

Although we have discussed gene expression data previously for the analysis of cellular interactions, this family of methods is useful for tissue-level visual exploration and to identify biomarkers in cell aggregates when performed *in situ* and paired with imaging data. These cell aggregates are often identified through clustering methods [468], and can then provide molecular and cellular resolution maps of the body, e.g., of embryonic tissue development in the first trimester [41]. Subsequent exploration of such tissue maps provides an opportunity to discover emergent properties at this scale. Mul-

teesum [339] exemplifies this interplay between exploration and analysis, where comparing similar expression profiles of aggregates of cells allows researchers to form hypotheses about gene relationships and location. Numerous domain approaches also use visualization mainly to confirm hypotheses, e.g., the direct visualization of digital histology slide data to quantify the progression of liver tissue damage in fibrosis [306].

Communication-oriented visualizations of tissue dynamics may take the form of adjustable simulations with easy-to-use interfaces and simple graphics that appeal to both researchers and a broader audience [563]. Animation of 3D models is also useful as an educational tool for showing the process of organ development, e.g., of the developing heart [442]. Lastly, hand-crafted illustrations that describe models of tissue growth, as shown in Fig. A.13, are invaluable to clearly and succinctly share models with peers [313].

Summary. Visualization of tissue dynamics is often geared first towards exploration to familiarize oneself with the data, as data at this scale are typically complex and high-dimensional. Comparison tasks between groups are then common, where experts wish to identify parameters or biomarkers that define certain tissue behaviors or functions. Tissue dynamics are challenging to visualize *in vivo*, with approaches often using underlying processes such as gene expression or the presence of other biomarkers to characterize tissue functional properties. Simulations provide the means for visualizing truly dynamic growth processes in healthy and disease conditions. However, they often are abstracted from reality, with visualizations that expose only the final part of a multiscale story that is rooted in the molecular scale and with limited interactivity. Devising methods to enable fully interactive exploration and analysis of tissue dynamics, whether purely at the tissue scale or extending across scales, is a grand challenge and opportunity in visualization research.

A.7.2 Tissue Interactions

In this section, we focus on the interactions between different tissue types that allow for the passage and exchange of nutrients, as in tissue perfusion, or for the passage of electrical signals, as in signal propagation.

The function of blood flow on the microscopic scale is to supply, or perfuse, tissues in the body with oxygen, nutrients, and hormones and to transport waste products away into the appropriate "recycling" centers such as the lungs, kidneys, and liver. Different tissues have different perfusion rates, and visualization can be a powerful tool in profiling tissues based on these data. Perfusion data are particularly useful in identifying the extent, composition, and metabolic profiles of tumors.

Most use cases to visualize **tissue perfusion** are highly clinically-motivated with a particular set of analysis questions already in mind, although many methods incorporate a degree of exploration. These approaches incorporate structural visualizations of tumors and the surrounding tissue to provide context, and use derived multi-parametric imaging data to classify and visualize key physiological parameters. Approaches use simple color overlays with parameters mapped to color channels [22] to quickly quantify values, or allow for user interaction and exploration that incorporate time intensity curves [161, 319, 382], radar plots [347] as shown in Fig. A.14, and scatter plots with




Figure A.14: Paraglyder is a tool for the visual analysis of tumor metabolic profiles [347]. Image provided by the authors and reproduced with permission.

glyphs encoding further information [348] that allow experts to identify and compare features of tumor physiology.

In the human central nervous system, information is processed by signal transmission and propagation between neurons, where the extracellular space plays an important role in transmission and signal propagation.

In visualizing **signal propagation**, particularly in simulations, experts wish to understand the mechanism, path, and timing of these propagation events. Straightforward visualizations that plot spikes in signal propagation are relatively common in the domain literature, such as in Rhodes et al. [416]. Microscopy data often provide a structural foundation for visualizing simulations of signal propagation between neurons or in a multi-neuron network [68, 283]. We show an example of a multineuron simulation network from BioDynaMo in Fig. A.15 that is realized through procedurally-generated surface models. Dimensionality reduction methods can facilitate exploratory visual analysis of signal data, and allow users to identify patterns that signify, e.g., key points of a behavioral task [69]. Abstracted 2D plots, such as L-plots proposed by Dunin-Barkowski et al. [113], allow experts to observe and compare neural signaling patterns. Signal propagation is influenced by several factors, e.g., the distribution and density of glycogen around a synapse, which is the space where two neurons meet. Ab-




Figure A.15: Agent-based simulation of large-scale pyramidal neuron cell growth with the BioDynaMo platform [68]. Reproduced with author permission and under Creative Commons CC BY license.

stractocyte [345] is mainly designed for the visual analysis of astrocyte structure and distribution around neurons, but its pipeline includes glycogen distribution analysis.

Simulations of signal propagation are not limited to neural tissue. Visualizing simulations of the electrical conduction system in cardiac muscle tissue is of interest for experts to understand the timing and rate of signal propagation in different phases of the cardiac cycle [36, 181].

Summary. Unlike perfusion of tissue, the majority of work we found for spatially visualizing signal propagation comes from simulation data. As hardware and software continue to advance and support more complex simulations, there will be a corresponding increasing need to provide visual methods to explore, analyze, and communicate these data to various stakeholders. In tissue perfusion, visualization research that further supports exploratory analysis to identify complex biomarkers is an ongoing challenge. Visualization approaches for both tissue perfusion and signal transduction tend to have strong analytical components, especially in the case of perfusion, where clinical diagnostic improvements are the driving need for these applications. Communication-oriented research works are limited at this scale. The clinical motivation for understanding physiology becomes even more apparent at the organ scale, which deals with the interplay between different tissue types on a larger scale.

A.8 Organ Function

The last scale that we review is that of visualization for organ-scale processes. Organs are discrete units in the body that perform a function or a group of functions [472]. In the following, we cover visualization for four areas that are well-known in medical visualization: the dynamic properties of **blood flow** (hemodynamics), and the **functioning of the heart, lungs, and brain**. We also include a brief discussion of other visualized organ functions to give a sense of further opportunities at this scale, e.g., skeletal muscle function.

Data. Typical data inputs for visualization of physiological data at this resolution include a host of imaging modalities that can be time-resolved alongside, or separate to, simulation data. In many instances, structural data provide context to the visualized dynamic process. Information about anatomical structures can also be based on measures of physiology. For example, while diffusion-weighted and diffusion tensor imaging (DWI/DTI) methods use the diffusion of water molecules to capture the fiber architecture of the brain, these data are primarily used to probe white matter microstructure [399]. As such, the visualization of DWI/DTI data alone is out of the main scope of this work. The same is true for other such structurally-focused modalities. 4D computed tomography (CT) and a range of MRI sequences, e.g., dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) and phase-contrast MRI (PC-MRI) [260], are frequently used to assess various organ functions. Specialized ultrasound (US) methods are adapted to capture particular processes, e.g., Doppler ultrasound for hemodynamics, or electrocardiography [567] and echocardiography [317] for heart function. Computed tomography angiography is useful in assessing hemodynamics and heart function [334]. For a discussion of the strengths and weaknesses of these modalities in hemodynamics imaging, a hot topic in visualization research, we refer to Jennings et al. [233], Markl et al. [323], and Sengupta et al. [455]. Typical

imaging modalities measuring brain function include electroencephalography (EEG), functional magnetic resonance imaging (fMRI), and PET. We refer to Pfister et al. [399] for a detailed discussion of each of these modalities. Electromyography (EMG) [44] and motion capture data are common sources for assessing muscle function.

A number of these modalities alone have insufficient spatial and/or temporal resolution to capture an organ process of interest. These are often augmented with simulations, or simulations are developed from these imaging data. The purpose of simulation can also be to correct issues with the acquisition, such as motion-related artifacts. Approaches include computational fluid dynamics (CFD) for blood flow [420], statistical heart and lung motion models, and large-scale simulations of signal propagation for brain function.

Related Surveys. Surveys that discuss aspects of organ physiology on a broad level are typically motivated by the medical domain. Preim et al. survey the use of medical animations for organ-level processes [401] that tend to focus on communication-oriented tasks. Birkeland et al. survey works that fit in the ultrasound visualization pipeline [50], where the end-user task often is to explore the data and to identify specific features within the data related to, e.g., blood flow and heart function. Many visualization approaches at this scale combine modalities to overcome individual modality limitations. Lawonn et al. [287] provide an extensive discussion on multimodal visualization. Tory et al. provide a brief overview of methods for MRI in combination with dynamic SPECT data [506]. In general, visualization tasks at this scale focus on giving experts, whether in medical research or more directly in the clinic, tools to explore and analyze physiological features for improved diagnosis and treatment.

A

A.8.1 Blood Flow

While we previously looked at blood flow from the lens of how it supplies nutrients to tissues (Sec. A.7.2), researchers often are interested in the dynamics of blood itself as it travels through the heart and vessels of the body. Understanding patterns of blood flow can help researchers and clinicians make better decisions about patient health, such as when to operate on an aneurysm.

Related Surveys. This is a mature area with several surveys and state-of-the-art reports available. For further details on visualization techniques and challenges for this topic, we refer to reports by Markl et al. [324], van Pelt et al. [521], Vilanova et al. [529], and Stankovic et al. [477]. For further reading on visualization methods specific to PC-MRI blood flow, see Köhler et al. [260]. Most recently, Oeltze-Jafra et al. [384] survey trends and challenges in visualizing medical flow data, where the primary focus is on blood flow data. These surveys highlight a mix of exploratory and analytical visualization tasks, where tasks are highly motivated by domain experts' needs to locate and identify flow features that impact patient health. General flow visualization techniques are commonly used, e.g., glyphs, textures, integral curves, line integral convolution (LIC), colored cut planes, extraction to surface models, e.g., streamlines. Contextualization of blood flow dynamics using image slice or volume rendering of the surrounding anatomy is key in nearly all blood flow visualization scenarios. Advanced visualization techniques often incorporate multiple interactive views with facilities for validation, filtering of key parameters, and uncertainty analysis.




Figure A.16: Method of aorta straightening for occlusion-free comparison of flow over time [20]. Reproduced with author permission.

Exploratory tasks generally aim at obtaining an overview of flow patterns and often precede a quantitative workflow. These may visualize simulation data [371, 420], use techniques that combine different modalities to create the visualization [53, 139, 397], or use single imaging methods [71]. The containing structure of interest, e.g., a vessel or the heart, is useful to preserve for context. The paths and direction of flow data can be presented as pathlines [102], streamlines [49], arrows [206], and pathlets [21], which can also be used for analysis. Newer approaches have experimented with effects like smoke or dye to better visualize time-varying flow patterns with indications of uncertainty [104]. Interactive elements, such as virtual probes [519], aid in expert exploration of flow patterns prior to quantitative analysis.

Although illustrative techniques are often associated with communication-oriented tasks, most illustrative techniques employed for this topic are aimed at experts to facilitate exploration of flow data. These approaches can, e.g., reduce occlusion from vessel walls through adaptive surface visualizations [152, 286] or focus+context flow lens treatment [151]. Other illustrative approaches facilitate exploration of wall thickness relative to flow properties [285].

Many analytical approaches from the domain are limited in interactivity. These approaches again include a structural context, with similar visual representations for flow data as for general exploration. Visual representations also often use heatmap overlays and/or glyphs of the same styles mentioned for exploration, e.g., streamlines, arrows, to indicate flow velocities at particular points. This allows for feature quantification directly from imaging data [206, 492]. Experts also often wish to quantify and compare flow rates in simulations relative to time-resolved imaging data [282].

Analytically-focused tools and methods developed from collaborations between domain experts and the visualization community often take more experimental or abstracted approaches to aid analysis tasks. For example, Angelelli et al. [20] flatten 3D tubular flow to 2D to compare flow patterns over time, as shown in Fig. A.16. Semi-automatic classification and clustering methods are also common to aid expert identification and comparison of vortices and shapes in the data [123, 335, 381, 383]. Interactive linked views, particularly when including simulated and acquired data, can help experts better evaluate hemodynamic patterns and model accuracy [296]. Multi-view visual analysis tools often rely on an interplay between exploration and analysis for users to get a sense of the data and browse interesting regions before identifying and comparing features to understand, e.g., aneurysm rupture risk [337, 338].

Approaches from the research community to present blood flow for education for a more general audience are limited. These methods may hint at the original flow data through animation but instantiate red blood cells to indicate to a broader audience what the flow represents [148], or employ more fanciful metaphors to show the passage of blood in a cardiac cycle [91].

Summary. Despite the extensive work on hemodynamics, further challenges and opportunities remain. The ultimate aim of many of these works is to make visualization of hemodynamics available in a clinical setting to aid in rapid and accurate identification of life-threatening flow behaviors. Studies that assess the possibility of real adoption of these techniques in clinical routine are an interesting avenue to explore. Communicating these data then to patients, in a way that is both understandable and minimally-alarming, is essential and remains an open challenge in visualization.

A.8.2 Heart Function

Heart function is well-characterized in physiology and visualization research. In this section, we focus on visualization related to the (1) mechanics of the heart as a pump and (2) the cardiac conduction system of the heart, which is an electrical network that controls heart rate and rhythm [180]. Diseases related to these aspects of heart function include (1) myocardial ischemia, where the heart tissue does not receive enough blood from its supplying arteries, (2) heart failure, where the heart is unable to pump blood effectively, and (3) atrial fibrillation, a dysfunction of the cardiac conduction system that leads to irregular heart rate and rhythm. These diseases provide strong clinical motivation and drive many of the visualization use cases in this topic.

Related Surveys. Nazir et al. survey the visualization of various aspects of heart function from the medical domain, focusing mainly on analysis for use clinical routine [364]. Walton et al. [538] provide a broad overview of the methods and challenges in visualizing cardiovascular magnetic resonance imagery for clinical research before presenting a prototype approach for visualizing this type of data. Generation of surface models and volume renderings of the heart, paired with time-lapse video to describe deformation, are key visualization techniques for this topic. Heatmap visualizations commonly indicate parameters of interest, and the bull's eye plot is ubiquitous for the visual analysis of perfusion data to understand heart function.

Exploration-centered visualizations provide an overview to experts of general features and parameters related to shape changes, e.g., for specific chambers, valves, or the entire heart, in a cardiac cycle. These works visualize phases of the cardiac cycle from simulations on surface models (which typically are abstracted from acquired data), such as the LFX Virtual Cardiac Model [238] or the constrained Multi-linear Shape Model [218]. Patient-specific approaches visualize models in combination with acquired data [29, 562], or only acquired data. Some works build predictive models, and preserve links between the simulation and the original data to understand the mapping procedure [486]. Although we discuss blood flow extensively in Sec. A.8.1, for completeness, we note that several approaches visualize patterns of blood flow to explore questions related to heart function, e.g., Kulp et al. [276].

Approaches geared towards a combination of exploration and analysis, or focused more purely on analysis, often favor colormaps applied to mesh or imaging data to

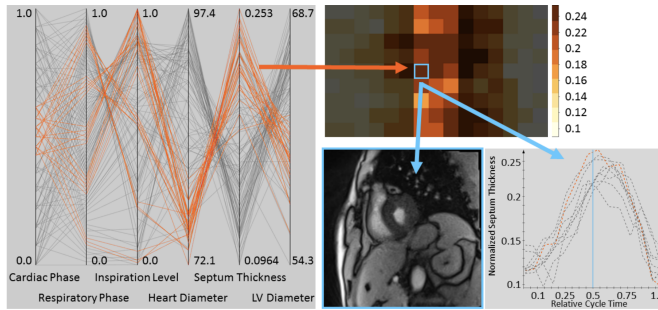


Figure A.17: Visual analysis for quantification of septum motion from real-time cardiac MRI [494]. Reproduced with permission. © 2017 The Author(s). Eurographics Proceedings © 2017 The Eurographics Association.

quantify parameters of interest. Rainbow colormaps are common, especially in the medical domain. These works often use multiple interactive views to link simulation, imaging, and derived statistical information. The main purpose of the visualization is to evaluate particular parameters, e.g., strain rate [188, 391]. Evaluating the movement of particular landmarks can be aided by additional plots, such as parallel coordinates and heatmaps, to compare different motion parameters [494], as shown in Fig. A.17.

Perfusion data, which we visited previously in Sec. A.7.2, are often used as a basis to determine heart functionality. Most use cases are tied to experts with a need to identify and quantify particular tissue properties in the context of a pathology, and while exploration is a component, it is usually not the main task. Many approaches incorporate a bull's eye plot in cases related to myocardial ischemia, where heart tissue does not receive the blood and nutrients it needs to function. This visualization is familiar to clinicians to quantify the extent of the damage to the heart tissue, often alongside structural representations of the coronary arteries [380, 425, 496, 497]. The bull's eye plot can be further adapted for targeted visual analysis of the motion of the left ventricle of the heart over time [462]. Approaches are often interactive to provide an exploratory element for the user. Glyphs can be incorporated, as by Meyer-Spradow et al. [342], to quantify local tissue perfusion across the whole heart. Statistical, e.g., PCA, and aggregate measures may be used to reduce the complexity of the data [379] and include representations of uncertainty [425].

Visualization is also used to evaluate the accuracy of simulations against acquired data, where heatmaps [334] or juxtaposed line plots [182] indicate shape prediction accuracy. Other methods focus heavily on patient-centered care and outcomes [566], such as simulations of surgical procedures, e.g., mitral valve clipping, with quantitative evaluation with heatmaps to help predict patient outcomes [317].

Some heart simulations have been developed, not only for expert exploration, but for use in education and surgical training. These include Dayan et al.'s 3D animation of the dynamics of a simulated mitral valve [100] and the virtual reality (VR) simulation of radio frequency ablation by Pernod et al. [396]. Pernod et al.'s approach uses a heatmap to show membrane propagation potential, a tissue-scale process, on the heart surface. VR has also been used for patients in a biofeedback scenario to manage stress [171].




Figure A.18: A coupled model of tissue deformation and network airflow that enables predictions of dynamic flow properties [251]. Reproduced with permission.

Summary. While full cardiac models of heart function for detailed exploration and analysis are of high interest to the medical community, visualization research efforts to aid these tasks are relatively limited in comparison to the efforts dedicated to blood flow. In the analysis of perfusion data for whole-heart pathologies, a higher volume of works focus on solving tasks for clinicians. However, studies on their actual adoption and utility in a clinical setting are limited. Communication-focused works that visualize heart function, while numerous in the field of medical illustration, remain a comparatively limited topic in visualization. A number of recent works to visualize heart function rely on multiscale models. We discuss these works in Sec. A.9.

A.8.3 Lung Function

The main function of breathing, i.e., respiration, is to provide oxygen to the tissues in the body and to remove carbon dioxide [180]. We focus here on the in- and outflow of air from the lungs and on limited cases where research includes other organs affected by lung movements.

Experts interested in learning about features of lung deformation during breathing often rely on simulations, commonly from statistical modeling, to visualize this process [119, 258, 435, 487, 564]. These generally integrate with imaging data to provide spatial context, often via surface modeling or volume rendering techniques. More recent approaches have used neural networks to reconstruct lung deformations as surface meshes from 4D CT data [541].

Experts are also interested in understanding patterns and features of dynamic airflow. Interest may be in visualizing airflow patterns within bronchial tubes [469] or on a larger scale. Kim et al. present a coupled model of tissue deformation on the level of the whole lungs alongside network airflow, enabling predictions of various dynamic flow properties [251]. The model is multiscale, but, as shown in Fig. A.18, does not necessarily provide spatial visualization of cell-level air exchange and lacks visual interactivity. Wiechert et al. [555] couple tissue- and organ-scale processes to allow visual exploration of tissue regions locally and the whole lung and airway system globally.

In more analytically-focused cases, visualization enables the comparison of movements of an object of interest against an acquired signal [293]. Similar to heart function, visualization is also used to evaluate model accuracy against acquired data and to assess and compare the magnitude of lung deformation using heatmaps [359].

Respiratory visualization models, in addition to parameter exploration, may also be used to compare characteristics of air flow in healthy versus diseased patients [251]. In other instances, experts use lung function to indirectly provide information in the analysis of other organs. For example, breathing exerts force on the kidneys. The degree that the kidneys are compressed during the breath cycle can be used to evaluate kidney fibrosis [483].

Summary. The bulk of literature we found from visualization and domain research focuses heavily on expert exploration of airflow and lung deformation over breath cycles, with limited tools for analysis and even more limited work in visualizing lung function for communication. These works primarily come from outside the visualization community and represent an open opportunity to develop visual methods to support experts in better understanding and analyzing lung function. These tasks are particularly important with the advent of COVID-19, as experts work to understand the long-term effects of this disease on lung function.

A.8.4 Brain Function

The brain is part of the central nervous system that contains more than 100 billion neurons. It is the primary seat of control for any process occurring within our bodies. Understanding brain function provides a key to understanding human behavior as well as neurological diseases and disorders. While our discussion of signal propagation in Sec. A.7 focused on the propagation of action potentials between cells, we now discuss signal propagation and functional neural connections over the entire brain. This is known as functional connectomics, where the brain is modeled as a network [473].

Related Surveys. Margulies et al. [318] and Pfister et al. [399] provide an overview of ways that the human connectome, both structural and functional, can be visualized for different exploratory, analytical, and communication tasks. Node-link diagrams, scatter plots, dendrograms, and heatmaps are common techniques in the application domain to visualize synchronous activity between brain regions. Structural models often provide spatial context for functional connectivity and can be depicted as image slices, surface models, volume renderings, or, in the case of DTI data, through advanced techniques that include ellipsoid or brush-stroke glyphs [277] and superquadric glyphs [252].

Experts are interested in learning how different regions of the brain functionally connect and in exploring patterns of brain activation in response to the presence, or absence, of certain stimuli. This represents a broader-scope view of the questions experts have when studying signal propagation. Heatmaps superimposed onto imaging data or derived surface models to show activation regions is a common approach that allows experts to explore and evaluate functional imaging data [134, 426, 550]. Interactive exploratory methods include dynamic querying for structural and functional connectivity using DTI and fMRI data [463].

Visual analysis methods help experts identify functional parameters and connections of interest. Standard methods often use a correlation matrix to identify functional connectivity [399]. Interactive analysis approaches often incorporate linked views that incorporate structural data alongside plots containing additional functional information, e.g., time plots [227, 310] or radar plots [320]. Color-coded isosurfaces from

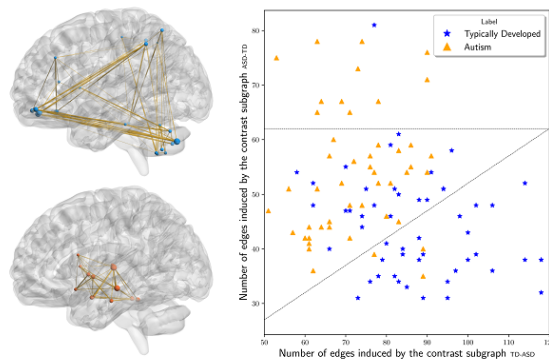


Figure A.19: Approach for classifying brain networks based on extracting contrast subgraphs, i.e., a set of vertices whose induced subgraphs are dense in one class of graphs and sparse in the other [280]. Reproduced with author permission.

structural imaging, with network and scatter plots to show functional connectivity and correlations [107], are another example of interactive visual analysis interfaces. These interactive, multi-view approaches may facilitate hypothesis generation in addition to confirmative analysis and integrate several data types and plots [241]. Clustering methods can group structural fibers from DTI data into to functionally-meaningful bundles. These can then be color-coded to aid identification and comparison of functional groups [153], or to identify and compare resting state networks that are again color-coded [516]. Other classification approaches to aid analysis use contrast subgraphs, shown in Fig. A.19, as a primary means to compare between groups [280].

While the bulk of visualization research for brain function is targeted at experts, some works that use illustrative techniques can be used for communication to a broader audience. The work by Jainek et al. [225] exemplifies such an instance for their use of glows and a soft, harmonious color palette to show brain activity.

Summary. Visualizing brain function is challenging. This is owed to the high computational power needed to simulate activity over an entire brain, and the fact that many visualizations of brain activity are driven by imaging data that only indirectly indicate brain activity and function. An additional challenge is that functional networks and structural connections do not always overlap. Visualization research to depict these uncertainties can aid neuroscience researchers. Furthermore, tools tend to focus on tasks related to exploration and analysis for domain experts. The development of visual methods targeted toward clinical rather than research use to identify aberrant patterns of brain function is an ongoing challenge. Finally, as previously discussed, further research into communication-oriented approaches that facilitate doctor-patient communication and patient understanding are essential for raising the bar of health literacy and public health. For example, data-driven approaches can communicate public safety stories, such as the impact of alcohol on brain function when driving a vehicle.

A.8.5 Other Organ Function

Although we focus our survey mainly on blood flow and functions related to the heart, lungs, and brain, we briefly highlight other organ functions that are representative of further opportunities in visualization research for physiology.

Skeletal muscle plays a crucial role in body movement and defining anatomical shape. Lee et al. [291] review visual approaches, which often use 3D surface models, for modeling muscle deformation and simulation of skeletal muscle functions at this scale. Experts are chiefly interested in understanding the shape deformation that skeletal muscle undergoes during contraction and relaxation [44]. Visualizations represent muscles at different degrees of abstraction to serve different objectives in a simulation, e.g., a single action line to show the axis of movement [365] or reconstructing and simulating a subset of muscle fibers that capture the overall shape of the muscle as it deforms [263, 264, 423]. Visual analysis approaches are often interested in quantifying muscle properties during contraction, such as muscle stiffness [464] or muscle speed in contraction [25]. These data can be captured in multi-view visualizations that combine structural models with line plots that describe displacement, velocity, and acceleration over the course of the simulation [573].

The stomach and liver are part of the gastrointestinal system whose main functions are (1) to take in food and liquids and break them down into a usable form and (2) to remove waste from the body. The stomach is a highly elastic organ that serves as a temporary holding place for food and is responsible for its initial breakdown. To understand the stomach's changing dimensions over time, Gilja et al. [158] present a multi-view system that includes a planar map of a region of the stomach for experts to better understand these temporal changes. The liver is the centerpiece for many essential bodily functions, including blood and nutrient storage, with a complex physiology that could greatly benefit from visualization. Lin et al. [302] visualize fluid transport through the liver and its vasculature at the organ and tissue scales, using simulations in combination with imaging and animation techniques. Their aim is to understand how the liver absorbs and metabolizes substances.

Metabolic activity, as discussed in the context of molecular pathways in Sec. A.5.3, can be visualized on the organ level to assess for normal organ function. Approaches can be analytically-focused, as in Nguyen et al.'s method to visualize uncertainty from PET kinetic modeling [369]. Ropinski et al. [424] exemplify a more exploratory approach to visualizing organ-level metabolic activity. Their approach is also designed to facilitate doctor-doctor communication through *interactive closeups*, whereby users can adjust view layout and composition to best fit their communication agenda.

Summary. Organs function as the result of a chain of processes that begin at the molecular level and extend through the cell and tissue scales. While blood flow, heart, and brain function are especially well-covered in visualization research, the lungs and other organs that we briefly highlighted in this section have received less attention, although experts clearly benefit from visualization tools to aid their exploratory and analytical questions. As in other scales, we observe a comparative lack of visualization research oriented to communication tasks. This represents an open opportunity for future work. Furthermore, while visualizations of organ-scale processes are often highly clinically-motivated, relatively limited research investigates visual methods to aid practicing clinicians which are usable in a time-crunched environment.

A.9 True Multiscale and Beyond

Cakmak et al. define multiscale visualizations as those that allow users to present, navigate and relate data across multiple abstraction scales [79]. In this section, we highlight examples and trends of true multiscale visualization for physiology, where the spatial and temporal representations and associated user tasks span three or more scales within our taxonomy.

The Visible Human Project is one of the earliest multimodal data initiatives to visualize the human body in its entirety. Although the project mainly focuses on anatomical structures, one of its key goals is to link structural image data with text-based physiological data [3]. Mathematical modeling initiatives to simulate multiscale human physiology like HumMod [197] are a rich resource for physiological models. The HumMod Browser, built from empirical data from peer-reviewed physiology literature, relies on grouped word clouds to allow experts to explore hierarchical and causal relationships of whole body physiology [565].

Works that are similarly exploratory, but include spatial information in their visualization, include Insley et al. [216]. They present a multiscale, multiphysics simulation spanning cell to organ scale of the formation of a blood clot within a cerebral aneurysm. Their method allows the visualization of individual red blood cells, platelets, and solvent particles. It expands further to visualize large-scale flow patterns with streamlines, enables the observation of platelet aggregation along the aneurysm wall, and shows this phenomenon in the context of the surrounding vasculature. Miller et al. [343] present a multiscale, although primarily structurally-focused, brain map that spans molecule to tissue scale. They include functional information from spatial transcriptomics data to describe pathological Tau proteins as well as signal propagation information. Leggio et al. [294] present MorphoNet, an open-source online tool allowing users to interactively explore the anatomy and dynamics of biological entities from molecule to whole-organism scale. It furthermore allows for genetic data to be overlaid onto these models. Primarily focused on developmental processes, this tool uses 3D color-coded surface models and is targeted at research and education, as shown in Fig. A.20.

Qutub et al. [411], whose work we briefly discussed in Sec. A.7.1, present a review of multiscale modeling approaches, from molecule to organ scale, for angiogenesis. Although many of these models are multiscale, the visualization result is often not multiscale. An exception to this, although still limited and not interactive, includes Mac Gabhann et al.'s [312] multiscale muscle model. This model includes muscle fibers (muscle cells), the microvascular bed that supplies oxygen to these cells and tissues, and the associated molecular pathways for angiogenesis.

While several truly multiscale models for the heart range from the level of ion channel opening to the heart's contractions over a full cardiac cycle, the visualization output often is limited to an organ-scale mesh representation. A heatmap then encodes electrical activity mapped to the surface mesh and static, non-interactive charts display the changes that occur at the cellular, molecular, or tissue scales [29, 175, 430]. Gil et al. [157] expand on this typical representation by incorporating myocardial fibers, extracted from DTI data, with simulation data to understand ventricular muscle tissue structure and connectivity. Chabiniok et al. provide a review of multiscale cardiac modeling methods with an eye toward their integration for analysis in clinical practice [84]. This shows a clear interest in understanding such models from the domain.

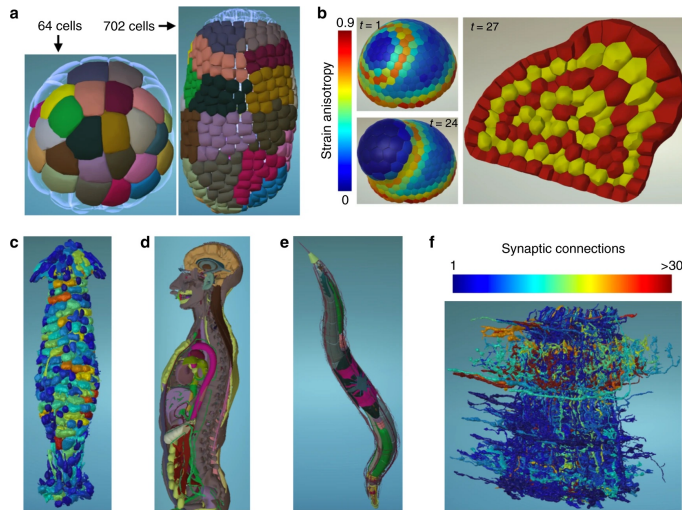


Figure A.20: MorphoNet is an open-source online tool to explore and communicate the anatomy and dynamics of biological entities from imaging and genetic data [294]. Reproduced under Creative Commons CC BY license.

Some multiscale lung function models span the molecular to organ scale, such as those by Burrowes et al. [75], but lack integrated visualizations across all scales. As presented in Sec. A.8.3, Kim et al. [251] present a partially integrated visualization that captures the airflow dynamics of their model but lacks spatiality across levels. Furthermore, the visualization, as in many of these systems, does not allow for direct user interaction.

Multiscale models for brain function are becoming increasingly common as computational power increases. Spanning cell to organ scale, such models capture the depolarization of a single neuron, signal propagation through brain tissue, and the effects that this signal has on brain function [120]. However, the visualizations of these models are often limited in the scales they depict. For example, in the Cognitive Computation Project, Ananthanaryanan et al. [15] simulate a cat brain and visualize parts of the model with a 2D heatmap that plots groups of neurons with similarly-timed firing rates in a cortical area. A topographic plot provides a detailed view of the first spike of signal within each neuron group.

Increasing works target the developmental stages of entire organisms through communication-oriented tasks that present physiology for education and outreach. Sorger et al. [470] use visual abstraction techniques to show, at a molecular level, the transitional stages of an HIV-virion. On a larger scale, the OpenWorm and its various subprojects [437] aim to visualize and simulate multiscale anatomy and physiology of *C. elegans* with numerous tools designed for public outreach. On a system-specific level, full-body virtual anatomical models of the human musculoskeletal system are a lively research topic, not only for domain experts [314] but for visual effects animators as well [423]. Multiscale modeling approaches, such as those by Rzepecki et al. [428], propose to combine multimodal structural, physiological, and biomechanical

data sources in an interactive viewer that scales from visualizations of cartilage tissue porosity up to a simulation of human gait.

A.10 Discussion

This survey is intended as a guide for visualization researchers interested in understanding common approaches and challenges to visualizing physiology from a spatio-temporal and task-oriented perspective. In this section, we discuss general themes along with lessons learned.

vitaLITY. Our literature collection approach utilized traditional search methods and leveraged new visual analysis tools from the community to facilitate this process [361]. Using vitaLITY was a huge help to identify holes in our search. Although vitaLITY does not span the space of literature that we covered, its coverage of the visualization literature is comprehensive, and helped us identify whether any holes in our search were due to issues with our methodology or due to the lack of visualization literature for a given topic. For example, using this combination of search methods, we discovered that the spike in publications in 2010 could be attributed in part to the rise in interest in the visualization of omics-related data, with four surveys published on the topic that year [154, 370, 377, 378] along with the running of funded research projects like the Physiome, IllustraSound, and PhysioIllustration projects in Europe during this time frame. Close to 40% of the papers included in this survey were found using vitaLITY. In many cases, we used papers found in this tool as seed papers in a standard search methodology. Using this combination of tools helped give us confidence that our search methodology was even, in spite of an apparent bias towards, e.g., organ-scale functions. Through vitaLITY, we were able to sanity-check that these are heavily-weighted topics in visualization for physiology.

Spatio-Temporal Distribution. Classifying literature along a spatio-temporal axis uncovers a few interesting patterns in Fig. A.5. Most salient is the dark grouping of works related to organ function that is positioned symmetrically on both spatial and temporal axes. The data acquisition methods to capture these processes (brain, heart, lung function, and blood flow) are well-established with strong clinical motivations. It is not surprising to see an abundance of work in this region, given these two factors.

Another, though less dark, region we observe ranges spatially from large molecules to cellular substructures and occurs over minutes. This corresponds to the abundance of work where the key process of interest in the visualization is understanding gene expression and where the majority of experimental methods, e.g., next-generation sequencing and proteomics methods, lack the temporal resolution to capture the active process of translation and transcription [105]. These capture the result of the process, which occurs typically in the span of minutes (for a single gene) [344]. This density of work reflects the comparatively recent developments in experimental methods for measuring gene expression, which similarly accounts in large part for the increase in works over the last few years, as we observe in Fig. A.4. We expect to see a continuing increase in such works.

The light region corresponding to a temporal resolution of a few seconds (10^1 sec) represents a few possibilities. One relates to the temporal resolution of experimental methods for gene expression: the majority of next-generation sequencing methods are,

at best, on the scale of minutes [105]. Another reason is that, while a number of processes bridge over this time range, for human physiology we found limited processes that are confined to this range. Diffusion of molecules across a cell can occur in this time span [344], but these are typically of interest when bundled into the greater context of a molecular pathway or in the dynamics of an entire cell, which encompass a larger time frame.

Smaller-scale processes, particularly at the molecular scale, have a much broader potential temporal range than the organ-scale processes our study examined. As many works in our survey allude to, this broad temporal range is both an enormous challenge and opportunity for visualization to aid experts in exploration and analysis of their data when events of interest are easily lost in temporal noise. A major, related, ongoing challenge is in integrating into a visualization the larger spatial scales that build from, and are affected by, the molecular process of interest. The discussion of the differing breadth of time scales for the different spatial scales introduces another point: while temporality generally increases with biological complexity [452], living organisms are not bound to a system based on powers of ten, but rather to a roughly 24h cycle known as the circadian rhythm [346]. This explains, in part, the temporal ranges of functions such as gene expression and the heartbeat and breath cycles that we observe in Fig. A.5.

Cell & Tissue Function. We found comparatively few cell- and tissue-scale visualization works. This reflects the history and trends in the availability and technological advancements of the source data. Data that can truly visualize dynamic, living cell processes on the scale of molecule, cell, and tissue have only just recently become available and accessible, and we see in the publication dates that cell-related visualization works are on the rise. Computational power is also steadily increasing to the point where whole-cell visualizations are becoming a reality, while the additional complexity inherent in tissue-level physiology visualization is still a challenge. Furthermore, many application domain approaches use only basic visualization techniques, often simply reviewing microscopy imaging data, to explore or quantify features or behaviors of interest. This may show a lack of trust in abstraction that is, on some level, unavoidable when processing data to visualize through other methods. Approaches that incorporate raw imaging data in multi-view interactive tools alongside uncertainty quantification are useful directions to continue to investigate.

Imaging and Simulation Data Across Scales. Visual exploratory and analytical tasks at the organ scale and, to a lesser degree, at the tissue and cell scales are typically closely tied to imaging data. In cases using a model, the model usually is compared or validated against imaging data. Hence, visualization tasks at larger scales are often to understand a given process from imaging data. This is not necessarily the case in real-time visualization of dynamic processes at the molecular scale, and in many instances at the cellular scale. Here, the visualization of a model often serves as the primary means to understand a given process. As technology improvements lead to higher resolution imaging methods, or if simulations come to be seen as more accurate and trustworthy than the imaging methods themselves, we may see this balance shift.

Task Distribution. Task distribution between exploration, analysis, and communication differs slightly across all scales, with one major trend consistent across all: communication-oriented visualization works for physiology are limited relative to exploration or analysis-oriented works. Most visualization works we surveyed have con-

crete expert collaboration partners with specific data types and specific goals to understand these data. This works at a much lower level than communication. There are some possible explanations for this disparity in task distribution. One is that the data being visualized is often cutting-edge, and the domain scientists producing these data may not yet fully understand it themselves—visualization is needed to discover and analyze, and the communication element comes after understanding. A second possibility relates to data-related permissions. Patient and research subject data is often heavily protected and, in many instances, may not allow visualization beyond internal, domain-specific analytical use or requires significant processing efforts to anonymize the data before use. Lastly, visual communication of this type of data simply is hard. The data are highly complex and multifaceted, and visual communication requires a high degree of understanding on both the side of the domain expert and the visualization researcher to distill this information into a clear narrative.

Application Domain Adoption. Many of the visualization approaches discussed in this report have not fully permeated the application domain. These are often highly-specific algorithms or techniques, while widely-adopted approaches in the domain are often more generally applicable. Furthermore, such production-ready, i.e., stable, solutions are continually maintained and designed for ease of use. For example, in visualizing molecular dynamics, tools such as VMD [210] allow for easy visualization of simulations through a movie-like series of time steps. Across all scales, visualization methods adopted in the domain remain relatively restricted to direct visualization of imaging data in a time-lapse sequence, as in light microscopy or medical imaging, e.g., fMRI. Techniques can extend into volumetric rendering with limited features for exploration and analysis, with the option to create surface meshes from these data. These visualization methods are easily available in tools like Amira [476], 3DSlicer, [131], and ParaView [9]. Additionally, across all scales, basic visualizations such as bar or scatter plots are common to identify the frequency or distribution of physiological biomarkers under study. These are often created in tools like Microsoft Excel. Such charts are often limited in interactivity, e.g., filtering, and do not apply across multiple views. Coordinated efforts with funding agencies to establish initiatives for the deployment of advanced visualization techniques can enable broader access from the application domains. Furthermore, researchers can consider opportunities to develop advanced visualization approaches as plugins to existing domain tools rather than as stand-alone solutions. This is an opportunity to enable greater domain accessibility, and can mitigate the resource limitations introduced by stand-alone tools.

A

A.11 Research Outlook

Physiology is a challenging, complex domain that visualization can do much more to contribute to in the coming years. Increasingly sophisticated modeling and data acquisition methods can capture physiology at finer spatial and temporal resolution but, in exchange, produce even higher volumes of complex data. In addition, the multiscale, multisystem, multidisciplinary nature of physiology needs visualization to help bridge gaps, not only in exploration and analysis of data between scientists but in communicating this science to a broader audience.

Numerous recent technological advances in imaging and experimental methods pose exciting opportunities for visualization across several scales. While light microscopy was previously limited to a maximum resolution of 200 nm, Pulsed Interleaved MINIFLUX with a standard microscope has increased this resolution to 1 nm, allowing visualization of metabolites and other small molecules *in vivo* [327]. The boundaries of computed tomography have been similarly expanded with hierarchical phase-contrast tomography (HiP-CT) to allow true multiscale imaging from the organ down to the cellular level [536]. Other data have not received much attention from the visualization community, such as the suite of methods used to assess cell biomechanical forces [38]. Pioneering experimental methods to observe protein translation occurring in real-time in living cells, such as nascent chain tracking (NCT) [105] are available, but visualization of these methods is limited. Expanding visualization research to collaborate in such areas to develop new methods for experts to engage with these data is an enormous opportunity.

To answer questions left by gaps in systems biology and integrative physiology, research is shifting to focus on the human organism as a complete integrated network. This considers not only the scales that have been the primary focus of this report but also the coordinated efforts between organ systems and sub-systems. This study of the human organism as an integrated network is termed network physiology, with a set of grand challenges in this new discipline published only recently [223]. Visualization methods and tools that can meet the exploratory, analytical, and communication demands for this area of study are exciting opportunities.

Multiscale computational models are increasingly ubiquitous with advancements in computational power and parallel processing. However, while these multiscale models exist, corresponding multiscale visualizations are often lacking or exist in unlinked, unintegrated forms. This is another opportunity for visualization research.

True multiscale and multisystem approaches necessitate multidisciplinary collaborations. A dearth of visual methods and tools facilitate knowledge transfer between domains or communicate physiology to the public. As the last two years of the pandemic have demonstrated, clear and accurate visual communication of physiology for public health is critical at all levels of society.

A.12 Conclusion

This survey offers a broad overview of visualization trends and opportunities for physiology and aims to provide a foundation for discussion and future research directions in this area. From a mixed-methods literature search approach that uses state-of-the-art visual analysis tools, we embed our discussion of these approaches in a spatio-temporal context that focuses on the core tasks that drive the visualization: exploration, analysis, and communication. Our report demonstrates an abundance of work at the organ scale, particularly for hemodynamics. Molecular visualization, particularly related to visual analysis and exploration of actors in molecular pathways, is a growing research area driven by the advent of new technologies. These new technologies hold immense promise for visualization research that incorporates multiple data types to span the true multiscale nature of human physiology, from molecule to organ scale and beyond.

Acknowledgments

This work was supported by the Visual Data Science for Large Scale Hypothesis Management in Imaging Biomarker Discovery (VIDI) funded by the University of Bergen and the Trond Mohn Foundation in Bergen (813558). Parts of this work have been done in the context of CEDAS, UiB's Center for Data Science. We are grateful as well to Juraj Pálenik, Bara Kozlíková, Bernhard Preim, Renate Grüner, and Eduard Gröller.

A.13 Cited Literature in Report

This section provides an overview of the literature cited in this report on visualization for physiology, organized according to spatio-temporal scale with glyphs encoding high-level visualization task classification. To reduce repetition of information while still providing a table overviewing the literature cited in this paper, we have modified this from the original published version.

Scales: Molecular Cellular Tissue Organ Multiscale Survey

Tasks: Exploration Analysis Communication

**Color fill amount corresponds to primary task(s) each work addresses*

Molecular Function Works:

[13] [45] [62] [64] [73] [74] [78] [76]
 [77] [81] [94] [97] [98] [114] [138]
 [145] [147] [146] [149] [311] [172] [176]
 [177] [178] [179] [183] [194] [209] [210]
 [229] [228] [232] [237] [242] [243] [246]
 [248] [254] [255] [266] [265] [271] [279]
 [290] [295] [355] [356] [297] [299] [300]
 [304] [325] [340] [341] [89] [350] [363]
 [368] [388] [422] [439] [444] [445] [447]
 [448] [454] [458] [460] [499] [500] [512]
 [524] [528] [578]

Cellular Function Works:

[26] [35] [54] [58] [99] [101] [118]
 [126] [125] [124] [127] [130] [156] [160]
 [162] [166] [196] [202] [203] [204] [224]
 [236] [240] [256] [281] [292] [322] [386]
 [389] [398] [405] [410] [414] [431] [432]
 [446] [449] [467] [474] [515] [534] [537]
 [539] [544] [552] [569]

Tissue Function Works:

[22] [36] [41] [65] [68] [69] [86] [88]
 [113] [115] [161] [181] [192] [201] [212]
 [283] [306] [313] [319] [339] [342] [343]

[345] [347] [348] [373] [380] [379] [382]
 [411] [416] [425] [442] [468] [496] [497]
 [501] [502] [563]

■ Organ Function Works:

[9] [16] [20] [21] [23] [25] [29] [39]
 [40] [44] [49] [53] [59] [71] [83] [91]
 [100] [104] [103] [102] [111] [119] [123]
 [129] [131] [134] [139] [140] [142] [151]
 [152] [153] [155] [158] [171] [218] [182]
 [184] [188] [206] [225] [227] [238] [241]
 [258] [262] [264] [263] [270] [276] [280]
 [282] [286] [285] [284] [293] [296] [302]
 [310] [317] [320] [333] [334] [338] [337]
 [335] [359] [365] [367] [371] [381] [383]
 [391] [522] [396] [397] [420] [424] [426]
 [430] [435] [451] [462] [463] [464] [469]
 [483] [486] [487] [492] [494] [495] [516]
 [107] [520] [519] [541] [550] [555] [560]
 [562] [564] [566] [567] [573]

■ Multiscale Works:

[15] [63] [75] [157] [175] [197] [216]
 [251] [294] [312] [314] [423] [427] [428]
 [434] [437] [470] [472] [476] [565]

■ Related Surveys:

[10] [11] [12] [50] [79] [84] [95] [132] [133] [154] [169] [170] [199] [200] [214] [239]
 [260] [261] [267] [268] [272] [287] [291] [318] [324] [323] [326] [330] [331] [364]
 [370] [372] [374] [376] [378] [377] [384] [393] [521] [399] [401] [403] [406] [433]
 [440] [441] [452] [466] [477] [484] [506] [523] [529] [538] [574]

A

Paper B

Interactive Visual Exploration of Metabolite Ratios in MR Spectroscopy Studies

Laura A. Garrison^{1,3}, Jakub Vašíček², Alexander R. Craven^{5,6},
Renate Grüner^{3,4}, Noeska N. Smit^{1,3}, Stefan Bruckner^{1,3}

¹ Dept. of Informatics, Univ. of Bergen, Norway

² Dept. of Informatics, Masaryk University, Brno, Czech Republic

³ Mohn Medical Imaging and Visualization Centre, Dept. of Radiology, Haukeland Univ. Hospital, Bergen, Norway

⁴ Dept. of Physics and Technology, Univ. of Bergen, Norway

⁵ Dept. of Biological and Medical Psychology, Univ. of Bergen, Norway

⁶ Dept. of Clinical Engineering, Haukeland Univ. Hospital, Bergen, Norway

Abstract

Magnetic resonance spectroscopy (MRS) is an advanced biochemical technique used to identify metabolic compounds in living tissue. While its sensitivity and specificity to chemical imbalances render it a valuable tool in clinical assessment, the results from this modality are abstract and difficult to interpret. With this design study we characterized and explored the tasks and requirements for evaluating these data from the perspective of a MRS research specialist. Our resulting tool, SpectraMosaic, links with upstream spectroscopy quantification software to provide a means for precise interactive visual analysis of metabolites with both single- and multi-peak spectral signatures. Using a layered visual approach, SpectraMosaic allows researchers to analyze any permutation of metabolites in ratio form for an entire cohort, or by sample region, individual, acquisition date, or brain activity status at the time of acquisition. A case study with three MRS researchers demonstrates the utility of our approach in rapid and iterative spectral data analysis.

B

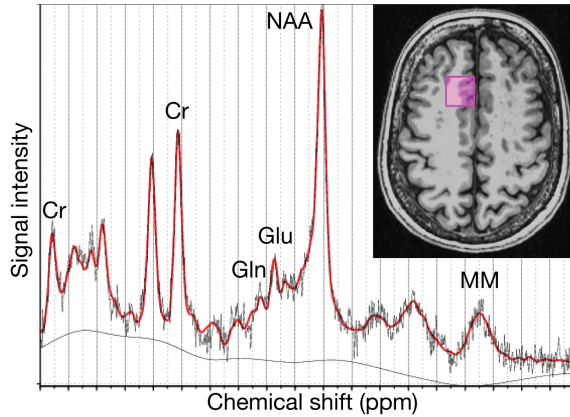


Figure B.1: Plot of the typical visual output for spectral quantification by LCModel, with processed data in black and model fit in red. Structural localization and annotations of key metabolite peaks have been added.

B.1 Introduction

Magnetic resonance spectroscopy (MRS) is an *in vivo* non-invasive biochemical technique used to estimate the concentrations of certain small molecules, known as metabolites, in a tissue region. When paired with high structural resolution MR imaging (MRI), it has shown clinical potential for improving diagnosis and treatment monitoring of numerous diseases and disorders of the central nervous system [517]. However, its clinical adoption remains limited. Translation from the metabolite signals acquired from MRS into clinically useful biomarkers is an open challenge in spectroscopy research. Optimization and tuning of parameters for consistent, isolated metabolite acquisition is one such area of research, while another branch of research aims to identify patterns of the subtle disease effects on multiple metabolites [489]. In this paper, we explore the application of visualization techniques to identify ratios and patterns of multiple metabolites.

While recent technology improvements in MRS acquisition have enhanced data quality and resolution [525], visualization of MRS data remains a largely unexplored area. MR spectroscopy produces a vastly different readout than MR imaging. Rather than a greyscale image of recognizable anatomical structures over many voxels, it acquires an abstract spectrum per single voxel. This spectrum consists of a series of peaks (resonances) that represent signal intensities as a function of frequency, as depicted in Fig. B.1. Metabolites may consist of single peaks, as in the case of N-acetylaspartate (NAA), or multiple peaks, as in Creatine (Cr). Most tools used to quantify single voxel spectral data, e.g., LCModel [408] produce only rudimentary visual output, such as the spectral graph in Fig. B.1. Recognizing the metabolites that correspond to these graphs is challenging. Although it is important to see the spectral graph as a means of quality assurance, metabolite concentrations are the most clinically relevant output from this method. These concentrations are most often output to a simple table in standard domain tools. This does little to advance interpretation, understanding, or to facilitate rapid comparison of metabolites between acquisitions.

This paper expands upon our previous design study [150] in building a general tool for the interactive visual analysis of all permutations of spectral metabolites, in ratio form, for a small cohort. While we previously emphasized rapid visualization of metabolite ratios directly from spectral input data, this work allows visualization of complete, and more complex, metabolic signatures via an integrated pipeline with Tarquin [557], an open source spectral quantification tool. Our specific contributions include:

1. We provide a detailed review of MRS data characteristics and abstraction of spectral analysis tasks identified from domain expert collaboration.
2. We present a refined pipeline that integrates spectral quantification and fitting to allow multi-peak metabolite analysis.
3. Our visual exploratory analysis tool provides an extended interface for linking of structural, spectral, and patient data, including group creation and uncertainty communication.
4. We introduce a tiered system of visual encodings depicting layers of aggregated metabolite ratios that can be partitioned by key attributes.
5. We present a clinical case study and feedback from three MR spectroscopy research experts.

Using SpectraMosaic, MR spectroscopy researchers are able to rapidly identify patterns at different layers that may be of interest for deeper clinical exploration.

B.2 Related Work

A key challenge in visualizing spectroscopy data is that each spectrum is in itself a multivariate dataset. We draw inspiration from tools such as InSpectr [14], which utilizes multiple linked views and comparative visualization techniques [163] from multimodal data sources (x-ray computed tomography and x-ray fluoroscopy) to provide insights into composition of a multivariate sample. SpectraMosaic similarly combines imaging techniques (MRS and MRI), but for a different domain and with a different focus. Isosurface similarity maps defined by Bruckner and Möller [72] were applied to spectra in Spectral Similarity Maps, an extension of the InSpectr framework [141]. In this approach, correlations between spectra are shown as an intensity map. We adopt a similar concept in our tool, but rather than mapping energy correlation we instead map metabolite ratios.

Prior visualization approaches for MRS data have been limited to the analysis and visualization of a subset of metabolites at a time. SDDS (scale driven data spheres) presented by Feng et al. [136] provide a 3D representation of metabolites within a voxel. This application was later extended to include scatter and parallel coordinate plots for a subset of metabolites [135]. SpectraMosaic remains in the abstract visualization space, but allows comparison of all metabolite ratios. Nunes et al. [373] presented a visual analysis framework combining ComVis [328] and MITK [559]. Brushing and linking mechanisms allow for the definition of a biological target volume with its corresponding metabolite values. However, this work was developed specifically for radiotherapy

treatment visualization. Retention of spectra was not the focus of the application and it provided limited functionality for metabolite comparison. SpectraMosaic extends the flexibility of metabolite ratio calculations, and displays additional MRS data attributes (spatial, individual, temporal, and brain activity status) in an overview and detail visual representation. Marino and Kaufman [321] implemented direct volume rendering (DVR) to represent male prostate anatomy from MRI data combined with PET and MRS in prostate tumor delineation. However, this application was focused on a single metabolite ratio, and could only present an individual in a single time slice. SpectraMosaic retains an abstract visualization format, but offers broader insights into metabolite relationships over time and between individuals. Jawad et al. [229] developed a system for the analysis of segmented brain tissue composition to identify the metabolic signatures of brain tumors—this tool was optimized for multivoxel data, and focused on statistical outcome measurements. SpectraMosaic works at a more generalized level in spectral analysis. Further work by Jawad et al. [228] presented an approach for the comparative analysis of single voxel spectroscopy in cohort data, focusing primarily on violin and parallel coordinate plots to convey spectral metabolite relationships. Our approach uses a similar range of data inputs and processing tools. However, our tool focuses on simultaneous comparison of all metabolite ratios, using a nested visual design linking multiple MRS data elements.

First introduced by Bertin [47], numerous solutions have leveraged small related graphics series to visualize multivariate data. We base SpectraMosaic on this concept, but extend this by including a second layer of nested visual encodings. This is inspired by ATOM [390], a grammar for unit visualizations where individual data items are represented by unique visual marks (units) in a visual encoding system. PivotTable, subsequently trademarked by Microsoft and extended by Polaris [482], enables exploration and analysis of multidimensional data with the flexibility to modify visual encodings, graphics, and table configuration for visualization. Klemm et al. [257] built on this concept for linked visualization of image-centric heterogeneous cohort data. Our approach is related in that we allow on-the-fly reconfiguration of our matrix inputs. Although the cohorts our application focuses on are not large, we share similar considerations with heterogeneous and multivariate data inputs.

While our prior iteration of the SpectraMosaic application focused on the rapid analysis of single-peak metabolites directly from spectral graphs [150], this work expands the tool to allow full, precise spectral analysis in an integrated pipeline with robust MR spectroscopy quantification tools. This permits analysis of metabolites with more complex metabolic signatures; these are encoded to bar and box plots for ease of interpretation. We further increase the practical usability of the tool with new facilities for analysis group creation and additional means for conveying the underlying data distribution. These features arose from additional working sessions and discussions with spectroscopy researchers.

B.3 Background

MRS is an advanced spectroscopic technique used to non-invasively describe the biochemical composition of living tissue. While MRI shows the spatial distribution of atomic nuclei with high spatial resolution, MRS trades spatial resolution for detailed

chemical information, using the same hardware. For example, where MRI may be used to identify the extent of a tumor, MRS can help to identify the type of tumor [190]. For each measured voxel, MRS produces a spectrum of signal intensity as a function of frequency. Intensity peaks at different resonance frequencies are described as chemical shifts. These chemical shifts, expressed in parts per million (ppm), arise from fundamentally different nuclear properties of the chemical structures being measured, and represent metabolites in the acquired voxel [513]. The most commonly measured signal comes from hydrogen atoms; this is known as proton MRS (^1H -MRS). This technique is capable of detecting metabolites in concentrations 50,000 times lower than that of fat or water as imaged in conventional MRI.

MRS acquisition techniques include single voxel spectroscopy (SVS) or chemical shift imaging (CSI). CSI is essentially a slab of multiple smaller single voxels. It covers a much larger spatial area than SVS, but suffers from a reduced signal-to-noise ratio. CSI produces a low-resolution image for each metabolite, being in that way similar to conventional MRI, while SVS is more abstract and cannot be visualised in a conventional way. Since SVS acquisition techniques afford more detailed spectra for analysis, we focus our work on this technique. The majority of acquisitions by our collaborators are collected at single time points, i.e., in longitudinal studies, but may also be captured as time-resolved concentrations within a single examination, i.e., functional studies. In the latter approach the subject can also be asked to perform tasks, such as tapping fingers during the acquisition (active brain state), and alternately resting (resting brain state).

Following acquisition, data are output to a vendor-specific format that contains raw data and a header file containing all experimental parameters. Subsequent preprocessing and quantification steps follow to map spectral peak intensities to metabolite concentrations in the measured voxel. In a final fitting step, a model based on prior information is fit to the acquired spectrum; in many approaches, this is effectively a linear combination of basis sets consisting of simulated or measured metabolite signatures. Metabolite concentrations are typically calculated relative to a stable reference, often water or creatine. This allows for a direct comparison of relative metabolite concentrations, assuming the same acquisition hardware and protocols are used. While a more comprehensive discussion of all steps is beyond the scope of this paper, interested readers can refer to Stagg et al. [475] for a detailed overview. A number of existing tools can be used to perform these steps: LCMoDel [408] is one such widely-used commercial tool, while jMRUI [479], TARQUIN [557], SIVIC [93], OXSA [409], and Gannet [352] offer open source solutions. Equipment manufacturers also supply basic tools to facilitate simple analyses on the scanner console. Our collaborators typically use LCMoDel or Tarquin; we utilize Tarquin in our pipeline for its ease of use and open availability. The output from these steps includes the experimental parameters used for the acquisition as well as the fitted data and quantification information for each metabolite.

B.4 Task and Requirement Analysis

We developed SpectraMosaic over the course of one year. We met weekly with our domain collaborators, two of whom are coauthors of this paper. Collaborator backgrounds

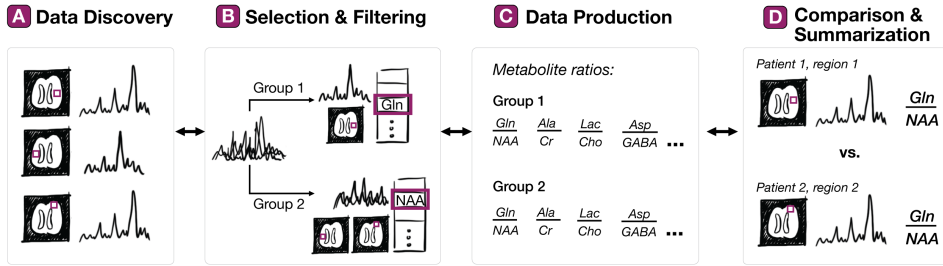


Figure B.2: Typical task flow for MRS data analysis. Users begin with data discovery (A) to review spatial voxel position, associated spectral graphs, and relevant acquisition parameters. (B) continues with data selection and filtering, where spectral voxels of interest are selected and divided into groups. Data production (C) calculates all possible ratios of selected metabolites, e.g., Glutamine (Gln) to N-acetylaspartate (NAA). In (D) ratios are compared and summarized between, e.g., Gln/NAA for different patients or different brain regions. Each of these steps may be revisited.

included two MD/PhDs in radiology, eight PhD researchers in MR imaging, and three MR engineers. The weekly meetings went through three distinct phases. The first phase focused on domain evaluation, identification of key challenges and where visualisation could potentially help overcome them. Ultimately, the output from this phase was agreement on core tasks and requirements. The second phase explored the design space for these tasks/requirements with discussion and interface prototypes. These were refined and narrowed down to a single option. Our third phase reviewed and refined an alpha application. Basic use case testing alongside individual and group evaluation feedback ultimately helped us settle on the version we present in this paper.

B.4.1 Task Analysis and Abstraction

We frame the analysis tasks identified in phase one of our collaboration in the context of Brehmer and Munzner's multi-level typology of abstract visualization tasks [67]. This abstraction was useful for our development process, as it allowed us to more objectively frame the challenges experienced by our colleagues. These tasks form a generalized workflow shown in Fig. B.2. The first step, data discovery, provides a general overview of the input components for spectral analysis. Following user selection of components for analysis, a data production step calculates ratios from all inputs. Ratio comparison and summarization follows.

T1: Data Discovery. The first set of tasks relates to data consumption for discovery and verification of key MRS data aspects (Fig. B.2A). Spectra, anatomical reference images, and associated subject data are reviewed together in an initial overview step. Researchers visualize spectral graphs to establish a general sense of the data quality and to form initial hypotheses. Supplemental parameter information, such as the echo time (TE), during the acquisition can be used to verify validity of experimental comparisons. Researchers additionally validate their assumptions about the spectral graph against its sample location. This serves two purposes: (1) as a second quality assurance

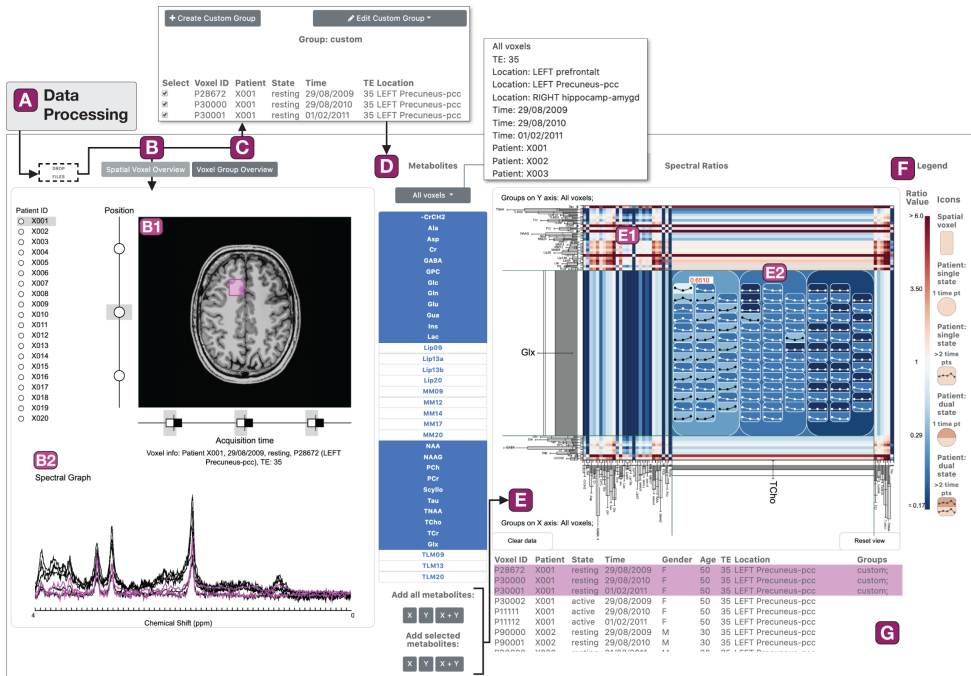


Figure B.3: SpectraMosaic application workflow overview. Raw spectral data are first processed in an offline step (A), then loaded into the application. In (B) the user visualizes the anatomical image with voxel placement for each acquisition (B1) and the associated spectral graph (B2). In (C) users may create custom groups for analysis. Metabolites may be selected (D) for analysis from custom or preset groups in a drop-down list, and selections assigned to the x- or y-axis of a ratio heatmap (E). The ratio heatmap is divided into a cell grid (E1) based on the number of metabolite inputs to each axis. Detailed inspection of a cell (E2) shows the ratios in a series of nested glyphs representing spatial region, individual, individual brain state, and individual time acquisitions. A legend at the right provides a reference for heatmap glyphs and colors (F). A table (G) shows relevant metadata for each voxel.

measure to check whether the data were sampled in the correct region, and (2) to provide initial validation for graph differences between spatial regions. This is because a normal spectrum in one area of the brain may be aberrant in another region with a different tissue composition [475].

T2: Selection and Filtering. Following an overview, researchers next select and filter the data (Fig. B.2B). In both medical and clinical research studies our collaborators often wish to select a subset of spectra or metabolites for further analysis for a variety of different reasons. For instance, researchers may wish to look only at the variation in metabolite concentration ratios for a single time acquisition in a longitudinal cohort study, e.g., pre-operative patients in a tumor cohort, or to analyze only female subjects within a study. Furthermore, some metabolites may be uninteresting to include for certain clinical studies, e.g., lipids and macromolecules are not usually relevant outside of

certain oncological studies [475], and are useful to exclude on-demand.

T3: Data Production. Spectra can vary considerably between acquisitions. This can occur due to different acquisition parameter settings or simply between different scanners. Ratios and correlations calculated from metabolite concentrations are two standard methods to understand spectroscopy data [315]. The use of ratios to determine metabolite concentrations is a core critical task for any MRS application for two reasons, (1) as a method to correct for inhomogeneity across the sample and (2) to account for varying tissue composition. Following selection of interesting metabolites for analysis, a data derivation step takes as input the Tarquin-processed and quantified metabolite values and outputs the metabolites in ratio form (Fig. B.2C).

T4: Comparison and Summarization. Following data derivation, researchers then wish to summarize and compare metabolite ratios (Fig. B.2D). For example, researchers studying oxygen deprivation (hypoxia) in newborns are interested in comparing the metabolic differences between healthy and hypoxic newborns. This can be achieved by evaluating ratios of the same metabolites between both groups, e.g., NAA healthy vs. NAA hypoxic. Furthermore, researchers would like to understand the metabolic profile of hypoxic newborns on a spatial and individual level. For instance, the basal ganglia region of the brain is known to be sensitive to oxygen deprivation, so it is clinically relevant to compare this region to a less sensitive region. Within a given region of interest researchers then wish to compare individuals to identify clinically relevant outliers in order to answer questions such as “*How does Lactate/Choline compare for Patient X versus Y?*” Moreover, oxygen-deprived newborns who survive often experience developmental disabilities later in life. Longitudinal MRS studies allow researchers to understand how the metabolic profiles of affected individuals change over time relative to healthy individuals. In a different scenario, researchers studying schizophrenia are interested in comparing the metabolic profiles of individuals when their brains are *active* relative to their *resting* brain state. Different metabolites present in different concentrations in these states, and identification of these differences may help progress understanding of this disorder.

Following comparison of interesting metabolite ratios, researchers often wish to refine their hypotheses and revisit metabolic input data. This task sequence then repeats, following an iterative analytical approach to hypothesis exploration and verification in MRS data.

B.4.2 Design Requirements

Following the identification of tasks important for our collaborators in MRS analysis, we developed the design requirements for our application. First, on a technical and infrastructure level, our colleagues often switch between hospital workstations while accessing sensitive patient data. Thus, for practical utility it is critical to provide a tool that enables a *machine-independent workflow* (**R1**) that adheres to *patient data restrictions* (**R2**).

As discussed in **T3**, for a combined analysis of spectra acquired from different scanners, or with different acquisition settings, it is necessary to calculate *metabolite ratios* (**R3**). Furthermore, as implied by **T1** and **T2**, *visual linking between input*

data (voxel placement, spectral graph, patient- and acquisition-specific information) and *calculated metabolite ratios* is important for many analysis questions (R4). For our collaborators, the most important patient and scanner-specific information to retain include patient age, gender, and echo time (TE).

Based on the types of questions outlined in T4, users must be able to *compare metabolite ratios of interest* (R5). This should be accomplished for any permutation through four key attributes: *spatial region, individual, time point, and brain activity state*. Additionally, appropriate mechanisms to compare ratios over time as well as between spatial regions and individuals are critical for longitudinal or single-run studies. Furthermore, for functional MRS studies it is important to support comparison of metabolite ratios in an active relative to a resting brain state.

B.5 SpectraMosaic Workflow and Interface

We provide an overview of the SpectraMosaic interface in Fig. B.3. Following an offline processing step, data are loaded into the web tool (Fig. B.3A). Data of interest for analysis can be explored, selected, and added (Fig. B.3B-D) to a spectral ratio heatmap for deeper inquiry and hypothesis verification (Fig. B.3E). A legend provides information on the encodings used in the tool (Fig. B.3F). A table below the heatmap summarizes salient acquisition information (Fig. B.3G).

Data Processing and Loading. We first perform an offline processing step that automates spectral processing and quantification from Tarquin and MATLAB [329]. We utilize MATLAB to process the structural imaging files, which includes patient data de-identification (R2). The resulting output contains a structural image to localize the voxel sample, the spectral graph, quantified metabolites, and associated metadata; these data remain semantically linked in the visual tool. We use a custom data format because the DICOM standard is not universally or consistently adopted for MRS data.

Visual Inspection of Voxel Positioning and Spectral Graphs. Following data loading (Fig. B.3A), the spatial voxel overview panel (Fig. B.3B) is used to review the spectral graph, associated anatomical image, and included metadata for each acquisition. This panel consists of a set of images for each patient. In each structural image, a fuchsia rectangle indicates the voxel sample region for the MRS acquisition (Fig. B.3B1). To the left, a position selector consists of small filled nodes, each of which indicates an acquisition for the selected patient. Using the standard CPK color convention for atomic elements [92], we represent ^1H spectral metabolites with a white-filled node. A light gray bar behind the disks shows the active selection image, while the node becomes filled in fuchsia to indicate image linkage to a spectrum that is selected in the spectral heatmap panel (Fig. B.3B2). Users can access different images via these position nodes or time acquisition nodes (horizontal axis). A selected node shows the structural image with localized voxel, associated spectral graph, and supplemental metadata, such as TE setting, patient age and gender, stored with that voxel (R4). These data are stored hierarchically, where each voxel sample with spatial information is first sorted by individual identifier and associated metadata, then by time of acquisition, and finally

by brain activity state during the acquisition.

Group Creation and Metabolite Selection. Following visual inspection of voxel position and spectral graphs, the user may then create custom groups of spectral voxels for subsequent analysis (**T3**) in the Voxel Group Overview panel (Fig. B.3C). Custom groups may be edited at any time. Membership in a custom group is listed in the metadata table at the bottom right region of the interface (Fig. B.3G). Our application additionally creates preset groups for each echo time, spatial region, individual, brain state, and time point. These may be immediately accessed in a drop-down list in the Metabolite panel (Fig. B.3D).

Following a group selection from the metabolite drop-down list, all quantified metabolites from the offline processing step are displayed. Users then have the option of adding all metabolites in the list to the x-axis, y-axis or both axes of a spectral ratio heatmap located to the right of this list (Fig. B.3E). Alternatively, only a subset of metabolites may be added to the heatmap axes. Groups may be flexibly added or removed from either axis at any time. Metabolites populate along heatmap axes in alphabetical order; we discussed a number of ordering options with our domain collaborators, settling on this ordering method for consistency and pattern recognition between studies.

Ratio Exploration. Following loading of metabolite groups onto each axis, we determine ratios for all metabolite permutations for display in the heatmap panel (Fig. B.3E). This serves as the primary visualization component of our tool, as shown in Fig. B.3E and which is described in detail in Section B.6. In this view, users can compare average (Fig. B.3E1) or individual metabolite ratios at different levels of detail (**R5**). Users may interactively expand a cell to reveal key attribute details (Fig. B.3E2), as inspired by Bertifier [395]. The background of the cell remains visible behind individual ratio elements for all expansions to preserve context of the aggregated value during navigation. This subtle context preservation was deemed useful by experts in our development process.

A legend at the far right (Fig. B.3F) serves to indicate hue and glyph meaning. Hovering over a cell or glyph correspondingly highlights linked data elements in fuchsia, including the associated spectral graph, patient anatomical image, and associated metadata (**R4**), as depicted in Fig. B.3G.

B

B.6 Spectral Ratio Heatmap

In the heatmap panel we divide MRS data elements into tiers of visual priority (**R3-R5**):

Tier 1 Quantified spectral data

Tier 2 Derived spectral data

Tier 3 Spectral metadata

Tier 1 has primary importance; it consists of relative metabolite concentrations which are the result of pre-processing and quantification steps from the raw spectral acquisition. **Tier 2** comprises the complete set of metabolite ratios. It is used for comparison between user-defined groups as well as the following key attributes: spatial region, individual, brain activity state, and time point. Spatial region indicates the

voxel sample position within the brain. Individual refers to a given patient included in the analysis. Time indicates either the number of separate spectral acquisitions performed on an individual over a study period, as in a longitudinal study, or recorded metabolite values within an acquisition session, as in a time-resolved MRS study. Finally, brain activity state indicates if the subject was in an *active* (task-explicit) state or *resting* (task-negative) state during signal acquisition. **Tier 3** includes metadata important for context and selection that are unnecessary to include as explicit encodings in the visualization: gender, age, and acquisition settings can have varying impact on the resulting concentrations and ratios of metabolites [475, 568].

Tier 1 Encoding. Visual perception research has shown that encoding position along a common axis is the most effective visual channel for communicating quantitative information [90]. Box plots are a simple, ubiquitous and descriptive means of visually encoding statistical information about a dataset [556]. Since each MRS spectrum is essentially a multivariate set, where each metabolite is a variable, each metabolite in the spectrum then is tied to its own set of unique statistical information. We chose box plots over violin [198] or summary plots [400] to visualize **tier 1** data, as our goal with this tier is to provide clean, quickly readable insight to the input value range. Our use of box plots is additionally inspired by Blumenschein et al. [52], who used bars to encode aggregate dimensions in their work on table visualization. Bars and box plots are additionally well-recognized and easy to interpret; use of elements that were familiar to our target user group was an important design consideration. Furthermore, since box plots are only applicable when a dataset consists of five or more members, we introduce three variations depending on the number of inputs as illustrated in Fig. B.4. For any of these variations, we first flatten the voxel hierarchy described in the spatial overview panel, and split the data into one voxel array per axis. In each array, we calculate the mean for each metabolite. In the case of a single spectral input, we use design variation A, which utilizes bars only, where height encodes the concentration of each metabolite (Fig. B.4A). We calculate median, minimum, and maximum for two or more metabolite values on an axis. This corresponds to variation B, where height encodes the median value and whiskers encode the minimum and maximum metabolite concentration value, respectively (Fig. B.4B). For five or more metabolites on an axis we additionally calculate the interquartile range. The box and whisker plot in variation C is utilized in this case, and shows the median, first and third quartiles, and the minimum and maximum value (Fig. B.4C).

Tier 2 Encoding: Overview. In **tier 2**, we visualize ratios between the mean along the x- and y-axes in a heatmap matrix (**R5**), as shown in Figure B.3E. This effectively trades the low spatial resolution of MRS data for abstract resolution, focusing on biochemical concentrations in detail for a small region of interest. Each cell shows the aggregate ratio of the metabolite on the x-axis position to the corresponding y-axis metabolite, for instance mean Glutamine (Gln)/mean N-acetylaspartate (NAA), as illustrated in Fig. B.5. We map the ratio value to a diverging red-blue colormap [185] inside each heatmap cell, as this color scheme is a familiar sight to our collaborators. In instances where the ratio is less than 1, we invert the ratio and switch the sign. To obtain a cleanly symmetric, divergent mapping structure we drop all values by 1 so the diagonal of the heatmap matrix is 0, rather than 1. Our aim is to draw attention to large input differences; this was identified as important for spectroscopy researchers. Red indicates a higher x-axis metabolite input while blue indicates a higher y-axis metabo-




Figure B.4: Box and bar plots encode metabolite concentrations (tier 1 visual encoding). (A) utilizes bars where height encodes the concentration of each metabolite for a single spectral input. (B) is utilized for two to four spectral inputs on an axis, where height encodes the median value and whiskers encode the minimum and maximum metabolite concentration values, respectively. A box plot (C) is utilized for five or more spectral inputs on a given axis.

lite input. Equivalent inputs map to white. If an input value is 0, we map the cell color to dark grey. We originally thought to exclude such values from the heatmap, but on further discussion with our collaborators felt these were useful to include in order to preserve context. This heatmap view provides a means to visualize otherwise undetectable patterns in a rapid overview. To aid color interpretation and perception, our application includes a colormap legend to the right of the heatmap (Fig. B.3F).

B

Tier 2 Encoding: Attributes. Through a series of group interviews and individual shadowing sessions to the MR scanners we identified that, following an overview of all aggregated metabolite ratios, researchers are most interested in *comparing and summarizing* (T4) individual metabolites. For a given metabolite ratio, researchers first are interested in comparing *brain spatial regions*, as this can provide the most context for understanding ratio differences, e.g., in a tumor cohort study where voxels are acquired in the tumor region and in a healthy region of the brain. With spatial context, researchers can easily compare ratios between *individuals*. Assessing *brain activity state* is then most relevant in the context of the individual. After comparing the difference in active vs. resting brain state for an individual, the researcher may review the difference in these values over a cohort. Similarly, *time points* are best assessed first within a given brain state, then between states of an individual, before comparing between individuals.

In order to support experts in better identifying unexpected source ratios in a study, they thus need to evaluate four key attributes: (1) *brain spatial region*, (2) *individual*, i.e., *patient*, (3) *brain activity state*, and (4) *time point*. Furthermore, through each of these analysis stages we found that researchers prefer to maintain context between attributes to better understand sources of variation. This helped drive our development of a detailed metabolite ratio view that nests within each heatmap cell. Many MRS studies, particularly proof-of-concept research studies, by our collaborators often include around 20 subjects. They may sample up to four brain regions (although two is more typical), include up to three time points, i.e., pre-operative, post-operative, and long-term follow-up, and measure either a single or dual brain activity state. This space of attributes and approximate study size produces a set of 16 possible case scenarios to account for in our detailed comparison view.

Tier 2 Encoding: Detail. Given the low number of key attributes, we found a simple glyph representing each attribute to be the most conducive to user analysis. Our glyph




Figure B.5: In nested ratio calculations, the cell background (A) is first mapped to color based on the average of input metabolites on the x-axis divided by the average of input metabolites on the y-axis. Within the cell (B), the value of each input metabolite for all patients, at all time and brain state collections, is averaged and compared as a ratio for each spatial region. Within a spatial region, the average of each metabolite is compared for each patient, then for the brain state of each patient. The innermost step, time, takes a single metabolite input for both the numerator and denominator.

choice and design was mainly inspired by findings from unit visualization research, mainly the ATOM grammar by Park et al. [390], for this method's demonstrated strong intuition and interaction properties. Since our target study sizes are typically relatively small, we avoid issues with display and perceptual scalability from which unit visualizations often suffer. To maintain important context in the analysis flow, we nest glyphs to mirror the order of analysis preferred by researchers. Our glyph nesting design is inspired by dimensional stacking visualization techniques pioneered in XmdvTool and N-land by Ward et al. [542, 543]. Since nested glyphs can form complex shapes, we chose glyphs that were simple and familiar to our collaborators to reduce interpretation difficulties. Although we discussed different stroke styles for glyphs, for simplicity and clarity our ultimate design uses a solid hairline stroke for each of the four attributes. Experts felt that changes in stroke weight or style was distracting and overemphasized elements; this may bias conclusions.

The visual design for this detailed view is mapped from a series of nested ratios. Inside each cell we flatten the data to a single voxel array, skipping any duplicate voxels. We then determine ratios for each of the key attributes, where available, in a nested fashion that mirrors the preferred order of user analysis: the ratio for each spatial region (using the average of all individuals for this region), each individual (using the average of all states for the given individual in a given region), each state (using the average of all time points for a given state of an individual from a given region) and each time point, as shown in Fig. B.5B. These nested values then map to the appropriate glyph.

We represent spatial regions as rounded rectangular glyphs. We chose rounded corners to distinguish spatial glyphs from the square shape of the heatmap overview cell. Furthermore, the rounded corners leave space to reveal the heatmap cell color, thereby subtly preserving context within the detail view. In each cell, we evenly divide the space vertically by the number of distinct regions sampled. Individuals are presented

Case	region	individ.	state	time pt	visual	Case	region	individ.	state	time pt	visual
1	single	single	single	single		9	single	multiple	single	single	
2	single	single	dual	single		10	single	multiple	dual	single	
3	single	single	single	multiple		11	single	multiple	single	multiple	
4	single	single	dual	multiple		12	single	multiple	dual	multiple	
5	multiple	single	single	single		13	multiple	multiple	single	single	
6	multiple	single	dual	single		14	multiple	multiple	dual	single	
7	multiple	single	single	multiple		15	multiple	multiple	single	multiple	
8	multiple	single	dual	multiple		16	multiple	multiple	dual	multiple	

Figure B.6: Key tier 2 visual attributes include: brain spatial region, individual, brain state, and time point. We assign a unique glyph to each of these four attributes. Brain state is defined as active or resting; in absence of a classification we assume a resting state. All remaining three attributes may have single or multiple recordings. This produces 16 possible scenarios for spectral analysis. A sample visual is included for each scenario.

B

as filled disks when only shown in a single time acquisition (e.g., case 9), expanding to rounded squares when time series data are incorporated (e.g., case 3). This shape change permits a spark line to move evenly across the space without going outside the border of the enclosing glyph. Shapes scale to fill space within their frame. In instances where different brain activity states are analyzed, we divide the shape in half horizontally (e.g., case 2). This feature was important to include for our collaborators who perform time-resolved spectroscopy, as this is not available in other tools. Finally, we encode different time acquisitions as points connected via a spark line, inspired by Meyer et al. in their work, Pathline [341]. This spark line is nested into the relevant glyph: if a multi time step series is captured in a study analyzing different brain states, the spark line is placed within each state half-moon glyph (e.g., case 4). If analysis is only for a single activity state, the spark line nests inside the individual glyph (e.g., cases 3, 11), or inside spatial glyphs for a single patient (e.g., case 7). The remaining cases comprise different permutations of these spatial region, individual, brain state, and time point arrangements.

For example, consider an instance of scenario 16: two patients are sampled in two regions of the brain four times in a year. During two acquisition times the subjects were asked to perform a task (active brain state), while the other two times were asked to relax (resting brain state). This produces a total of 12 unique measurements, 6 per patient. The overview cell is calculated by averaging the 6 values of Gln for patient 1 and the 6 values of NAA for patient 2, and dividing the result of NAA into Gln. Inside the cell, we compute this ratio as a series of nested averages for each of the four key attributes, as depicted in Fig. B.5: (1) spatial region, (2) individual, (3) brain state, and (4) time point. For each, we average the metabolite concentrations before computing the ratio. For additional detail view images and example tasks of each scenario in a more complex dataset, we refer interested readers to the supplementary material *SpectraMosaic Detail Case Scenarios*.

Hovering facilities display the ratio value for each cell or attribute of interest (**R5**). Displaying this numerical value provides a safeguard against possible distortions of color perceptions that may occur with our chosen glyph nesting structure. This value is displayed in red text if one metabolite input exhibits an uncertainty above 15% (Cramér-Rao lower bound) [82]. This information may be used to assess both the quality of the measurement and the accuracy of the spectral processing and quantification steps.

Tier 3 of MRS data consists of metadata information used for context and selection. We depict this information in a table below the heatmap. Gender, age, and echo time comprise other important patient attributes to track because the shape of the spectrum can vary considerably with these factors—for example, the lactate peak is virtually undetectable in healthy babies [568], but is nearly always measurable in healthy adults with increased neural activation [475]. Acquisition settings are also important, as different echo times will yield a vastly different spectral representation for the same patient.

B.7 Implementation

SpectraMosaic is a web-based application implemented with HTML, CSS, Javascript, as well as the D3 [60], P5, and gridster Javascript libraries. It was developed as a web application to allow for easy integration and use within the hospital network (**R1**). A Python back end integrates MATLAB [329] and Tarquin [557] components in the pre-processing steps. Assets are stored on the client and fetched on-demand. Our visualization tool code is open source and is publicly available at <https://github.com/mmiv-center/spectramosaic-public>.

B.8 Case Study

We evaluated the utility of SpectraMosaic as a research tool using a giardiasis MRS case study. Giardiasis is a parasite-borne disease affecting the small intestine caused by drinking water contamination. The metabolic byproducts of this disease are subtle, but have been shown to be detectable by MRS [117, 526]. The goal of this study is to explore and identify possible metabolic indicators for infection using our tool.




Figure B.7: Heatmap inspection in a two-patient, multi-voxel study acquired at two TEs: 35 ms (x-axis) and 144 ms (y-axis). Investigating the Alanine (Ala)/Alanine (Ala) cell reveals a higher measurement in the TE 144 ms group. However, the tooltip indicates that this ratio may be unreliable due to a poor model fit.

Collected in Bergen, Norway, study data comprised two patients imaged some months apart in three different regions of the brain at a single echo time (TE 35 ms). For one region (prefrontal region) two different TE parameter settings were used (TE 35 ms and TE 144 ms). These data were analyzed by three volunteers recruited from the fMRI/MRS research group in Bergen. All three provided feedback on earlier interfaces of the SpectraMosaic application, and are not co-authors of this work. User A is an MR physicist specializing in development and refinement of spectroscopy protocols for clinical studies of neuropsychiatric and developmental disorders. User B, also an MR physicist, uses ^{31}P and ^{13}C -labeled pyruvate timecourse data to study real-time metabolism. User C is a cognitive neuroscientist who uses MRS in conjunction with fMRI in research on neurodegenerative and developmental disorders, e.g., Parkinson's disease, stroke, and stuttering. We processed the data in advance to focus evaluation on the visual web tool; this step included de-identification of patient-specific information.

B

Case Workflow Feedback. After a brief introduction to the tool, users analyzed this case following a “think-aloud” protocol [298]. We conducted follow-up interviews after the analysis was complete, which we summarize and discuss in this section.

All three users began with an overview of the spectral graphs and voxel position for each imaged brain region (T1). User A investigated spectral graphs by region, while B and C explored by patient. Users A and B commented that this overview provides an important quality assurance check for each acquisition. Since all three users are familiar with MRS, they agreed with our decision to exclude labeling of spectral peaks; they felt this would have been unnecessary and distracting to include. All commented that the hippocampus region spectra looked strange, which could be indicative of either a pathology or acquisition problem. They noted that this region is particularly difficult to image well, and requires deeper investigation.

All users then explored available group presets and experimented with creation of custom groups for analysis (T2). They agreed that the presets particularly improved the practical usability of the tool, stating that these were comprehensive and largely removed the need to make custom groups. All users experimented with adding a subset of basis set metabolites (Fig. B.3D) to the heatmap view, although they felt that analysis of all metabolites is a useful first step for exploring new hypotheses. However, they agreed that subset metabolite analysis is useful as hypotheses are refined to a narrower metabolite set.

Feedback was positive for the alphabetical ordering of metabolites on heatmap axes. User C strongly felt that any statistics-based ordering method would make interpretation too difficult because they would spend too much time locating metabolites along the axes. All users agreed that the representation of metabolite relative concentrations as whisker bar or box plots was extremely useful, as it offered additional insight into unexpected values observed in the heatmap. User B stated: *“Checking the range on the metabolite inputs helps me as a first check; a huge range could indicate a [brain region] area effect or a bad acquisition. I can easily then verify this by checking the spectral graph in the other panel.”* All users noted a massive range for Gamma-Aminobutyric acid (GABA) in this study, and were able to quickly conclude that the acquisition technique used is not effective for this metabolite. For this study and others acquired on the same scanner, through the same technique and parameters, this representation allows for a straightforward relative comparison of metabolites before ratio computation (T3).

In the spectral ratio heatmap, user A was primarily interested in exploring ratios at different echo times (TE) (T4). We see this exploration in Fig. B.7; TE 144 ms voxels are placed on the y-axis while TE 35 ms voxels are placed on the x-axis. This user focused on the diagonal of the matrix, and primarily on examining known metabolites implicated in giardia infection, e.g., Alanine. Although this ratio shows relative similarity, we note that the model fit for this metabolite is outside the accepted range. This requires further investigation. User B also compared different echo times, but over the entire matrix space for any unexpected dark color regions. For each unexpected cell, the user noted whether this could be pathology, or an acquisition problem.

All three users were also interested in comparing ratios of metabolites between patients for each of the three measured brain regions (T4). They first filtered out TE 144 ms acquisitions, then arrayed each patient on opposing axes. Assuming both patients are healthy, we would expect that the patient glyphs for all spatial regions would show similar values. All three users noted an unexpected, relatively large difference for Lactate/Total Creatine in the hippocampal region (Fig. B.8). To investigate this disparity, users A and B first verified whether the value met the threshold for each patient. The value did not meet the threshold for the female patient, indicating an unreliable fit. Users then reviewed the spectral graph of the hippocampal acquisition for this patient, noting its abnormality—users concluded that this merits deeper investigation, and likely requires an new acquisition.

Summary Feedback. All three participants felt that SpectraMosaic was useful and could augment their standard workflow for deeper insights into spectral data. User A noted that the visual feedback on the model fit in each ratio provides invaluable data quality information. User B stated, *“The linking between the glyph ratios, the spectra, the table, and the images is incredibly useful for us—whenever we look at metabolite*

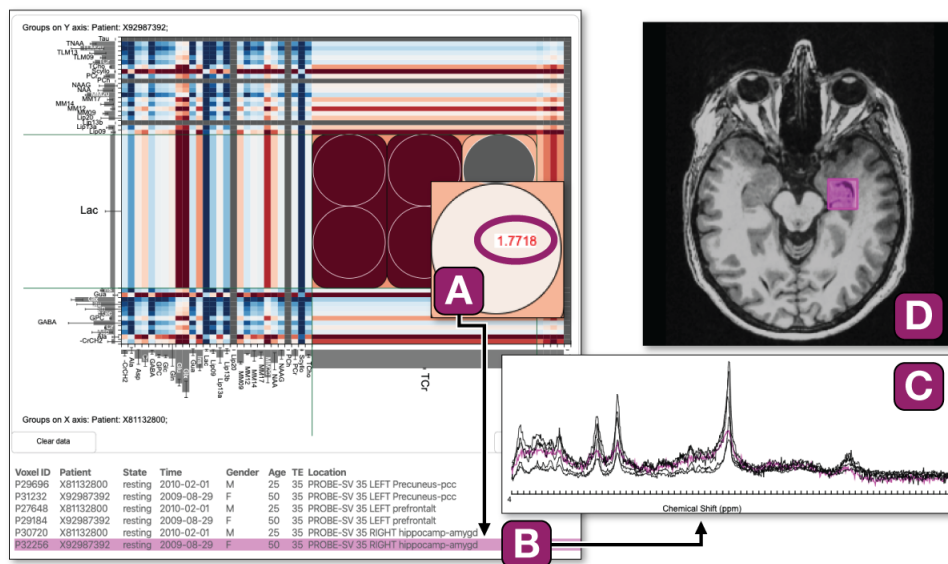


Figure B.8: Following overview of the metabolite ratios between the individual patients, the user inspects Lactate (Lac)/Total creatine (TCr) ratio between two patients for all three regions at TE 35 ms. The user notes a high lactate measurement for the female patient relative to the other measured regions (A). Subsequent inspection of spectral metadata (B), the spectral graph (C), and the brain region in which this measurement was acquired (D) help establish reasons for this difference.

results we always want to go back to the raw spectra and see if this makes sense, and if the quality is good, and this makes it really easy to see. I see this tool as being useful to verify assumptions I have going in to the study, and to explore the entire range for quality checks that might affect the results that I'm expecting."

B

All participants felt the nested glyphs were integral elements of the metabolite ratio exploration process. The detail glyph view provided a means to quickly drill into an unexpected ratio and identify the possible source(s), while easily retaining contextual information from the surrounding heatmap. User C noted: *"This [spectral heatmap] overview and detail glyph feature is useful to have a closer look at, for example, neurodegeneration [in Parkinson's] with the loss of dopaminergic connections, as seen with concentrations of glutamate or GABA... and it is ideal for testing new protocols against established protocols."* Furthermore, experts agreed that the glyph design and nesting structure was intuitive and clear in all case scenarios, even in larger, more complex studies. All three stated that interpreting these glyphs was not difficult, particularly when compared to the very steep learning curve to interpreting spectral data through their standard approach. They agreed that the inclusion of the legend was helpful when first familiarizing themselves with the system, but that they had little need to reference it after the first few minutes of heatmap exploration. However, two experts commented that our mapping of vertical time points could be scaled differently to more clearly demonstrate relative ratio value changes, which were at times difficult to recognize. For detailed expert feedback on clarity and interpretability of the nested glyph

structure in all 16 possible case scenarios in a larger dataset, we refer interested readers to *SpectraMosaic Detail Case Scenarios* in the supplementary material.

All users indicated interest in an option to extract spectral heatmap visuals and data for subsequent statistical analysis; user A expressed interest in seeing this output to the hospital PACS for access by radiologists to aid in more rapid interpretation of spectroscopy data for more widespread clinical use.

B.9 Discussion and Limitations

In the case evaluation of SpectraMosaic we found that our tool provides new, interesting insights on metabolic profiles at different aggregation levels.

Our task analysis showed that experts were particularly interested in large metabolite differences. Although our diverging color mapping approach in the heatmap is effective in demonstrating large differences between metabolites, subtle differences are less obvious. Investigation into fine grained color mapping options or user-defined color map scaling may help more clearly highlight these instances. This extends to our plotting of time points, where subtle ratio changes could benefit from a logarithmic axis scaling approach to highlight such changes to users.

While our decision to sort metabolite inputs in a consistent order limits the ability for pattern recognition *within* a study, this approach allows for pattern recognition between studies, where users can begin to observe a typical “footprint” for certain acquisition techniques.

Although this is uncommon for our collaborators, we also note that if data are not acquired from the same scanner and same parameters, the utility of the bar and box encoding becomes more limited. This is because different scanners and different parameters can vastly change the metabolite concentrations; in this case the ratio heatmap becomes the primary tool for comparative analysis.

Our visual design, particularly with reference to the nested glyphs in the detail view, was guided by collaborative discussions with research experts. These relatively small study sizes are conducive to nested unit visualizations, and in this iteration of the application were not designed to scale to, e.g., hundreds of patients. With respect to the scalability of groupings within our planned design, we conducted a preliminary assessment of nested glyph interpretability for each case scenario using a larger study. We provide the results of this assessment in the supplementary material (*SpectraMosaic Detail Case Scenario*). Our collaborators even indicated that they could envision this approach scaling beyond 20 patients for some scenarios. Additionally, we could incorporate clustering or an additional design layer for automatically- or user-generated groupings for further scalability (T4).

Lastly, while our glyph system covers all main use cases, we found that echo time is varied in research studies more often than initially expected. This frequency of use may imply that this attribute should be encoded at the second priority visualization tier, rather than its current third level. However, comparison of different echo times beyond an overview level is of less clinical interest than the four attributes we have discussed. Inclusion of a fifth glyph would require careful consideration.

B.10 Conclusions and Future Work

In this design study we contributed a characterization of the data, task, and design requirements for the development of SpectraMosaic, followed with an expanded tiered visual encoding system and pipeline. We performed case studies with three domain experts to validate our tool in spectroscopy clinical research and protocol development. MR spectroscopy is a ripe area for continued visualization research.

The flexible design of our tool allows for a number of possible extensions; this may include investigation into additional statistical measures relevant for comparative analysis, e.g., correlation. Although this paper focuses on ^1H -MRS, ^{31}P -MRS and ^{23}Na -MRS analysis may also be integrated to our tool. While we offer basic mechanisms for uncertainty visualization, exploring additional means for uncertainty feedback in the heatmap cells and glyphs can offer deeper insights into the data. Finally, although typical MRS cohort studies are relatively small, exploration of methods to extend our visual encoding system to successfully manage larger cohorts may further increase tool usability.

Automatic interface adjustment based on acquisition technique offers a valuable investigation of parameter space analysis in MRS. Exploration of the most salient features to reveal for, e.g., PRESS versus MEGA-PRESS, may help experts more effectively identify interesting ratios for further investigation. Beyond the medical domain, an additional interesting line of inquiry would be to explore the adaptability of our abstracted tasks paired with our visual encoding system in other areas facing similar challenges with heterogeneous multidimensional data, such as meteorology or geophysics.

Acknowledgments

We thank the UiB VisGroup and our collaborators at the Mohn Medical Imaging and Visualization Centre (MMIV) and Haukeland University Hospital: Lars Ersland, Karsten Specht, Frank Riemer, Hauke Bartsch, Kenneth Hugdahl, and Gerard Dwyer for their invaluable insights and feedback. This work is supported by the University of Bergen and the Trond Mohn Foundation in Bergen, Norway (#813558, Visualizing Data Science for Large Scale Hypothesis Management in Imaging Biomarker Discovery; and #811255).

Paper C

DimLift: Interactive Hierarchical Data Exploration through Dimensional Bundling

Laura Garrison^{1,2}, Juliane Müller³, Stefanie Schreiber^{3,4}, Steffen Oeltze-Jafra^{3,4}, Helwig Hauser^{1,2}, and Stefan Bruckner^{1,2}

¹ Dept. of Informatics, Univ. of Bergen, Norway

² Mohn Medical Imaging and Visualization Centre, Dept. of Radiology, Haukeland Univ. Hospital, Bergen, Norway

³ Dept. of Neurology, Otto von Guericke University Magdeburg, Germany

⁴ Center for Behavioral Brain Sciences, Otto von Guericke University Magdeburg, Germany

Abstract

The identification of interesting patterns and relationships is essential to exploratory data analysis. This becomes increasingly difficult in high dimensional datasets. While dimensionality reduction techniques can be utilized to reduce the analysis space, these may unintentionally bury key dimensions within a larger grouping and obfuscate meaningful patterns. With this work we introduce *DimLift*, a novel visual analysis method for creating and interacting with *dimensional bundles*. Generated through an iterative dimensionality reduction or user-driven approach, *dimensional bundles* are expressive groups of dimensions that contribute similarly to the variance of a dataset. Interactive exploration and reconstruction methods via a layered parallel coordinates plot allow users to *lift* interesting and subtle relationships to the surface, even in complex scenarios of missing and mixed data types. We exemplify the power of this technique in an expert case study on clinical cohort data alongside two additional case examples from nutrition and ecology.

Dimensionality reduction techniques are frequently utilized to reduce the complexity of high dimensional data by projection to a lower dimensional space. However, when used alone and monolithically, these techniques can emphasize strong, uninteresting patterns in the data and hide important variations. For example, although cardiac risk is well-known to correlate with waist measurement, a more interesting, though subtle, relation to gender or smoking may be relevant for a clinician to see. Visual analytics leverages the strengths of powerful statistical tools, including dimensionality reduction, in tandem with user knowledge. However, while some visual analysis tools have been developed to create expressive dimensional groupings, they do not easily allow for incorporation of user knowledge for faceted hypothesis generation. Furthermore, connecting the results of the dimensionality reduction back to the original data for interpretation and relation to subsequent steps, e.g., decision making, can be difficult.

For instance, in clinical cohort studies medical experts are chiefly interested in untangling interesting and relevant measures of a given disease, e.g., cerebral small vessel disease (CSVD), for diagnostic purposes. Biomarker discovery is a complex and challenging process, and dimensionality reduction techniques provide a means to reduce the analysis space. However, these techniques may produce groupings that are not interesting to the expert for particular subcohorts, e.g., a specific gene expression level grouped with test results for young patients. Our method, which utilizes iterative dimensionality reduction to extract subsets of dimensions that contribute similarly to the variation of a dataset, allows for flexible user-driven restructuring of subcohorts and subsequent groupings to support exploratory hypothesis generation. For example, adjusting the previous subcohort to include middle-aged patients with high blood pressure may be done to support a new hypothesis that high expression of a particular gene in combination with a certain range of test scores, such as high blood pressure, can act as a set of indicators for CSVD in middle-aged patients. Similar such scenarios occur in many areas of science and engineering. These domains are interested in exposing patterns in subsets of large, complex populations, and benefit from this style of visual reasoning.

The rapid identification of interesting patterns and relationships is key to the analysis of complex high dimensional data. Achieving this requires effective integration of statistical methods with user knowledge to reduce the space to salient dimensions. Core to our approach is the concept of *dimensional bundles*: statistical- or user-driven groupings of dimensions that are accessible as a unit or at the component level. Our statistical approach utilizes factor analysis of mixed data (FAMD) [289], a dimensionality reduction technique applicable to complex, mixed-type data. We run this algorithm in multiple iterations over the data; each iteration captures and extracts a set of dimensions, so called *dimensional bundles*, which contribute similarly to the variance within the dataset. This avoids a monolithic treatment and instead produces hierarchical bundles of dimensions that retain the expressivity of the original dataset. While previous approaches to dimensional grouping have focused on clustering or dimensionality reduction methods that converge to an ideal representation of a high dimensional dataset [269], our approach facilitates dynamic visual navigation and composition of high dimensional data to *lift* subtle, interesting features to the surface. A simple example of *dimensional bundle* restructuring is shown in Fig. C.1, where we explore the relationship between cardiac risk and gender.

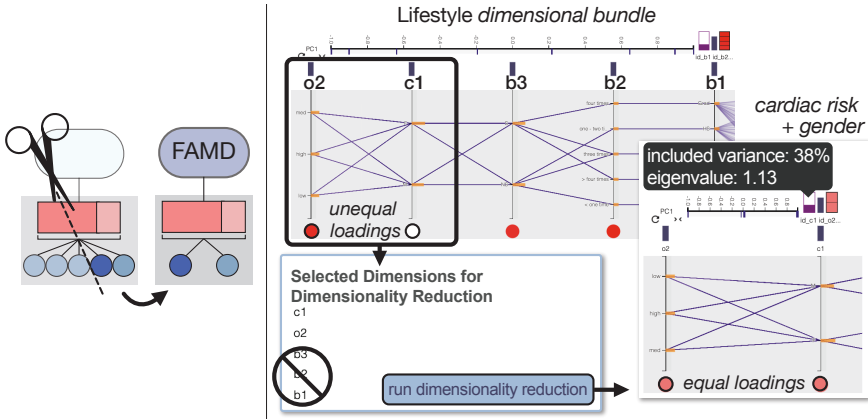


Figure C.1: We inspect a *dimensional bundle* comprised of lifestyle dimensions, e.g., education (b1), workout frequency (b2), smoking (b3), cardiac risk (o2), and gender (c1). We suspect a correlation between cardiac risk (o2) and gender (c1), so then *lift* these dimensions to better **target and test our hypothesis** by removing all other dimensions from this bundle. With an eigenvalue above 1 and changes in contributions/loadings indicated by hue at the bottom of the axes, we note a **subtle correlation** that was previously undetectable. This is conceptually illustrated on the left.

Expressive visual mapping techniques from visual analytics have been developed to represent dimensional hierarchies, e.g., parallel coordinates plots with embedded cluster diagrams [207]. However, these solutions typically expect data in a pre-established hierarchical structure, or offer limited interaction methods for restructuring groups during analysis. By contrast, our method proposes to semantically connect and track data transformations through visual mappings and interactions that allow on-the-fly recomposition of *dimensional bundles*. This provides a flexible solution to swiftly adapt perspectives on high dimensional data with the potential to rapidly identify relevant relations, even when overshadowed by well known or less interesting trends. Our concept extends to complex mixed-type and incomplete data. Following a review of related work and a description of our methodology, we demonstrate the power of our approach in the context of three scientific datasets, one of which is a case study with domain experts in clinical neurology.

C.1 Related Work

Visual analysis of high dimensional data is a grand challenge in the visualization research community, with applications across numerous domains. Discussion of efforts in this general area are beyond the scope of this paper, but are detailed with a survey of advances in recent years by Liu et al [307]. Our work expands on the idea of simultaneous dimensions and items analysis for exploratory hypothesis generation, described as the Dual Analysis approach, by Turkay et al. [510]. This work, along with a follow-on clinical application study [511], describes a visual analysis workflow where users seamlessly move between analysis of dimensions through comparative de-

scriptive statistics and item comparison to identify outliers and correlations of interest. Brushing and linking mechanisms provide clear visual feedback during the analysis process. This approach has since been extended to incorporate mixed data (continuous and categorical) with facilities for visualization of missing data by Müller et al. [351]. DimLift expands further on the reciprocity between dimension and item space by introducing *dimensional bundles* for analysis of high dimensional data. The SIRIUS system, presented by Dowling et al. [110], explores the interplay of dimension and item space while incorporating a nonlinear dimensionality reduction technique. This approach shows MDS projections for both dimension and item space in linked views to demonstrate correlations in high-dimensional data. Our approach similarly utilizes dimensionality reduction to aid correlation exploration of high-dimensional data, but adopts a linear technique to better track between the original and newly-produced dimensions.

Dimensionality reduction is used ubiquitously in visual analytics. Sacha et al. [429] provide an overview and classification of dimension reduction methods as used in visual analysis. Our work incorporates dimensionality reduction into a subset of the interaction scenarios they identified: data selection & emphasis, data manipulation, and feature selection & emphasis. Similar works in this space include the work of Tatu et al. [493], who utilize an interestingness-guided subspace search algorithm to identify subspace sets for subsequent visual analysis, although tools for user-driven subspace composition are limited. The DimStiller workflow by Ingram et al. [215] guides users through the dimensionality reduction process to find a single global optimal composition; our approach by contrast does not emphasize a single global optimum, and is designed for a variety of complementary perspectives onto relations between relevant subsets of the dimensions. Yuan et al. [576] combine a Dimension Projection Matrix, an extended scatterplot matrix, with a Dimension Projection Tree to explore data and dimension subspaces. Our approach tackles a similar goal of dimensional subspace analysis at both item and dimension levels with related interactions. However, our visual approach enables direct correlation comparison between multiple dimensions and items in a parallel coordinates view, and is targeted specifically at user-driven hypothesis exploration.

Although dimension reduction methods project relevant data features into low dimensional space, the results are often difficult to comprehend. Principal component analysis (PCA) [394], although a well-known and broadly applicable method utilized in dimensionality reduction, suffers from this interpretation gap. Müller et al. [357] present a general discussion of design solutions to clearly visualize the connection between data inputs and results from PCA. However, many of these solutions do not scale well with high dimensional data. Our visual interactive approach offers one method for bridging this intuitive gap into high dimensional data spaces. iPCA [235] is one other such solution designed to connect PCA results to source data. It uses multiple coordinated views to depict PCA results with interaction facilities for the user to adjust dimension contributions within any principal component—any adjustments update visuals for the final PCA results. Our approach similarly uses visual elements and interactions to connect the raw data to the results of a linear dimensionality reduction method, but rather than using visualization to understand the semantics of PCA, our approach uses similar results as a tool in hypothesis formation.

Parallel coordinates are a well-known method for representing multidimensional data [189]. Nested or hierarchical plots, adapted from the traditional flat parallel coordinates plot, are used to visualize and evaluate structural relationships of the data. Numerous solutions present data aggregation by item relatedness into parallel coordinates as a means to reduce clutter and noise in the plot [24, 80, 143, 144, 418, 518, 533]. Each of these methods focuses on the hierarchical construction of sets of items, while our approach aims one level above this on the hierarchical construction of sets of dimensions. Several approaches have utilized parallel coordinates to visualize dimension-level aggregation, created either through algorithmic methods or pre-defined data hierarchies. These methods provide varying degrees of interaction to the user. For example, Wang et al. [540] and Dunica et al. [112] use parallel coordinates to visualize results of a single-run PCA, where each axis represents a different principal component. While we similarly incorporate principal components, we instead take an iterative algorithmic approach to produce principal components of selected subsets of particularly related dimensions. These subsets form the *dimensional bundles* in our method.

Approaches to **parallel coordinate dimension hierarchies** often incorporate other views on the data, integrated either separately or directly into the parallel coordinates. Huang et al. [207] create hierarchical clusters of dimensions in parallel coordinates using dendrograms that attach to each axis. DOFSA [571] and InterRing [572] are connected tools that allow interactive visual exploration and modification of hierarchical data. These modifications are made on InterRing and linked to other panels, e.g., parallel coordinates. By contrast, our method does not divide user attention over different graphical interfaces. Furthermore, the DOFSA hierarchy itself is flattened in parallel coordinates, and its order is informed by the hierarchy constructed in InterRing. Our approach does not flatten the hierarchy in this manner. Weidele [548] recently presented the conditional parallel coordinates method, which ties and reveals additional dimensions to the range of a given parent dimension only if certain conditions are met. Perhaps most similar in principle to our visual approach, Brodbeck & Girardin [70] and Andrews et al. [18] create aggregated dimension axes for parallel coordinates plots, which may then be expanded to reveal the contained dimensions. In contrast to these methods, we do not expect pre-defined hierarchies, instead allowing flexible regrouping as hypotheses evolve.

C.2 DimLift Approach

Key to complex high dimensional data analysis is the rapid identification of interesting dimensions. While dimensionality reduction is a core tool for high dimensional data analysis, nonlinear methods do not allow for an easy link back to the original dimensions, which Sedlmair et al. [453] identify as key tasks for users interested in finding important original dimensions (as opposed to purely gaining insights on the dataset structure). Sedlmair et al. also identify the need of users to compare, or unmap, original dimensions to newly-created dimensions; nonlinear methods are also of limited use for this task. Lastly, Sacha et al. [429] identified user interactions as critical components of an exploratory visual analysis pipeline utilizing dimensionality reduction. While numerous solutions in this space have incorporated a human into the loop, many aim to help the user to better understand the results of the algorithm, or to guide the user

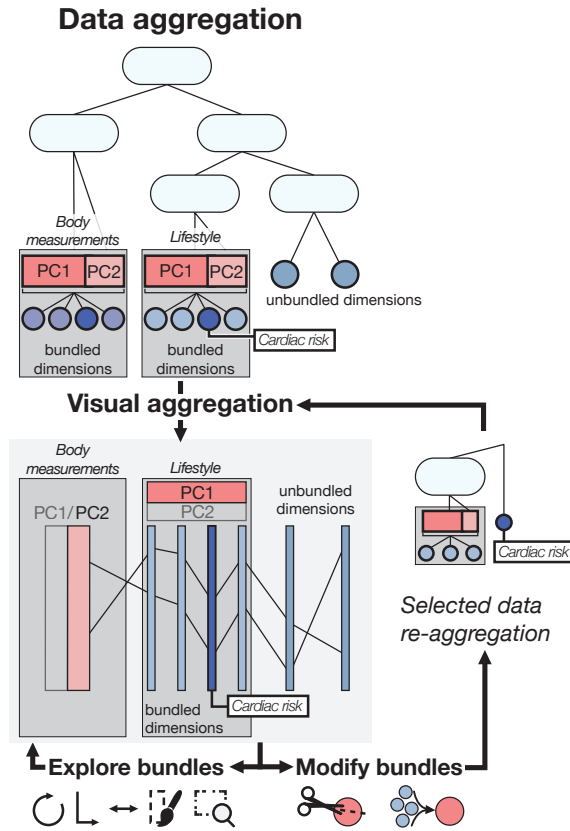


Figure C.2: Conceptual pipeline of *DimLift*. Factor analysis of mixed data (FAMD) is applied iteratively to produce *dimensional bundles* (body measurements, lifestyle). Data are mapped to a layered parallel coordinates plot for users to explore and structurally modify. FAMD is re-run on any structurally-altered *dimensional bundles* before visual remapping. We highlight the path of the dimension cardiac risk in one possible interaction flow in our approach.

to identify a single global optimum of reduced dimensions, as we discussed previously in Sec. C.1. These solutions are less effective for an exploratory approach where the user develops multiple new hypotheses over a single session. For each newly-formed hypothesis, the user needs to identify interesting, important original dimensions.

In contrast, our *DimLift* approach utilizes dimensionality reduction and user-driven methods to produce similarly-contributing groups of dimensions, i.e., *dimensional bundles*, that serve as the primary unit of exploration and interaction. These bundles reduce the analysis space while allowing the user full control to discover interesting relationships that may otherwise go unnoticed. Inspired by Elmqvist and Fekete's [121] principles for the visualization of aggregate hierarchies, we show our analytical workflow in Fig. C.2. The user initiates algorithmic construction of *dimensional bundles* with a linear dimensionality reduction technique. Subsequent visual analysis allows the user to explore the degree of bundling of their data, which offers insights on the degree of correlation within the data. Our choice of a linear dimensionality reduction algorithm

allows users to visually inspect bundle contents to identify important original dimensions that are now mapped to the new bundles. As new questions form, users may reconstruct bundles to emphasize and lift interesting patterns for detailed exploration. This series of steps may be repeated as new hypotheses and insights are continually formed.

In the remainder of this section, we present our methodological approach alongside a synthetic dataset containing human lifestyle and body measurements, organized as follows: In Section C.2.1, we detail our method for automatically generating *dimensional bundles*, Section C.2.2 discusses our visual encodings, and Section C.2.3 describes our interaction facilities for lifting expressive dimensions. We conclude with a discussion of our treatment of mixed and missing data in Section C.2.4.

C.2.1 Creating Expressive *Dimensional Bundles*

High dimensional data analysis typically involves producing a low dimensional projection of the data. Common dimensionality reduction techniques automatically treat a dataset monolithically, and may obfuscate subtle but relevant characteristics. In our simple example, smoking (b3) is an important indicator for cardiac risk (o2), but a standard dimensionality reduction does not easily show this relationship. It instead buries these dimensions in all five principal components (Fig. C.3, top). In contrast, our iterative dimensionality reduction approach extracts subsets of dimensions that contribute similarly to the variance within the dataset. We define these subsets as *dimensional bundles*. For example, our approach places the lifestyle-related dimensions education level (b1), workout frequency (b2), and smoking (b3), together with the similarly-contributing variable cardiac risk (o2) (Fig. C.3, bottom). In the following, we describe our algorithmic process to creating *dimensional bundles*. This is additionally described in pseudocode in Algorithm 2.

Algorithm 2: *Dimensional bundle* creation for two or more dimensions

```

1 initialize pool = all dimensions in dataset
2 do
3   mark all dims in pool as possibly contributing
4   initialize new bundle
5   perform FAMD on pool
6   for all dimensions in pool
7     if PC1 loading  $\geq$  contribution threshold
8       move dimension from pool to new bundle
9     else
10      mark dimension as non-contributing
11 while pool contains dimensions marked as non-contributing
12 for all bundles
13   perform FAMD on bundle
14   store PC1 and PC2 for bundle

```

Bundle Creation. Prior to analysis, we standardize all input dimensions; this ensures equal weighting between dimensions comprised of items on different scales. We then run a factor analysis of mixed data (FAMD) [387] on all dimensions, provided two or

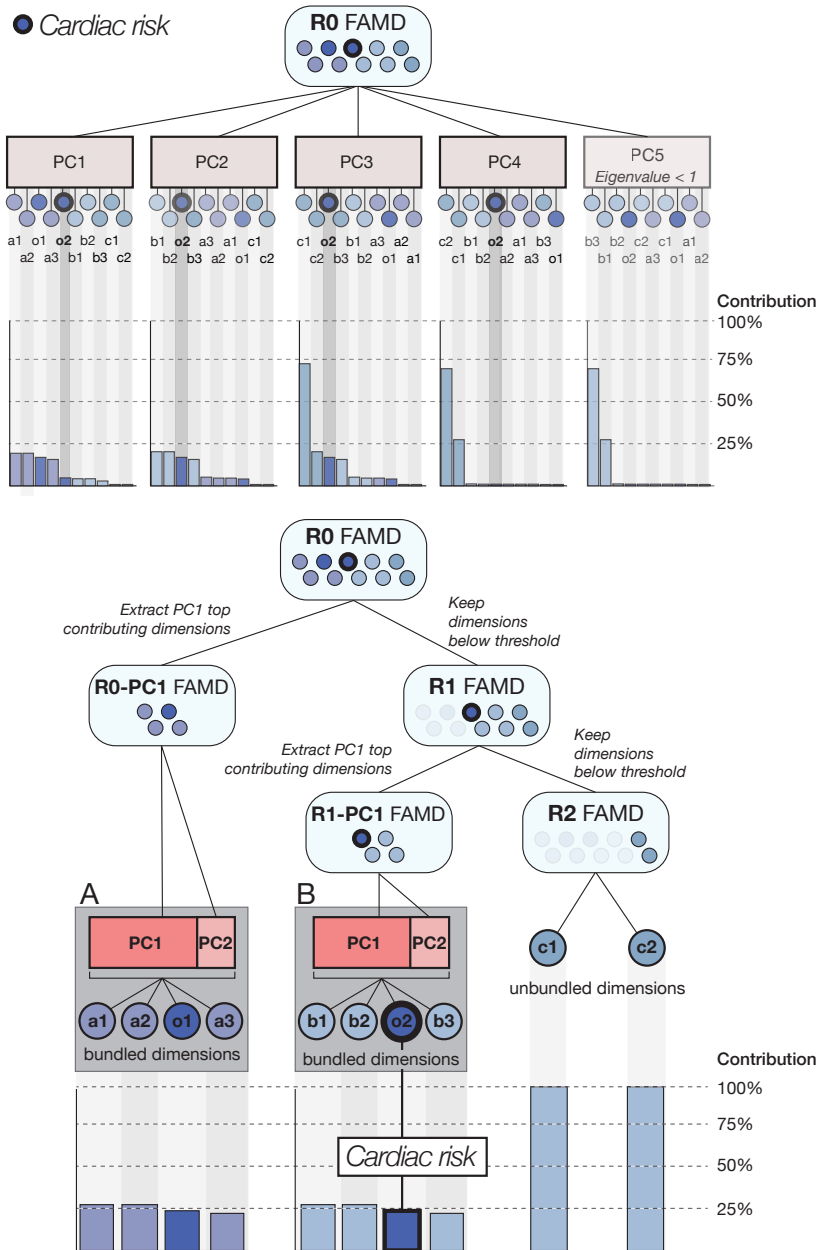


Figure C.3: We contrast our iterative algorithmic approach (bottom) with a standard approach (top) using a synthetic ten dimensional health and lifestyle dataset comprised of four quantitative [height (a1), weight (a2), waist circumference (a3), BMI (o1)] and six qualitative [education level (b1), workout frequency (b2), smoking (b3), gender (c1), eye color (c2), and cardiac risk (o2)] dimensions. A standard approach contains all ten dimensions in each principal component (PC), e.g., cardiac risk (o2) is present in all PCs. In contrast, our approach (bottom) produces a pair of *dimensional bundles* (A: body measurement, B: lifestyle) containing only dimensions with similar variance contributions, where cardiac risk is bundled into B. Dissimilarly contributing dimensions, i.e., c1 and c2, remain unbundled.

more dimensions are available in the pool (line 5). The resulting correlation matrix is used to determine principal components (PCs), their respective eigenvalues, and the contributing dimensions to each PC. We focus on the first principal component (PC1) for the creation of each bundle (line 7), as this captures the largest variance within the data [394] and shows the most potential to create expressive bundles.

Formally, PC1 is defined to maximize the sum of squared correlation coefficients r^2 between itself and each dimension k :

$$\sum_k r^2(k, PC1) \quad (C.1)$$

Referencing the loading of each dimension, i.e., the correlation coefficient r that defines the factors by which the corresponding original attributes are multiplied so that they add up to the scores of PC1, we extract only those dimensions contributing above a threshold defined as $100\% / \text{number of input dimensions}$ [1] (lines 6-8). We use this threshold as a baseline heuristic for creating bundles of similarly-contributing dimensions; it defines whether the contribution of a given dimension exceeds the average contribution to PC1. Thus, we formally define the initialization of a *dimensional bundle* as the set of all dimensions with loadings greater than or equal to this threshold, with respect to PC1.

On all *dimensional bundles* we then run another FAMD and save: (a) the principal component scores, which are the computed representations of the individual items for the bundled dimensions, and the (b) contribution and (c) loading of each dimension (lines 12-14). For this second run we do not use the threshold, and instead keep all contributing dimensions. We preserve the second principal component (PC2) in this second FAMD run to provide additional structural context for PC1, and to further indicate the quality of the bundling. We found diminishing returns for including any further PCs, particularly since the full dimensional information is already provided with the first FAMD sequence. Preserving PC1 and PC2 at the *dimensional bundle* level conforms to a manageable mental analysis model and avoids visually overwhelming the user. Thus, these three elements: PC1, PC2, and their contributing dimensions, comprise a complete organized *dimensional bundle* (Fig. C.3A, B).

The dimensions that do not meet the contribution threshold remain in the original dimension pool (line 10). We recurse on this pool of dimensions (lines 2-11) until less than two dimensions remain, or until the eigenvalue of PC1 falls below 1, meaning that PC1 accounts for less variance than one of the original dimensions (Kaiser criterion [575]). These dimensions are left unbundled (Fig. C.3, unbundled dimensions). This produces the branching structure as shown in the bottom diagram of Fig. C.3. The results of this algorithm, both *dimensional bundles* and unbundled dimensions, serve as a meaningful basis for subsequent user-driven exploration and knowledge integration.

C.2.2 Visual Encodings

Projecting to a lower dimensional subspace in dimensionality reduction often creates a degree of disconnect from the source data [235]. If the analyst can identify interesting correlations leading to new discoveries in their bundles, but is ultimately unable to relate these correlations back to the original data, then this is not an actionable application of dimension grouping. To solve this issue, we preserve and map the semantics

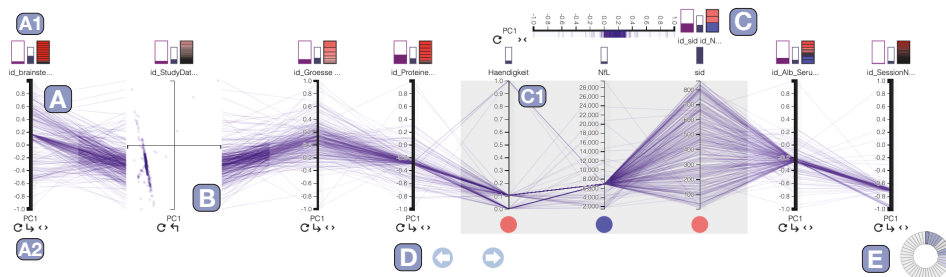


Figure C.4: *DimLift* is a mixed-initiative approach to creating and navigating *dimensional bundles*. They are defined as a subset of similarly contributing dimensions to the overall variation of a dataset, as computed from factor analysis of mixed data. Parallel coordinate axes (A) map to the first or second principal component (PC1, PC2) of a *dimensional bundle*. Glyphs (A1) provide feedback on variance contribution, missingness, and composition. View interactions (A2) allow users to pan (D) through the dataset, swap axes between PC1 and PC2, drill-down into a PC1 vs. PC2 score plot (B), or drill-down further to the *dimensional bundle* component dimensions (C) and their relationships (C1). A chart at the bottom right (E) provides an overview of all *dimensional bundles* and unbundled dimensions, a subset of which are visualized as plot axes.

of *dimensional bundles* produced through dimensionality reduction directly to visual elements. Our visual aggregation utilizes a modified parallel coordinates plot that mirrors the results of the data aggregation step. Our approach is guided by principles of unambiguous data depiction and visual-data correspondence, inspired by the algebraic method of visualization design [253].

The basic unit of the *DimLift* method, *dimensional bundles*, consist of two principal components (first and second) and their constituent input dimensions. This composition forms a hierarchy of data representations. Usually, each *dimensional bundle* has a number of sibling bundles, which are other bundles produced by our iterative FAMD approach. Our visual design is based on the following requirements, which we draw from the basic high dimension data analysis tasks that we discussed at the beginning of Sec. C.2:

- R1** Support the creation of *dimensional bundles*
- R2** Support the iterative modification of *dimensional bundles*
- R3** Allow rapid retrieval of item values in a given *dimensional bundle*
- R4** Lift dimensions of interest in a *dimensional bundle*
- R5** Provide information on the quality of each *dimensional bundle*
- R6** Allow for relation investigation between *dimensional bundles* and input dimensions

Parallel coordinates are a popular, generally applicable technique to visualize relationships and correlations in multidimensional datasets [189]. Furthermore, they have been shown as more effective in visual retrieval of data values relative to scatterplot matrices (SPLOMs) [273] (**R3**), and more performant than SPLOMs in solving tasks for higher dimensional data [366]. We utilize parallel coordinates but with some adaptations; although bifocal parallel coordinates presented by Kaur and Karki [244] visualize all dimensions simultaneously, this becomes overwhelming. We instead use an approach inspired by the perspective walls technique [421] to focus attention on bun-

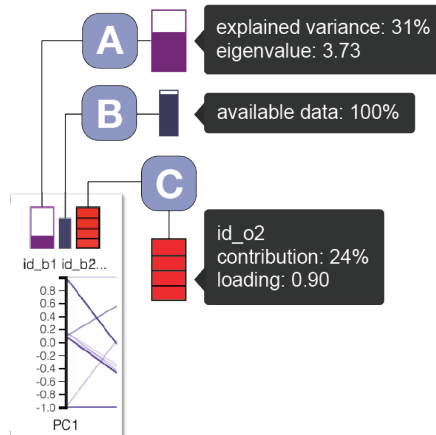


Figure C.5: Rectangular glyphs above each *dimensional bundle* axis provide information on bundle composition. These glyphs, with accompanying tooltips available on hover, display (A) the eigenvalue and explained variance, where height encodes the percent variance while the eigenvalue is revealed in the tooltip, (B) percent available, i.e., non-imputed, items, and (C) contributing dimensions and loadings in a given bundle, where bar height encodes the percent contribution while hue encodes the loading of each dimension.

dles relevant to the user. Our modified plot further supports three layers of nested visual analysis within and between each *dimensional bundle*. This nested approach is inspired by model-based reasoning methods described by Liu et al., where deeper data insights can be obtained by presenting information sets and supersets [308]. The result is shown in Fig. C.4 (**R1–R4, R6**). Beginning with the axes of a traditional parallel coordinates plot, we set the stroke-width of each axis relative to the number of contained dimensions, similar to Andrews et al. [18]. For each *dimensional bundle*, PC1 is depicted as the primary axis in the plot (Fig. C.4A). Items are plotted by their scores (**R3**). PC2 is included on-demand as a secondary axis expanding horizontally from the primary axis; this scatter plot shows item scores for both PC1 and PC2 to inform on the similarity of dimensions included in these components (Fig. C.4B). This approach is inspired by the natural orthogonality of the first and second principal components. Our nested plot approach is similar to previously suggested enhancements to parallel coordinates plots [187, 207, 226]. The innermost third level comprises all dimensions contributing to the principal components of the *dimensional bundle* (Fig. C.4C) and plots the original item values (**R3, R6**). It resides conceptually within each *dimensional bundle* parallel coordinates axis as a second parallel coordinates plot that is made visible on-demand (Fig. C.4C1), as inspired by the approach by Andrews et al. [18].

Rectangular filled glyphs, positioned above each axis, provide information on *dimensional bundle* composition and variance contribution (Fig. C.5) (**R5, R6**). Our glyph choice is driven by position-based principles from graphical perception research [90]. These glyphs display the eigenvalue and explained variance (Fig. C.5A), available, i.e., non-imputed, data (Fig. C.5B), and contributing dimensions with their respective loadings (Fig. C.5C).

Understanding the explained variance alongside the eigenvalue is a critical aspect of determining the utility of a given bundle in its ability to explain properties of the dataset. Using the Kaiser criterion [575], if an eigenvalue is below 1, we can conclude that the contributing dimensions are more informative when unbundled. The variance contribution gives an indication of the type of relationship between dimensions—a low overall variance may indicate more complex, non-linear relations. We provide this information for each bundle as shown in Fig. C.5A (**R5**).

Our approach also explicitly handles missing and imputed data. In particular, the proportion of non-imputed data items can provide feedback on the certainty of the bundles (**R5**). The amount of available, i.e., non-imputed data, is visualized by the glyph shown in Fig. C.5B. A bundle containing primarily imputed items, e.g., a mostly white/unfilled glyph, offers less certainty than a fully-filled glyph for a bundle or single dimension. We provide further details on our approach to handling of missing data and imputation in Sec. C.2.4.

To draw meaningful, actionable conclusions from an analysis the user must link back to the original data (**R3**, **R6**). The bundle composition glyph (Fig. C.5C) shows the percent contribution and correlation direction, i.e., loading, of each dimension to the bundle. The glyph is broken into segments by each dimension's variance contribution. We encode correlation direction using a diverging red-blue colormap, where red indicates a positive correlation while blue indicates a negative correlation. These encodings are additionally present in the nested dimensions parallel coordinates plot for each bundle (Fig. C.4C1) in circles placed under each dimension axis. This indicates to the user the relationship of item values to the principal component, and supports its interpretation. For instance, consider a synthetic dimensional bundle containing height, weight, and BMI: all dimensions are positively correlated and encoded with red at both plot levels. Brushing on any axis would highlight similarly high values in the principal component axis, showing that these values tend to increase and decrease together.

Both quantitative and qualitative data can be involved in relevant and interesting patterns. For instance, our synthetic dataset includes a cardiac risk outcome dimension which is comprised of quantitative body measurement and qualitative lifestyle dimensions (Fig. C.3-o2). To help reduce visual clutter and clarify relative occurrences, we utilize a horizontal bar chart extending from each categorical parallel coordinate axis (Fig. C.1), where bar length encodes the frequency of item occurrence in each category, as inspired by Hauser et al. [187].

C

C.2.3 Lifting Expressive Dimensions

A dimensionality reduction process that does not incorporate user interaction may overemphasize trivial aspects of the data. Key to the *DimLift* approach is *lifting*, an operation that changes the data hierarchy and structure of *dimensional bundles* for greater expressivity. For instance, our synthetic grouping shows cardiac risk as bundled in the automatic process with the lifestyle bundle (Fig. C.1). While useful for understanding that cardiac risk is, in our example, more closely correlated with lifestyle dimensions, we would like to lift this, and any other outcome-associated dimensions, to their own bundle for direct correlation assessment. Our method incorporating task-based user interactions allows for the flexibility to explore data at differing levels of granularity, and to reconstitute existing bundles to discover new and unexpected relationships. We di-

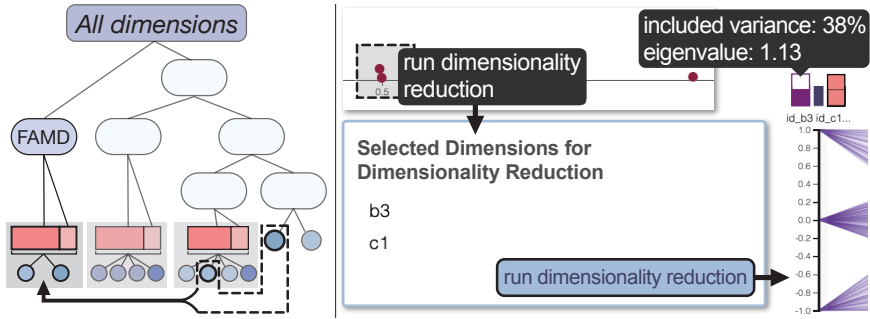


Figure C.6: *DimLift* structural interactions allow for the creation or modification of *dimensional bundles*. Using our synthetic health dataset we **create** a new bundle combining smoking (b3) with gender (c1); a resulting eigenvalue above 1 shows a fair grouping with equal dimension contributions. The left diagram provides a conceptual overview of this process.

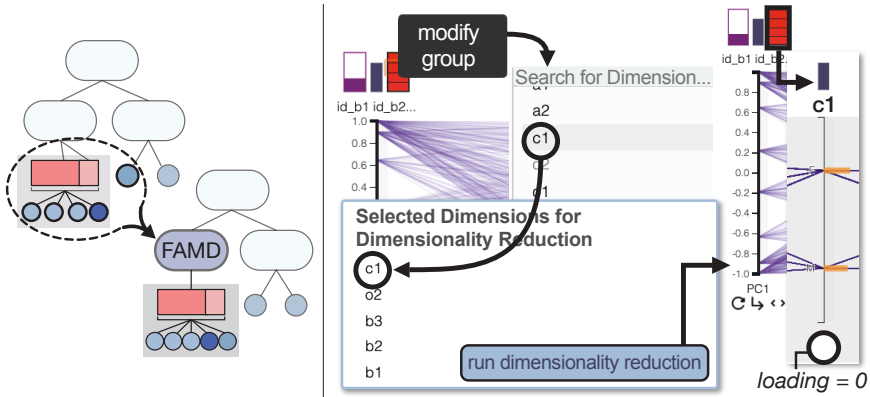


Figure C.7: In the analysis of a synthetic health dataset we may suspect gender (c1), to have interesting correlations with the lifestyle-related bundle, i.e., education level (b1), workout frequency (b2), cardiac risk (o2), and smoking (b3). We **add** gender (c1) to this bundle and observe that gender shows no contribution (loading = 0) to the bundle variance.

vide the interaction techniques for our approach into two classes. For a demonstration of the following interactions, we refer readers to the video included in the supplementary materials.

Structural interactions are operations that alter *dimensional bundles* (Fig. C.2E) by combining, adding to, or removing dimensions from these bundles. A linkage between layered parallel coordinates and a dimension scatterplot provides an easy mechanism for bundling interesting dimensions by similar statistical measures.

View interactions are inspired by the model for hierarchical aggregation interaction techniques proposed by Elmquist & Fekete [121]. These do not change the fundamental structure of the data hierarchy (Fig. C.2D), and include: pan, brush & subset, drill-down/roll-up, or swap levels in their exploration of the data space.

Structural Interactions

Dimensional bundles created automatically may not always be conducive to specific user analysis goals. As such, we introduce structural modifications that allow the user



to create entirely new, or modify existing, *dimensional bundles* to lift interesting dimensions for analysis.

Creating New *Dimensional Bundles*. During the analysis a user may wish to visualize the degree that a group of conceptually-related dimensions, e.g., all lifestyle input variables in our synthetic human measurements dataset, are correlated. Similarly, seemingly conceptually-unrelated dimensions may exhibit similar descriptive statistics, e.g., similar mode or diversity measures, that would be interesting to apply dimensionality reduction to for deeper correlation assessment. Figure C.6 demonstrates the workflow for creating a new *dimensional bundle* based on similar descriptive statistics, beginning with a marquee selection of a pair of dimensions positioned near each other. The user confirms their selection in the dimension grouping menu. Selection by descriptive statistics serves as a rough guide for the suitability of a possible bundle, which is then validated by applying the dimensionality reduction and visualizing the eigenvalue and contributing dimension attributes in the parallel coordinates plot. On creation, this bundle is briefly highlighted with a red underline in the parallel coordinates plot. When manually creating new bundles, redundant dimensions are not extracted from their original bundle to the new bundle—a single dimension can remain in multiple bundles. The reasoning is that it could be that one dimension is highly important in multiple bundles. For example, smoking (b3) is important semantically as a lifestyle variable and is logically bundled with other lifestyle variables, but it is additionally clinically interesting to bundle with, e.g., gender (c1), to assess for patterns or relationships between these two dimensions. Our approach allows the user to see this from both perspectives.

Modifying *Dimensional Bundles*. Algorithmically-created *dimensional bundles* may still bury an interesting dimension within, e.g., cardiac risk within a lifestyle bundle, or leave out a conceptually interesting dimension, e.g., gender from the lifestyle bundle. Rather than creating a new bundle, the user may simply modify the existing bundle and either add or remove dimensions in place. We demonstrate the workflow to add a dimension to an existing bundle in Fig. C.7; a right-click on the contributing dimensions glyph for the bundle of interest opens the dimension selector panel where bundle membership may be updated. The user may search by name or explore the list to add dimensions. Similarly, dimensions may be removed by entering the same panel (Fig. C.1). To remove a dimension, the user clicks on the dimension in the selected dimension list for immediate removal. After dimension addition or removal, the user can choose to run the dimensionality reduction algorithm on the updated bundle. As with bundle creation, the updated group is highlighted briefly in the parallel coordinates plot. All contribution information is updated in the glyphs, and correlations in the parallel coordinates are updated automatically.

These structural modification tools empower the user to reconstruct the dimensional hierarchy for open exploration. With flexible bundle composition and modification, along with feedback on their suitability in the parallel coordinates plot, users may rapidly form new insights about their data by lifting dimensions of interest from their original bundles.

View Interactions

With the *DimLift* approach users may visually navigate *dimensional bundles* via the previously described layers: the top-level parallel coordinates axes for between-bundle




Figure C.8: *DimLift* view interactions allow iterative exploration of *dimensional bundles*. Panning through the parallel coordinates plot allows the user to explore correlations between all bundles. Brushing over a bundle axis, e.g., lifestyle, creates a subset of moderately active, university education level smokers with moderate cardiac risk. This selection is adjustable in an adjacent panel. Drilling down to a plot of PC1 vs. PC2 item scores shows a distribution shape that is interesting to explore further. Drilling further to the contributing individual dimensions shows a definite correlation between b1 (education level) and b2 (workout frequency) (black rectangle). Swapping the bundle axis from the first (PC1) to second (PC2) principal component shows that b3 (smoking) contributes no variation to this component, while it contributes similarly to other included dimensions in PC1.



navigation (A), or nested scatterplots (B) and further nested parallel coordinates (C) for within-bundle navigation (Fig. C.4).

Pan. Pan operations are ubiquitous in visual analytics, particularly in aggregated datasets [121]. We utilize panning to bring *dimensional bundles* of interest into the field of view, as shown in Fig. C.8. Arrow buttons allow incremental panning while a small donut chart, used as it mirrors the panning-carousel nature of the parallel coordinates plot, provides a quick overview of the created bundles and individual dimensions while acting as an additional navigational aid [351]. It additionally serves to spotlight those bundles with subsets applied. Within this glyph, bundles are denoted as purple, while individual dimensions are grey. In addition to these manual controls, panning can be facilitated by axis reordering based on descriptive statistics, i.e., variance, standard deviation (and their qualitative analogs), diversity, modality, and percent missing values [351], or by order of extraction via the iterative FAMM algorithm.

Brush & Subset. Brushing and linking are commonly used in visual analytics to link data elements [510] across views. Our approach relies on this premise, but rather than only brushing and linking items or individual dimensions [351], we support brushing and linking of *dimensional bundles*. In our method the user may brush a dimension or dimension bundle in the dimensions overview plot (Fig. C.6) or in the layered parallel coordinates plot (Fig. C.8). In the latter, brushing creates subsets of *dimensional bundles*, as demonstrated in Fig. C.8, which may be subsequently adjusted.

Drill-Down and Roll-Up. Drill-down and roll-up are two primary methods for viewing data at multiple aggregation levels [121]. Since *dimensional bundles* comprise two PCs and raw dimensions, we utilize two different methods to access each data type in a bundle. The first branch explores PCs in a given *dimensional bundle*: with this method, the user may see the orthogonal axis of variation presented by the grouped data axis (Fig. C.8); this provides a greater sense of the bundling strength and reasonability.

The second method accesses constituent dimensions of PCs within a given *dimensional bundle*. To differentiate from the first method we drill-down/roll-up on a horizontal axis, e.g., expand/collapse (Fig. C.8). Expansion occurs in-place, and allows the user to assess correlations within and outside a given *dimensional bundle*.

Swap. Described as a flip operation by Elmqvist and Fekete, this allows the user to observe neighboring siblings in an aggregate hierarchy [121]. We can consider the first (PC1) and second principal components (PC2) as siblings in our aggregate structure. We view this operation as fundamentally different from drilling-down, as the swap does not add detail to the existing view but rather shifts to a different, related frame. Lifting the secondary axis of variation in the bundle to the surface permits visualization of interdimensional correlations through a different lens (Fig. C.8). This allows prioritization of second-level variation structures in the data to establish subsurface patterns.

C

C.2.4 Handling Mixed and Missing Data

Our algorithmic approach as described in Sec. C.2.1 may furthermore handle complex data, as characterized by missing items and mixed data types. In the instance of a purely qualitative dataset we perform all steps previously described, with a few alterations: we first use multiple correspondence analysis (MCA) to convert all qualitative dimensions to quantitative dummy variables [28, 436]. Then, rather than the squared

correlation coefficient criterion we instead use the squared correlation ratio to identify the PC1 leading to each *dimensional bundle*. In instances of mixed datasets, the algorithm simply uses the appropriate criterion to define each *dimensional bundle*.

Furthermore, data are often incomplete, as was true for one of our case studies which was 76% incomplete. While some solutions drop cases with missing data from the analysis, this can easily lead to an inaccurate picture of dimension correlations. Imputation of missing data is still a highly debated area of research, and is dependent on the analysis goals and the data itself. While our approach is flexibly designed to allow a variety of imputation methods, our default method is multiple imputation of chained equations (MICE) [554], a multiple imputation method, to minimize bias and reduce standard error. This default can be changed by the user. A key feature of MICE is that it can handle different variable types: quantitative continuous, binary, and ordered and ordered categorical data. As we aim our method to be broadly applicable to mixed datasets, this was a critical aspect of our decision process. It furthermore is widely used in epidemiology [108], a field known for its complex and highly missing data, and was chosen after discussions with our clinical collaborators on this paper.

MICE is applied by default to all dimensions with missingness of 78% or lower. We chose this default value experimentally, as this was the limit up to which MICE was typically still able to provide meaningful results. In the extreme case of dimensions with only a handful of total entries where multiple imputation produces meaningless results, e.g., 99% missing, we instead perform a single value imputation using the mean for quantitative variables and create a new "not defined" category for categorical data. The missing data glyphs serve as identifiers for the reliability of the data for patterns observed in these dimensions. We explore the impact of different imputation methods in the discussion section and supplementary material.

C.3 Case Studies

We implemented our approach as a web application using Javascript and D3.js [60]. Descriptive statistics computations and dimension groupings are performed in a Flask Python back end, and we use FactoMinR [289] to perform the FAMD in R. The full source code is available at <https://github.com/lauragarrison87/DimLift>.

Following initial analysis of the data via our iterative FAMD algorithm and visual aggregation, the user may explore the resulting dimension hierarchy. As part of the exploratory analysis process, users may flexibly adjust membership of *dimensional bundles* and construct a new dimensional hierarchy to lift interesting dimensions to the surface.

In the following, we demonstrate the value and versatility of *DimLift* applied to three data scenarios, one of which corresponds to an ongoing clinical collaboration. Analysis of data in these domains typically utilizes statistical analysis packages which are unwieldy when applied to open-ended data exploration [174]. Before discussing our clinical cohort case study, we introduce nutrient and plant ecology datasets to demonstrate our method's general applicability by reproducing insights from existing domain literature.

Cholesterol, Selenium, Vit. B12, Choline




Figure C.10: Analysis of FDA nutrient dataset using our approach. Selenium, vitamin B12, cholesterol, and total choline are bundled together automatically in our approach; this corroborates a known link in clinical literature between Selenium and Cholesterol. Subsetting to only the high values of selenium shows low values of cholesterol, although with values near 0 this finding needs confirmation with a larger dataset.

generated in a very short session, show promise for our approach in quickly lifting and establishing relationships that corroborate results from a complex plant ecology study.

C.3.2 USDA National Nutrient Data

We next demonstrate insights generated with our method using data from the USDA SR28 National Nutrient Database [514]. This database is the standard reference for food nutritional content in the United States; many of these variables correlate and thus this dataset lends itself well to dimensionality reduction. The subset we analyzed is predominately comprised of quantitative data and consists of 899 data items in 53 dimensions. Selenium is an essential micronutrient for effective thyroid hormone and reproductive function; when levels are sufficient in the body it has been shown to provide antioxidant and anti-inflammatory effects [413]. Cholesterol plays a known role in cardiovascular health; high total cholesterol levels are strongly linked to higher cardiovascular risk [478]. Chen et. al. [85], in their seven-year longitudinal nutritional cohort study, found that participants with higher levels of selenium exhibited a greater decrease in total cholesterol over the course of the study; this offers insights on selenium's possible mitigating effect of cardiovascular risk in elderly populations. Our goal in this exploratory analysis of nutrient data was to establish the possible ease and clarity of discovering this known clinical link.

Our approach rapidly lifts Selenium to the surface, placing it in a *dimensional bundle* alongside Vitamin B12, Cholesterol, and Total Choline; each contribute approximately 25% to the bundle variance (Fig. C.10). This immediate insight corroborates the correlation between these nutrients, while also providing an interesting line of inquiry on the additional relatedness of vitamin B12 and Choline. Creating a subset of high selenium values, our results are not as conclusive since we have a small popula-

tion sample in our dataset, but the results indicate the same link as shown in the clinical literature that we discovered in a short period of exploration. If we then remove VB12 and Choline from the bundle to hone in on the relationship between Selenium and Cholesterol, we find equivalent positive loading values for each. This further supports a positive colinear relationship between these two nutrients. If performed in a standard FAMD this link would have been difficult to identify, as the results would show all dimensions in each principal component. The subtle link between selenium and cholesterol would be buried beneath the stronger variance contributions of other nutrients, e.g., magnesium, folate, and calcium, to the data. Our approach allows this dimensional relationship to be immediately apparent.

C.3.3 Clinical Cohort: Cerebral Small Vessel Disease

The ultimate analysis goal for any clinical cohort dataset is to develop testable hypotheses that can lead to better treatment options and outcomes for the patient. One of the great difficulties with clinical datasets lies in the successful identification of interesting measures and patterns, particularly in diseases where the etiology is not entirely clear, e.g., in cerebral small vessel disease (CSVD). The current standard for analysing this type of data consists of queries with complex statistical analysis packages. Of such tools, our expert participants most frequently use SPSS. This and similar applications typically require extensive processing times, and are not generally conducive to an iterative, exploratory approach to the data. Having previously analyzed these data in SPSS, one expert noted that for an efficient analysis with SPSS they need to already have in mind the variables of interest and be familiar with the data characteristics prior to their assessment.

The data consist of 307 patients collected from clinical routine in the university hospital data management system. The data are mixed, consisting of 168 dimensions containing demographic, laboratory, education, and lifestyle information. 24 additional dimensions describe the volume of 24 brain structures, e.g., hippocampus and caudate, as derived from T1-weighted magnetic resonance imaging data. As is typical with this type of data, 76% of entries are missing due to, e.g., missed appointments, not all patients needing the same tests, and other criteria.

We performed two joint analyses of a clinical cohort for the study of CSVD. After a short presentation explaining the method and the prototype application, the experts explored the application themselves using a “think-aloud” protocol. Our primary goal with this study was to allow domain experts to freely explore their data in *DimLift* to assess ease and speed of iterating and forming new insights into possible CSVD-related measures and patterns. We highlight key aspects of their respective analyses; for further demonstration we refer the reader to the supplementary video.

Analysis 1. The first analysis was performed alongside one MD/PhD clinical neurologist specializing in cognitive aging and mixed cerebral pathologies, who is also a co-author of this paper, and one master-level engineer in neuroscience in a paired analysis session for one hour and 30 minutes. Both are experts in advanced statistical analysis of CSVD data; their workflow was particularly interested in the bundle contents and loadings. On loading the data, experts first browsed the number and contents of the created *dimensional bundles*. Noting from the glyphs above each axis that




Figure C.11: In exploring a clinical cohort dataset for cerebral small vessel disease (CSVD), experts select a bundle of primarily imaging data for closer examination and drill-down to observe two distinct clusters. Swapping the axis to PC2 allows subset creation of the top cluster; this corresponds to selection of non-imputed items within the bundle. Addition of APOE-related dimensions to the bundle allows for correlation assessment of these interesting dimensions within a single bundle.

many bundles suffer from missing data, the experts used hover features to assess bundle variance and dimension contributions. The experts were surprised that *lacunes* and *microbleedings* were bundled along with two Boston/STRIVE criteria dimensions and thus, lifted together. However, their bundling makes sense since these have been shown to correlate [392]. From this, they can hypothesize that lacunes and microbleeds in certain regions of the brain could be associated with a certain subset of Boston criteria. This has implications on bleeds in certain areas of the brain being indicators for aspects of CSVD. They stated, “*We would have probably not seen this in another framework.*”

Locating another interesting bundle containing primarily imaging data, as well as diagnosis and sex, they then drilled-down for further exploration. In this second level they observed two clusters. While variables contributing to PC1 are mostly imaging-related, PC2 contributing variables include Boston/STRIVE criteria at lumbar puncture, group, and sex. They swapped the axes and observed how the item distribution (Fig. C.11) and dimension statistical distributions are affected. They noted that brushing the top cluster selects individual dimension values that are complete, i.e., Boston/STRIVE criteria, Sex, and Group axes now exclude “not defined” items through this subset selection. Experts then added APOE-related dimensions to investigate the relationships within this bundle. However, they noted that the APOE dimensions do not provide strong contributions (the loading glyph on the original dimension is white in color)—this implies that these are not particularly correlated, and may indeed be better treated as separate dimensions or *dimensional bundles* (Fig. C.11). This allows them to reject their hypothesis that APOE genotypes are highly correlated to Boston/STRIVE criteria at lumbar puncture, group, and sex for this dataset. However, they note that this would be more interesting to explore in a larger cohort before fully rejecting this. They further noted that this subcohort is characterized by generally mid-to-high range white matter and

CSF volume values, but a broad range of volume data for other regions. From this, they hypothesized that these volume ranges of white matter and CSF can act as potential biomarkers for CSVD. This requires additional followup with a larger cohort and additional cognitive and clinical tests.

Analysis 2. Our second analysis session also lasted one hour 30 minutes, with a medical expert who has one year of experience in CSVD research and who is less experienced in statistical analysis. This user was primarily interested in generating a picture of the typical patient for each diagnostic group in the dataset. As such, they were less focused on the bundle loadings and differences in principal component loadings for each bundle, and rather interested in the composition and linear correlations within bundles. Their workflow generally went as follows: (1) *Bundle overview*, (2) *Subset within bundle to explore correlations*, (3) *Modify bundle contents*, and (4) *Repeat steps 1-3*.

In their overview of bundles they observed that a high proportion of the data was missing. They noted that this information is helpful because they know the statistics they explore have reduced power in hypothesis generation. In the bundle comprised largely of imaging data they were initially surprised that these were bundled, but reasoned that this was logical since these were tied to Boston/STRIVE criteria, which was also grouped in this bundle. The user then created a subset of CAA/HA/Mixed patients in the Group dimension, hypothesizing that if a patient has CAA that they are more likely to also suffer from seizures, stroke, and dementia in the pathology dimension. Direct correlation visualization between these parallel coordinate dimension axes allowed them to confirm this; such a finding also corroborates clinical literature findings. However, they stated that this would need to be verified in a larger and more complete clinical dataset.

On exploring other dimensions in the same bundle, the user indicated that some dimensions were not, in fact, particularly interesting to analyze, e.g., all of the imaging dimensions except for the hippocampal and white matter volume measurements. The ability to easily modify this bundle to remove these uninteresting dimensions for their current hypothesis was extremely helpful for them. In doing so, they were able to quickly note a slight positive correlation between these two dimensions that also related to the diagnostic group subset; this allowed formation of a second hypothesis: that hippocampal and white matter volumes correlate to this group of diagnoses in CSVD.

Having previously analyzed the data in SPSS, they noted that these preliminary trends and relationships they found interesting with over one hour of using SPSS could be found in five minutes using our approach, simply by drilling down into a bundle that contained already most of the dimensions of interest and subsetting to a certain diagnostic group, i.e., the CAA/HA/Mixed category.

Expert Feedback. Although clinical experts noted that our visual approach appears very complex in the beginning, they were able to operate it independently after ten minutes. They stated its usability to be very intuitive due to the visualizations, hovering facilities, and interactions. However, they felt that dimensions with freeform or prose text entry were not ideal for exploratory analysis with this tool. Not only is the variance on these dimensions massive because there are no defined categories (every entry can be unique), but readability may be problematic in the parallel coordinates view. They felt also that tracking patients and variable changes over a longitudinal

study would not be easy with this system, although they felt this approach serves a different purpose. Experts felt that SPSS provided means for a more direct and targeted analysis method, while DimLift takes a more open, discovery-oriented approach. As such, DimLift may be unnecessary to use if one has already identified target variables and wishes to perform specific statistical analyses of significance. However, they felt that for open exploration DimLift is a faster (e.g., Analysis 2 required 1.5 hours in SPSS compared to ten minutes in DimLift to identify a new hypothesis) and easier-to-use solution with visual aids that are neither readily or easily available with SPSS and R. To use the DimLift approach to its full potential experts agreed it is important to have a basic knowledge of statistics and dimensionality reduction techniques, otherwise the rationale for the algorithmic bundles may be difficult to appreciate. This level of statistical knowledge is common in clinical research. With our approach to high dimension space exploration and modification, all three experts were able to rapidly gain new insights into the data via the *dimensional bundles*, and to easily reflect the principal components back to the original dimensions. They stated that this tool is especially helpful for hypothesis generation, and recommended its usage within clinical research.

C.4 Discussion

Although we explored a number of possible algorithms to drive our technique, we ultimately chose factor analysis of mixed data as it is quite general and allows the analysis of mixed data by combining PCA and MCA. The broad applicability of this algorithm makes it a clear first choice for exploring this type of hierarchical creation for our visualization. However, this comes with an expectation for normally-distributed data, which is not always the case. An interesting avenue for future investigation is how our approach could be integrated with nonlinear dimensionality reduction techniques, although this presents other challenges in mapping back to the original dimensions.

Our algorithmic approach furthermore treats all dimensions as active, i.e., all dimensions are used in FAMD, and excludes the possibility for supplementary dimensions, i.e., dimensions that are not used directly in FAMD. While supplementary dimensions do not impact the eigenvalue of a bundle or dimension contributions, these can provide further insights by distinguishing correlations between active vs. supplementary dimensions. This is particularly interesting to explore further in user-driven bundle creation.

Our treatment of *dimensional bundle* labeling concatenates the names of all input dimensions, producing long names which are not fully visible at a glance. The ordering does not mirror the contributions of the dimensions, since we preserve the label for the first and second principal component to avoid confusion. Descriptive labeling of the new dimensions produced by dimensionality reduction poses an ongoing challenge in the community, and our approach could benefit from more advanced solutions.

Handling of missing data is an active area of research, and imputation methods are highly dependent on the particulars of the dataset. Our exploration of imputation methods involved a literature search, discussion with clinical collaborators, and testing of four selected imputation methods in our clinical cohort dataset, chosen as the test imputation dataset for its high proportion of missingness. The imputation methods we tested included overall mean imputation, hot deck imputation, principal components method,

Table C.1: Processing times required (MacBook Pro quad-core i5 processor) for three datasets using our approach.

Dataset	Items	Dimensions	Processing time (sec)
Plant	1758	13	132.4s
Nutrient	899	53	24.5s
Clinical	307	193	6.8s

using the missMDA R package [28], and multiple imputation of chained equations. In our tests we found that each imputation method created 10-11 bundles, with big trends or correlations generally preserved between each method, i.e., lacunes and microbleeds from various regions of the brain were mostly bundled together. Although naturally some differences were present between each method, the differences and bundling for these generally followed an outcome that made conceptual sense. Although we ultimately chose MICE as our default imputation method for its popularity in epidemiology studies, which are known for their complexity with mixed data types and missing elements [108], we have available as options the ability for the user to switch to any other imputation method as necessitated by the characteristics of their data and their analysis goals. We document in supplementary material details of our testing of these different methods. A strong benefit that we found in this exploration of the effects of imputation methods on our bundling is that it allowed us to discover more robust patterns within data we analyzed. The exploration of different types of imputation presents an exciting and challenging topic of research in visual analytics.

An additional challenge in missing data imputation is that there is not an established threshold of missingness in literature for which statistical analyses become no longer relevant [109]; this instead is highly dependent on the data itself. This problem was particularly relevant to our clinical case study, which in several dimensions were only 3% complete. We experimented with threshold settings for degree of completeness for each dimension, and used this threshold to determine whether we applied MICE or a more simple single-imputation method.

Currently, data preprocessing uses the existing implementation of FAMD in R, which provides adequate performance for moderately-sized datasets. We include Table 1, which lists the case study dataset number of items, dimensions, and processing time. We tested all cases on a MacBook Pro quad-core i5 processor. As we can see in Table 1, processing time is more sensitive to the number of items, rather than the number of dimensions. In order to more efficiently process high item datasets, a more optimized custom implementation would be beneficial. However, wide and shallow datasets, i.e., low item but high dimensionality, are processed relatively quickly.

Parallel coordinates as the base design for bundles may begin to suffer with a very high number of independent, uncorrelated dimensions in a dataset, as this would introduce a high number of axes that would then be prone to issues already known with visualization in parallel coordinates. Although we have explored the utility of our approach in datasets numbering into hundreds of dimensions, as befitting the data of interest for our clinical partners, we imagine a future work exploring the extensibility of this approach to even higher dimensional data. User interactions to structure their own bundles by conceptual relatedness paired with the described view interactions may still mitigate this dimensionality challenge; through a relatively small feature set we allow

a comprehensive analysis of the structure of the data with enough flexibility to explore and generate new hypotheses from this starting point. More complex interaction facilities that could perform a combination of steps in one would save the user time, but then run the risk of losing track of the semantics for the user to fully understand the consequences of their adjustments in the visualization.

C.5 Conclusions

We presented *DimLift*, a novel approach to creating and interacting with *dimensional bundles* that lifts interesting relationships to the user's attention. While prior approaches allow exploration of data in both item and dimension space, *dimensional bundles* provide an additional layer that reduces the analysis space in an expressive manner. Our method is driven by an iterative factor analysis of mixed data (FAMD) that produces expressive subsets of dimensions contributing similarly to the overall variance of a dataset. We provide a means to more transparently link data inputs and track transformations of *dimensional bundles* during the exploratory process through visual and interaction design elements grounded in a layered parallel coordinates plot. Through these interactions, expert users are able to explore possible dimensions of interest in the context of the structural hierarchy, and then proceed to dismantle and rebuild this hierarchy through different views and levels in the hypothesis generation process to meet their own hypotheses for bundle expressivity. We demonstrated our workflow in a study of ecological and nutrient data and in a paired clinical case study with medical experts. With each of these cases we were able to both corroborate existing findings, and establish new insights.

While statistics remains a necessary tool in high dimensional data analysis, statistical strength cannot itself dictate feature importance. User knowledge and semantics remain critical elements of this process. Furthermore, we draw a distinction between *analysis* and *exploration*. While analysis requires specific questions to leverage statistical techniques, exploration proceeds and utilizes statistics in a stepwise fashion, allowing the user to disregard irrelevant information and lift relevant items to the surface.

Dimensional bundles are a useful concept for interacting with high dimensional data. This opens the door for a number of areas of future research, including their possible connections to edge bundling for graph and network data visualization, as described by Holten and Van Wijk [205]. While *DimLift* primarily focuses on formation and interactions with *dimensional bundles*, a logical next step may allow for selections made in dimension subspace to additionally drive dimensional bundle formation. Similar to lifting of interesting dimensions, this could allow for interesting subspaces to be lifted to a primary view level. Other areas of future work may focus on evaluating *dimensional bundles* in other domains as a field study. This may bring further insights on the broad utility of such bundles in subspace exploration. We may also explore this in a controlled study uniquely adapted for exploratory tasks, although by nature this is quite challenging. Additional areas of future exploration includes the applicability of our approach in non-linear dimensionality reduction methods, where the interpretation gap presents a major hurdle to understanding the data, as well as integration of *dimensional bundles* with other interaction and visualization techniques.

Acknowledgments

We thank Frank Schreiber and Philipp Ulbrich for their invaluable feedback in our clinical case study. Parts of this work have been done in the context of the Center for Data Science (CEDAS) at the University of Bergen. This research is supported by the University of Bergen and the Trond Mohn Foundation in Bergen (#813558, Visualizing Data Science for Large Scale Hypothesis Management in Imaging Biomarker Discovery (VIDI)), and by the Federal State of Saxony-Anhalt, Germany (FKZ: I 88).

Paper D

An Exploration of Practice and Preferences for the Visual Communication of Biomedical Processes

Laura A. Garrison^{1,2}, Monique Meuschke³, Jennifer Fairman⁴,
Noeska N. Smit^{1,2}, Bernhard Preim³, and Stefan Bruckner^{1,2}

¹Dept. of Informatics, Univ. of Bergen, Norway

²Mohn Medical Imaging and Visualization Centre, Haukeland Univ. Hospital, Norway

³Institute for Simulation and Graphics, Otto-von-Guericke Univ., Germany

⁴Dept. of Art as Applied to Medicine, Johns Hopkins Univ., USA

Abstract

The visual communication of biomedical processes draws from diverse techniques in both visualization and biomedical illustration. However, matching these techniques to their intended audience often relies on practice-based heuristics or narrow-scope evaluations. We present an exploratory study of the criteria that audiences use when evaluating a biomedical process visualization targeted for communication. Designed over a series of expert interviews and focus groups, our study focuses on common communication scenarios of five well-known biomedical processes and their standard visual representations. We framed these scenarios in a survey with participant expertise spanning from minimal to expert knowledge of a given topic. Our results show frequent overlap in abstraction preferences between expert and non-expert audiences, with similar prioritization of clarity and the ability of an asset to meet a given communication objective. We also found that some illustrative conventions are not as clear as we thought, e.g., glows have broadly ambiguous meaning, while other approaches were unexpectedly preferred, e.g., biomedical illustrations in place of data-driven visualizations. Our findings suggest numerous opportunities for the continued convergence of visualization and biomedical illustration techniques for targeted visualization design.

D.1 Introduction

New technologies exposing novel aspects of science and medicine have increased demand for visual methods and tools for both experts [362] and non-experts. While numerous visualization works have been inspired by biomedical illustration [412], the demand for science communication has driven an increasing convergence of these two respective disciplines. For example, CellBlender [247, 480, 481], a molecular simulation plugin for Blender [51], can be used by both biomedical illustrators and visualization scientists for analysis and communication. Along with this increased demand for new visualizations and tools comes a need to understand their utility for different audience types. Differing values between audience types were apparent at the 2020 VCBM Workshop Image Competition, where the contest winner as selected by a jury of biomedical illustrators received one of the lowest rankings according to conference attendee popular choice. The two audiences clearly evaluated and prioritized different aspects of the visualizations in the competition. As a whole, our community lacks a clear understanding of the rationale behind differing audience preferences, and similarly lacks a complete view of the various scientific and illustrative techniques used to visualize biomedical processes.

Our goal is to gain insights into how visualization and biomedical illustration techniques are used and assessed by differing audiences for visual communication. In an interdisciplinary approach with biomedical illustrators and visualization scientists we explored the similarities, as well as differences, in common approaches to visualize biomedical processes. From this study we identify opportunities for further growth and convergence of techniques. The five topics we surveyed (signal transduction, constitutive activation, blood flow, aneurysm, and metastasis) span the micro- to macroscale and include patho- and physiological processes to serve as a proxy for the large space of representations of biomedical processes. For each topic, communication scenarios and assets are designed in conjunction with expert focus groups. This approach controls the design space while providing important in-depth insights on discipline-dependent visualization practices. Specifically this study contributes:

- Insights into the **design considerations** necessary to develop materials for communication of biomedical processes from both a visualization and biomedical illustration pipeline.
- **Curated assets** demonstrating typical techniques used to depict five common biomedical processes.
- A **qualitative survey** involving participants with diverse and creative expertise to evaluate visualization preferences for scenarios targeting (1) expert and (2) non-expert audiences.
- Reflection on **patterns observed in preferences** between different audience types with suggestions for further research.

D

D.2 Related Work

Our work is rooted in visualization design principles to communicate science through illustrative and data-driven means. We draw on prior ideas of abstraction spaces, with aspects of our survey modeled on the existing body of qualitative visualization research.

Purpose of Visualization. A number of theoretical frameworks guiding visualization are largely *data and task-centric*. Both Tominski & Schumann [503] and Munzner [354] frame the purpose of visualization as the exploration, description, explanation, communication, and/or presentation of data. For visualization task identification and validation, Brehmer & Munzner describe a multilevel task typology exploring the what, why, and how of visualization tasks [66]. Munzner’s nested model of visualization [353] provides a means for visualization scientists to evaluate their design choices at four distinct levels, from domain characterization to algorithm design.

Several works [168, 231, 232, 239] place an emphasis on visualization for communication, education, and outreach using *illustrative* techniques which often come from a practice-based perspective. Sousa et al. similarly include illustrative approaches in their illustrative visualization framework to help scientists approach and solve visualization tasks [471]. This parallels a traditional illustration pipeline of first receiving and recording information, then sketching and refinement, followed by rendering and addition of labels. Similar to these works, we take a broader view of visualization that includes illustrative and data-driven techniques aimed towards communication.

Abstraction in Visualization. Abstraction is inherent to visualization. Viola & Isenberg provided a formalization of abstraction in visualization [531]. Their definitions and updated formalization [530] of visual abstraction serve as the basis for the abstraction spaces in our study. Rautek et al. describe abstraction as a powerful visual communication tool which can lend additional insights to one’s data [412]. Andrews takes a similar view of abstraction from the perspective of biomedical illustration, discussing instances where illustration is an optimal medium to visualize certain concepts, e.g., to easily remove “visual garbage” or to superimpose structures [17]. This discussion is reminiscent of the data-driven principles of visualization stated by Tufte, e.g., avoidance of “chart junk” [509]. Abstraction, when fit appropriately to the task, lays the foundation for a successful visualization that can be evaluated empirically.

Empirical Visualization Studies. While empirical studies are increasingly considered as core elements of visualization research [87], the challenges to conducting a good empirical study are numerous [549, 580]. For example, use of expert reviews, rather than conducting a broader user study, is strongly dependent on the evaluated visualization and its development stage. Tory & Möller found value in conducting expert reviews particularly in evaluations of early prototypes [505]. Our survey targeted experts from diverse domains in order to focus our participant pool to those with sufficient knowledge to understand and provide high quality feedback on all presented assets and scenarios. Empirical visualization research may be conducted to understand the field of visualization as a whole, e.g., studying visualization research keywords [219], or specific terms, such as memorability [55, 56, 301]. Our survey design, and the use of keywords, is inspired primarily by the broader approach presented by Isenberg et al. [219].

Empirical studies on biomedical visualization are often controlled studies with narrow scopes, e.g., evaluating a specific technique. Such evaluation studies tend to focus on perceptual and cognitive aspects, e.g., Baca et al.’s study assessing efficacy based on usability, aesthetics, and iterability for a visualization of combustion [32]. Although we included both expert and non-expert audiences and also considered aesthetics as a

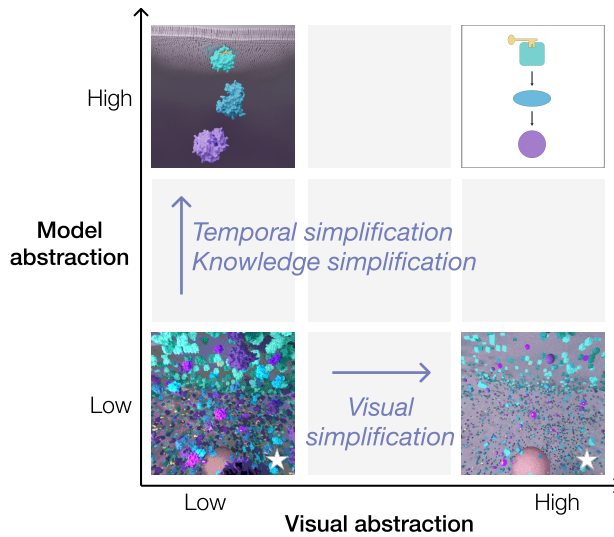


Figure D.1: Conceptual abstraction space. Model abstraction spans the relative knowledge precision, i.e., the creator’s mental model, of the source data and its temporality, while visual abstraction encompasses the relative visual simplification of the model (stars denote animated assets).

variable, we took a larger, qualitative scope. Comparative studies may examine traditional illustration methods, e.g., pen and ink, relative to computational renderings that mimic the traditional style [221]. Such an approach may also assess different computational techniques, e.g., semi-transparent structures in volume rendering [122], stylization and color adjustments to improve the aesthetics of surgical field imagery [48], or perceptually comparing aneurysm anatomy with embedded flow visualization [34]. Our survey focused on the comparison of assets produced using different visualization or biomedical illustration techniques. As Baer et al. [34], we asked participants to indicate personal preferences in their selections.

D.3 Abstraction Constructs

We apply two abstraction constructs to every asset: model and visual abstraction, as depicted in Fig. D.1. This creates a common foundation to compare audience preferences both within and between the five biomedical topics. We draw from the terminology and definitions of abstraction by Viola et al. [530]. The authors discuss abstraction of data representations and abstraction of visual representations as two distinct phases in the visualization process, beginning with entirely non-visual data representations. Here, the authors conceptualize data abstraction as the steps to achieve a desired sparsity of the dataset after acquisition, cleaning, and filtering. We expand on this data-driven notion to encompass the data representation and abstraction process for biomedical illustrations.

Model Abstraction. Rather than thinking of data only in the context of its attributes, we additionally consider the knowledge precision, i.e., the creator's *mental model*, of a given phenomenon. In addition, the temporal level of complexity plays a role in the level of abstraction in the resulting model. This accounts for understanding of the **details** and **dynamics** of a given biomedical process, e.g., signal transduction. These aspects constitute a generalized type of data abstraction that we term *model abstraction*. To illustrate model abstraction, consider the top-left and bottom-left assets in Fig. D.1. The bottom-left asset, a rule-based stochastic visualization, requires a higher degree of knowledge precision to produce than the top-left asset. With regards to temporality, this asset is less simplified, as it captures the naturally dynamic process of signal transduction more than the asset above with a reduced and static molecular environment.

Visual Abstraction. Visual abstraction can preserve and emphasize the most salient information to allow the viewer to extract meaningful information. We consider visual abstraction as the extent to which the underlying model is visually simplified. This includes shape abstraction, e.g., a molecule visualized from x-ray crystallography data has a low visual abstraction (Fig. D.1, left), relative to a shape primitive representation of the same molecule (Fig. D.1, right). Visual abstraction also applies to environments, e.g., the removal or simplification of background elements to draw attention to the desired elements as on the top-right of Fig. D.1. This is utilized in many focus+context techniques [186, 412].

Abstraction Space. We place each abstraction construct along an *abstraction axis*. Each axis describes a sequence of visual representations that incrementally depict degrees of reality [530]. These axes produce the *abstraction space* depicted in Fig. D.1 which provides the underlying basis for our survey design. We further segment each axis into non-expert relative categories from low to high abstraction. An asset that is high on both constructs is the most abstracted, e.g., Fig. D.1, top right.

D.4 Study Design

Our primary goal was to understand the differences in preferences between expert and non-expert audiences in visualizations of biomedical processes. We summarize our process in Fig. D.2. This study focused on spatial visual representations to enable a fair comparison of data-driven assets and illustrations. Prior evaluation studies in medical visualization have put less emphasis on illustrations, and have rather emphasized data-driven visualization works [404]. Our equal emphasis of both visual representation types allowed us to consider audience preferences in an expanded abstraction space. This approach included several challenges, the first of which was in establishing the boundaries of the design space with respect to visual representation and topic.

Design Space: Representation Constraints. The design space for depicting biomedical processes is enormous, and we do not intend our five topics to be comprehensive. They instead are meant to sufficiently cover the space of different criteria that an audience uses to evaluate a given topic representation. To constrain the design space, we first excluded interactivity; this has been explored elsewhere in a broader context [488].




Figure D.2: Three-phase study pipeline. **Setup**: define the design space, create audience scenarios and visual assets, and recruit survey participants, **Survey**: deploy survey asking participants to rank, quantify, and describe their top and bottom asset selections for each scenario, and **Results**: review survey results for patterns in selection abstraction space, scores, attribute rankings, and frequent keywords.

We included short animations to reflect the reality in our model abstraction construct that biomedical processes are highly dynamic. We included static elements that are often used to depict dynamic processes, e.g., glows and arrows [231]. We excluded animations that were only viewpoint changes, e.g., turntable animations, and focused on motion of the biomedical assets themselves. We also limit the abstraction space to typical representations of each topic without delving into stylistic methods, e.g., line, grayscale, or full color. This aspect of abstraction has been touched on elsewhere [220, 221, 288].

Design Space: Topic Constraints. Topics in biomedical processes also span a massive design space. Our aim was to evaluate the smallest reasonable topic set. Biomedical processes occur at all levels of magnification, from micro- to macroscale. They can be normal or pathological. To narrow the design space w.r.t. topic, we performed a literature review as well as interviews with visualization and biomedical illustration experts from both academia and industry. We also reviewed the Association of Medical Illustrators Online Salon [27] and several biomedical illustration portfolios to determine common topics visualized by both disciplines.

We chose two topics at the microscale: (1) signal transduction, a normal process whereby a signal is relayed between molecules in the body, and (2) constitutive activation, a process whereby one or more molecules in a signal chain is always switched “on” to create an never-ending signal relay. At the mesoscale we chose (3) normal blood flow and (4) an aneurysm. At the macroscale we chose (5) tumor metastasis, focusing on the movement of tumors from their origin site to other organs. The macroscale topic synthesizes concepts from both smaller magnification topics as it is driven through constitutive activation processes and travels through the bloodstream. Following topic selection we created audience scenarios for each topic that in turn guided asset production.

D.4.1 Survey Scenarios

We used scenarios to drive user comparison and selection, which we detail in Tables D.1 and D.2. This approach was inspired both by our expert interviews and by Lam et al.'s [278] findings that scenarios can effectively capture specific goals and research questions in a given domain. This corroborates well with biomedical illustration, where assets are most often created to fulfill the **communication objective** of a clearly-defined scenario. Our aim with these scenarios was to target relatively generic expert and non-expert audience use cases. We confirmed the validity of each described scenario with senior domain scientists, visualization scientists, and biomedical illustrators each with over ten years of experience. Our subsequent creation of visual assets was based on these audience scenarios. This workflow mirrors the standard approach to visualization production while also further constraining the design space.

Table D.1: **Expert Audience Survey Scenarios**

Topic	Scenario
Signal Transduction	An immunology researcher is publishing in an immunological venue on the newly-discovered pivotal role that a ligand plays in a signaling pathway. Their goal is to communicate the specificity of the activation pathway and its location in the cell with a visual supplement to their publication.
Constitutive Activation	An oncology researcher would like a visual supplement that demonstrates to the readership of an immunology journal the mechanism of disease in which a key molecule in the signal transduction chain is constitutively activated, which produces an unregulated positive feedback loop.
Blood Flow	A researcher studying vascular flow would like a visual to supplement their publication that explains the variation of laminar flow (i.e. smooth movement of fluid with no swirls), in normal hemodynamics (i.e., blood flow behavior).
Aneurysm	A researcher publishing in a medical venue would like to include a supplementary image or animation to describe the final shape of an aneurysm, resulting from abnormal hemodynamic forces (i.e., blood flow in helical or swirling patterns) and morphological properties of the vessel wall.
Metastasis	A radiation oncology researcher publishing in an oncology journal is focused on describing the metabolism and movement of metastatic tumors as the basis of validation for their novel radiation therapy approach.

D.4.2 Survey Visual Assets

We produced all assets via a series of topic-oriented focus groups to define the relevant design space and form consensus for each topic, following in part the framework for creative visualization-opportunities workshops described by Kerzner et al. [249]. Focus groups consisted of three to four people for each topic comprised of biomedical illustrators and/or visualization scientists. For each focus group we prepared sketches or concepts from our prior literature search and interviews to guide the discussion. Our interdisciplinary team of visualization scientists and biomedical illustrators enabled us to produce assets in-house to ensure consistency and to limit the number of variables in the survey. A key decision in our initial focus groups was to exclude labels and annotations, with the exception of occasional arrows when considered part of the model

Table D.2: **Non-Expert** Audience Survey Scenarios

Topic	Scenario
Signal Transduction	An introductory biology student is studying for an upcoming exam. Their goal is to understand how a “message” is relayed through a series of messengers inside a cell.
Constitutive Activation	The same introductory biology student is tasked with identifying where in the signaling pathway a molecule is constantly activated when it should not be. This causes the entire signaling pathway to be always switched “on.”
Blood Flow	A person with little/no prior knowledge on the topic is interested in learning more about their body. They visit a popular health and well-being website, e.g. WebMD, to understand how blood moves and delivers nutrients throughout the body.
Aneurysm	A person has recently been diagnosed with a cerebral aneurysm. Their doctor shows them a visual to communicate what aneurysms are and why they must be closely observed.
Metastasis	A patient recently diagnosed with cancer has been told by their doctor that their cancer may metastasize, meaning that the cancer may spread to a different part of the body from where it began. To help them understand this concept, their doctor shows them a visual.

abstraction axis, from all assets. This decision was made both to limit the variable space and to prevent distraction from the actual interpretability of the assets themselves. Our production pipeline included the Adobe Suite (Illustrator, Photoshop, AfterEffects) [5], Blender [51], and 3D Slicer [250]. Animated assets were produced as short, looping GIFs. The following briefly details the driving design concepts for each of the five chosen topics. For high resolution assets we refer the reader to the asset directory in supplementary material.

Signal Transduction. Signal transduction describes a cellular communication process in the body by which a sequence of molecules are activated or deactivated in response to an initiating signal. Visual approaches range from static to dynamic, from basic shape primitives to realistic molecular shapes taken from the Protein Data Bank (PDB) [43]. The environment may be simplified to only the main molecules up to fully immersive scenes with all molecules engaging in stochastic reactions with complex biomolecular assemblies [61, 125]. Glows, such as those utilized in CellPathway [415], are frequently used in biomedical illustration and less frequently in visualization to indicate the concept of activation. For further details we refer the reader to Kozlíková et al.’s survey of molecular representations [267].

We created 14 assets to represent common visualization options in this topic, shown in Fig. D.4A. Half of the visualizations use realistic molecular models extracted from PDB, e.g., C11, the other half use simple primitive shapes as often seen in biology textbook and journal figures, e.g., C1. We use a key icon in the primitive shape assets following a focus group discussion and our review of such illustrations in visualization literature, where a key is often used to indicate the special status of a ligand [407]. We illustrated half of the assets in a simplified context while the others show the main molecules in complete isolation. We used MCell to simulate molecule movement and stochastic interactions with CellBlender [247, 480, 481], a Blender plugin [51], to visualize our simulations. We excluded conformation changes in order to limit the design

space. These scenes served as representatives for robust data-driven models of the stochastic interactions in a real molecular environment. Although the simulation with realistic-looking molecules and interactions (C14) is the least abstracted of the set, we note that even this scene is heavily abstracted, as we just show the main molecules and include only a basic cell nucleus and membrane. Our color choices for the glows reference contemporary biomedical illustration trends to use a saturated color in the same hue range as the molecule base color.

Constitutive Activation. Constitutive activation describes a signal transduction process that is always turned “on”, meaning that the factors that keep a signal flowing between molecules are always present in the cellular environment. Although a number of processes in the body are naturally constitutively activated, mutations can cause a signalling pathway that is normally only conditionally activated to be constitutively activated. If left unchecked this process can lead to proliferation of tumor cells through uncontrolled cellular division. We created a corollary pathological variant that represents constitutive activation for each of the original 14 signal transduction assets (Fig. D.4B). We chose a generic mutation, showing a ligand that is not degraded or released from the first molecule in the chain after having activated the molecule. We followed conventions as indicated from our focus groups, showing the mutated molecule haloed in red with a red glow to indicate activation instead of the typical same-hue saturated color as in a normal signal cascade. We colored all other molecules and glows as in normal signal transduction, since they are not mutated. We kept all other scene aspects the same for assets C1-12. Since C13 and C14 included a more complex molecular environment with stochastic reactions, we factored in the effect of a constitutively-activated molecule where the result consists of many more activated molecules relative to normal signalling conditions.

Blood Flow. The flow of blood allows for delivery of oxygen and other essential substances to cells as well as the removal of waste products. While biomedical illustrators focus primarily on the appearance or on the constituents of blood cells, e.g., C4, C5, C11, and C12 in Fig. D.3, visualization scientists focus primarily on visualizing fluid dynamics that are linked to the acquisition modality, e.g., Phase-Contrast MRI (PC-MRI). Oeltze-Jafra et al. [384] provide a comprehensive summary of visualization techniques that are applied to blood flow. Our data-driven assets included streamlines, particles, streamribbons, streamtubes, and arrow glyphs using data from Berg et al. [42]. While hemodynamics are the focus, we rendered the vessel structure itself as translucent with ghosting of the mesh as exemplified by Baer et al. [34]. For closer alignment with the color palette of the illustration assets we used the inferno matplotlib color palette to render quantities.

Aneurysm. An aneurysm is an extensively visualized pathology caused by changes in the arterial wall and/or abnormal hemodynamics [457] with numerous methods developed to better understand aneurysm pathogenesis and rupture risk [384]. Unlike the microscale normal/pathological assets, the aneurysm/blood flow assets are not a 1:1 match. This was a conscious decision, as our goal for each topic was to produce the typical set of representations that would be used to convey the described scenario for an aneurysm. Some representations that are relevant for blood flow are irrelevant for




Figure D.3: Blood flow abstraction space. Assets are arrayed in the space by degree of model (y-axis) and visual abstraction (x-axis). Animated assets are denoted with a star glyph to the right of the asset name. Values in the four corners of each asset represent a weighted score for its selection frequency as the first, second, or third choice for an expert or a non-expert audience scenario (see 'How to read this chart,' left). Encircled regions indicate assets with scores in the 20th percentile of each scenario (see 'How to read this chart,' right).

communicating an aneurysm, e.g., the cellular composition of blood (Fig. D.3, C4). The external shape of a blood vessel (Fig. D.4C, C4) is a necessary and common visual representation to describe an aneurysm. While a number of the blood flow assets have an illustrative counterpart to the data-driven representation, in some cases such data are not available for aneurysms. For example, an aneurysm in the act of rupturing is difficult, if not impossible, to capture mid-rupture as in C6 of Fig. D.4C. We confirm from focus groups that this is a common illustration created to educate a non-expert audience on the risk of an untreated aneurysm.

Metastasis. Metastasis, when visualized at the macroscale, offers a synthesis and continuation of the lower scale topics: tumor proliferation is driven through constitutively activated signaling pathways, and tumors metastasize, i.e., spread, to other organs through the bloodstream. While we discussed using angiogenesis in early focus groups to represent tumor growth, the other four strongly movement-themed topics made metastasis, with its strong sense of movement, a more consistent choice. Our focus on the depiction of tumor spread exposed a notable visualization gap: medical technology does not allow for detection of the actual movement of tumors, so we cannot directly visualize this process. The closest available option for human subjects uses PET/CT data. This multimodal imaging strategy indicates regions of high metabolic activity, and is frequently used by clinicians to track metastasis.

Our illustrative assets demonstrate four levels of visual abstraction for tumor metastasis—half with highly abstracted tumor shapes while the others show realistic tumor shapes. We only included those organs and circulatory elements critical to telling the story in the most complex of the illustrative assets (C10-12 in Fig. D.4D), with gradual visual simplification of the organs at each step to the right of the visual

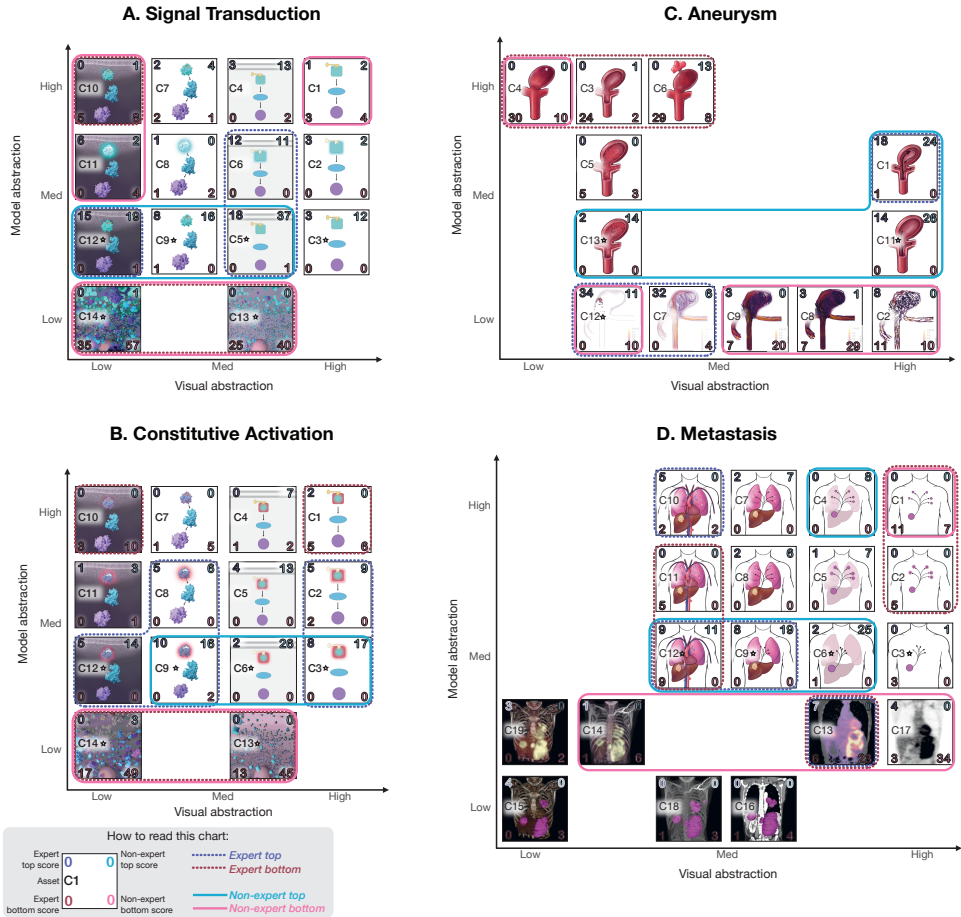


Figure D.4: Abstraction spaces for (A) Signal transduction, (B) Constitutive activation, (C) Aneurysm, and (D) Metastasis. Assets are arrayed in the space by degree of model (y-axis) and visual abstraction (x-axis). Animated assets are denoted with a star glyph to the right of the asset name. Values in the four corners of each asset represent a weighted score for its selection frequency as the first, second, or third choice for an expert or a non-expert audience scenario (see ‘How to read this chart,’ left). Encircled regions indicate assets with scores in the 20th percentile of each scenario (see ‘How to read this chart,’ right).

abstraction axis until in C1-3 they are entirely removed. The scientific visual assets follow the typical visualization techniques outlined by Lawonn et al. [287] in their state of the art report on multimodal medical visualization.

D.4.3 Survey Design Structure

We followed the principles for a comparative survey design laid out by Tory [504]. Topics are organized so that a healthy/normal physiological topic precedes a corresponding pathological topic. This format provides the necessary context for the pathology. We asked participants to rank only their top three and bottom three choices for each scenario to keep the survey scope manageable. The bottom choices are just as valuable as the top choices, as encouraging participants to explore negative aspects of a visualization can be illuminating. For the top- and bottom-ranked choices we subsequently asked participants to assign quantitative rankings of four variables: aesthetics, scientific accuracy, visual clarity, and communication success. Our variable selection was guided by works of Abdul-Ramen et al. [2] and by the judging criteria used for the Association of Medical Illustrators (AMI) juried salon. We additionally asked participants to select or enter their own keywords to describe the strengths and weaknesses of each of their ranked assets. We drew these keywords from the previously mentioned AMI salon judging criteria (see supplementary material). Lastly, we included an option for participants to add freeform comments.

We administered our survey via the Typeform [349]. Prior to deployment we conducted a pilot study with five participants to test our survey design. Following pilot study feedback we divided the survey into three segments by scale: micro-, meso-, macroscale to improve the overall completion rate. A second pilot study with three participants confirmed that the smaller segments kept average completion to 30 min.

D.4.4 Survey Recruitment

Our target participants included clinicians, biomedical illustrators, and domain and visualization scientists with familiarity in the selected biomedical topics. Our aim was to collect at least 20 high quality responses for each topic to adequately create a picture of audience preference. We recruited participants via the authors' respective professional networks. We collected only basic personal information, e.g., age, gender, and professional background. We additionally asked participants to report their expertise on each topic on a scale of 0 to 5, with 0 indicating "no knowledge" and 5 represents "extremely knowledgeable." We used this information to create two audience groups: (1) expert and (2) non-expert, where experts reported a **4 or higher** and non-expert audience participants reported a **3 or below**. We used the reported professions and expertise as a secondary check on the validity of their self-reported expertise level.

D

D.5 Study Findings

The survey ran for approximately three months, with each segment available to participants for one month. Participation was roughly gender balanced (M=male, F=female) for each topic (signal transduction: N=32, 16M, 16F; constitutive activation: N=28,




Figure D.5: Expert and non-expert attribute rankings for top and bottom choices for all five topics.

15M, 13F; blood flow: N=36, 20M, 16F; aneurysm: N=34, 19M, 15F; and metastasis: N=22, 10M, 12F). Participant backgrounds were mixed and included MR physicists, clinicians, visualization scientists, molecular biologists, and biomedical illustrators with training and background ranging from professors and program directors to executives to medical journal and agency staff. Self-rated expertise (E=expert, NE=non-expert) per topic varied (signal transduction: 12E, 20NE; constitutive activation: 7E, 21NE; blood flow: 25E, 12NE; aneurysm: 19E, 16NE; and metastasis: 8E, 14NE). The microscale and mesoscale segments contained two topics each and averaged 34 minutes to complete. The macroscale segment contained only one topic and averaged 18 minutes to complete. Participation falloff ranged from 3% to 26% over the course of a given segment. Higher falloff rates were likely due to a higher percentage of time-constrained clinicians who were unable to complete the survey. We dropped responses from participants who did not complete all questions for a given topic to avoid artificial biasing of asset choices.

For each of the five surveyed topics we report the following, with detailed per-topic results accessible at <https://public.tableau.com/profile/biomedsurvey2021>.

- **Asset scores:** Each asset received four weighted scores that represent the frequency that it was selected in the top or bottom three options in each scenario.
- **Average attribute ranking:** Average ranking values for aesthetics, scientific accuracy, visual clarity, and communication success for assets that were selected for each scenario (either as a top or bottom choice).
- **Keywords and comments:** Dominant keywords and representative comments used to describe the top- and bottom-scoring assets in each scenario.

Asset Scores. Asset scores are weighted such that $final\ score = 3s_1 + 2s_2 + s_3$, where s_1 , s_2 , and s_3 indicate the sum counts for an asset selected as 1st, 2nd, or 3rd for a given scenario. These scores are shown in the corners of each asset in Figs. D.3 and D.4A-D. We demarcate those assets falling in the top 20th percentile for expert top (dark blue), expert bottom (dark red), non-expert top (light blue), and non-expert bottom (pink) choice selections.

In all five topics we observe that the 20th percentile scores for both expert and non-expert top asset selections form clusters that often fall in the medium range of either one

or both abstraction axes. We see a dislike of the most extreme ranges of the abstraction space, with a few exceptions. For example, in the lower left corner that denotes both low model and low visual abstraction of the aneurysm abstraction space (Fig. D.4C), we see a cluster of expert top choices comprised of C12 (animated particle flow) and C7 (pathlines). For each topic selection, we see one or two clusters, or one cluster with one or more outliers. For example, blood flow in Fig. D.3 shows two separate clusters of expert top choices. Interestingly, in this case the split in clusters seems to be associated with the different professions. Clinicians/biomedical illustrators most often selected C11 (animated blood constituents) and C12 (animated red blood cells), while visualization/domain scientists selected C9 (streamtubes) and C10 (streamribbons) more often.

In all topics we see an overlap in preferences between audiences in the 20th percentile of top selections. With respect to expert top selections, we occasionally see a slightly larger spread in the abstraction spaces, particularly along the model abstraction axis. This is apparent in the blood flow (Fig. D.3), constitutive activation (Fig. D.4B), and aneurysm (Fig. D.4C) abstraction spaces. On the other hand, non-expert top selections that do not overlap with expert selections often fall into a higher abstraction space region. We see this in blood flow C4 (static blood cell components) in Fig. D.3, and in metastasis C4 (static abstracted tumors inside tinted organs) in Fig. D.4D.

We similarly see frequent overlaps in bottom scenario selections. Their spread in the abstraction space is also similar between audiences, with two exceptions. In constitutive activation (Fig. D.4B), we see a larger spread in bottom selections for the expert scenario, while in signal transduction (Fig. D.4A) and aneurysm (Fig. D.4C) the spread of bottom selections is larger for the non-expert scenario.

We additionally see occasional exceptions to top and bottom selection overlap for the expert and non-expert scenarios. For the aneurysm topic, C12 (animated particle flow) was selected as a bottom choice for a non-expert audience while also as the top choice for an expert audience (Fig. D.4C). Other interesting cases show selection overlap within an expert audience. In metastasis, both C12 and C13 (CT slice with colored PET heatmap overlay) falls into both expert top and bottom scenario selections (Fig. D.4D).

Attribute Rankings. Fig. D.5 shows the average attribute rankings (aesthetics, accuracy, visual clarity, and communication) for the top and bottom choices for the expert (top row) and non-expert scenarios (bottom row) for each of the five topics: signal transduction, constitutive activation, blood flow, aneurysm, and metastasis.

We show top selections in a blue hue (dark blue for experts, light blue for non-expert audience) and bottom selections in a red hue (dark red for experts, pink for non-expert audience). Attribute rankings over all four attributes average at 4.1 for expert top selections while bottom selections average at 1.8. Average rankings across all four attributes are similar for non-expert audience selections, with 3.7 for the top selection and 1.8 as the bottom selection. We observe similar average ranking assignments between top and bottom choices in the non-expert audience evaluation of asset accuracy for signal transduction, blood flow, and aneurysm. This makes sense, as a non-expert audience is unlikely to have the necessary expertise to determine the accuracy of a given asset. For the expert audience we see such similar ranking only in aesthetics in the top and bottom choices for aneurysm.

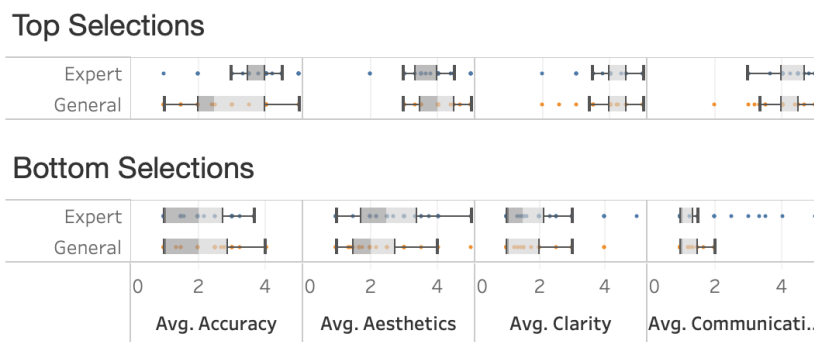


Figure D.6: Average attribute rankings to assets selected as either top or bottom for an expert (blue) or non-expert scenario (orange) for accuracy, aesthetics, clarity, and communication success.

The assigned attribute rankings in Fig. D.6 between expert and non-expert audiences are similarly distributed, although expert top selections often show a narrower distribution. Expert rankings for bottom choices show a long right tail, suggesting mixed perceptions of communication success for selected assets.

Keywords and Comments. Fig. D.7 reveals similar keyword preferences for both expert and non-expert audiences in their top selections, with *informative*, *easy to read*, and *clear* in the 20th percentile for both audiences. The only difference between the two audiences is the selection frequency of these keywords: experts prioritized *informative* over *easy to read*, while for a non-expert audience this order is reversed. We see a stronger difference in the 20th percentile of preferred keywords for bottom selections between audience levels. Experts used *confusing*, *simplistic*, and *pretty* most frequently to describe bottom choices. In contrast, the 20th percentile of keywords for non-expert audience bottom selections included *confusing*, *distracting* and *excessive*. Also intriguing is experts' frequent use of *pretty* to describe their bottom choices.

Scenario comments indicated a strong preference for the inclusion of labels, legends, and captions. Feedback on the use of arrows was also positive, although many participants felt that the positive feedback loop in constitutive activation was not effectively communicated and that a different approach was needed, e.g., an additional arrow that looped back from the last to the first molecule in the sequence. Comments were generally positive w.r.t. animated assets, with several comments indicating a preference for animated arrows particularly in non-expert scenarios. Comments related to data-driven assets, e.g., metastasis PET/CT and blood flow visualizations, often stressed that such assets were overly abstract for non-expert audiences, e.g., blood should not be perceived as composed of wires and tubes. At times such assets lacked an aspect of the stated communication objective. These included the lack of nutrients for blood flow scientific assets, lack of visuals showing real-time spreading of tumors for metastasis, or the lack of vessel wall layers and thickness for aneurysm.

Conversely, participant comments on illustrative assets that were expert top choices often indicated a desire for additional realism, e.g., more accurate motion, more accurately-sized cell or molecular components. One participant noted in their selection



Figure D.7: Word cloud of keywords chosen to describe top and bottom choices for expert and non-expert scenarios for all topics.

of the animated blood constituents asset (C11 in Fig. D.3), “The inclusion of multiple kinds of cells/molecules is helpful for accuracy. The animation could include more variability in flow among the objects for even more accuracy, but that could also potentially hinder the main communication goal if it becomes too distracting or hard to track.” Other assets were selected as bottom choices for being too misleading for the topic scenario, e.g., blood flow fluid illustrations in Fig. D.3 C8 and C13 “look too much like a clot,” or the removal of organs creating too much uncertainty for where tumors had spread in metastasis, “without any anatomy underneath, you have no way of knowing what the dots represent, or how deep into the tissue they are. Is it a rash spreading? Unclear.”

However, there was a clear limit to desired realism for either audience. Numerous comments focused on assets that were perceived as chaotic, noisy, and unnecessarily complex, e.g., the stochastic molecular interaction scenes included for signal transduction and constitutive activation (C13 and C14 in Fig. D.4A and B). This complexity made meaningful interpretability regarding the achievement of the communication goal impossible for both audiences. Assets with excessive realism occasionally veered into “scary” for non-expert audiences, e.g., the greyscale PET scan image with high metabolic activity regions (C17, Fig. D.4D).

D

D.6 Discussion

In the following we discuss the patterns we observed for audience preference and identify opportunities for improved visualization design for communication success, while also reflecting on the limitations of our study.

Preferred Abstraction. A meaningful visual abstraction eases visual processing and reduces cognitive load [530]. Our results indicate for both audiences that preferred abstractions often reside in a middle space of visual and model abstraction. They dislike either extreme realism or extreme abstraction. Initially we thought that experts would have a higher preference for these extrema for one of two reasons: (1) experts have such intimate knowledge of a subject that they do not need or want to see the complete picture, or (2) experts prefer completeness because their knowledge of a subject allows them to tolerate more complex information. Ultimately neither was consistently true. To some extent this corroborates previous works that found that the added value of dynamic visualizations is questionable and highly dependent upon the audience and communication objective [232, 402].

Selection Criteria. Interestingly, participant keyword choices indicate different selection criteria for bottom choices, but similar selection criteria for top choices. This matches our observations of the degree of selection overlap between the two audiences: top choices overlapped more extensively than the bottom choices. This indicates that participants may place equally high priority on positive visual clarity and communication-related factors, i.e., *informative, easy to read, clear*. However, their criteria to identify a poor visualization differ, and as does their idea for what constitutes *confusing*. Experts consider oversimplification to be confusing, while a non-expert audience reacts against overly distracting or excessive visualizations. The non-expert audience preference against confusing or distracting visualizations makes sense—without sufficient subject background, information-rich visualizations are often incomprehensible. Such information overload is exemplified in the molecular simulation assets (C13 and C14 in Fig. D.4A and B).

Aesthetics is not the only consideration in selecting a visualization. While the keyword *pretty* was selected often to describe both top and bottom choices, it was notably the third-most frequent keyword selected to describe expert bottom choices. For example, the bottom-most selection by experts to describe blood flow, C8 (static fluid visualization in Fig. D.3), was most described as *pretty*, but additionally as *simplistic, inaccurate, and misleading*. Thus it seems that clarity and communication may carry more weight for this audience type. This prioritization makes intuitive sense, as experts rely often on visualizations for technical information exchange. A quantitative control study focused on aesthetics relative to accuracy as perceived by different audiences would be an interesting follow-on work.

Background Biases. Background expertise and training play a large role in asset preferences, and likely affect our perception and understanding of a visualization. For example, in the blood flow topic when both components and hemodynamics were identified as important, experts with mostly clinical or biomedical illustration backgrounds prioritized the visualization of blood components (C11 and C12, middle region in Fig. D.3) over information encoding hemodynamic forces (C9 and C10, bottom region in Fig. D.3). In metastasis, we saw a similar background-based selection split for the PET/CT heatmap asset (C13, bottom region in Fig. D.4D). The experts selecting this as a top choice came from MR physics and visualization, while the experts selecting this as a bottom choice came from biomedical illustration or life sciences. The expert selection overlap with C12 in this topic is more difficult to explain. While background expertise likely plays a role, which we infer from one comment that it looks “too good

to be true,” its selection as both a top and bottom expert choice requires finer-grained information than captured in our study.

Our backgrounds can also influence our perception of the meaning of visual marks and channels, e.g., color. For example, while a clinician may be used to reading a PET/CT layered slice image with high metabolic activity regions as bright (C13) or dark (C17), someone without this background would interpret these differently, e.g., interpret the dark spots in C17 as dead tissue regions or the bright zones indicating a strange event in the body. A quick solution to disambiguate color meaning may involve labels and captions, but more immediately understandable solutions without this addition may be interesting to explore.

Stylistic Preferences. Stylistic elements are frequently used to emphasize a biomedical process. For example, while ubiquitously used in biomedical illustration, glows can mean many different things. Our focus group on metastasis discussed whether a tumor glow indicated pain, treatment application (radiotherapy), tumor metabolic activity, or was purely to draw attention. This lack of clarity became apparent in the survey, with one participant commenting, “It is unclear whether the glow in the tumors on the lungs is meant to denote a new growth or stylistic radiation treatment. If it is treatment, then perhaps there should be numbered steps or a device that provides the radiation.” At the microscale, the focus groups generally found glow indication to be meant to either draw attention or to indicate activity/aberrant activity. While this mixed meaning is convenient in our case, since we wanted to draw attention to areas of activity, it may quickly become problematic if that is not the communication goal. This suggests that glows should be used with care and their use reexamined in practice.

Study Limitations. We set a number of limitations and assumptions in this study given its large design space and broad topic range. For instance, our sampling of visualizations and topics was not comprehensive but representative of the massive space of creative and technological visualizations of biomedical processes. Additionally, the granularity of expertise in our survey is relatively coarse, and non-expert participants often had a higher basic scientific knowledge than someone from the broader public. A logical next step would be finer-grained surveys by expertise/target audience. This may introduce additional challenges in visual representation design, as many communication-oriented visualizations of biomedical processes that are aimed at the general public with no scientific background are heavily annotated or narrated, and often include multiple scales to orient the viewer, e.g., an initial view of the entire body is provided before diving inside an organ and on to the interior of the organ’s cell where a signal is passed between molecules in the cell. This type of visualization was out of the scope of our study, and given this we felt that including participants with a somewhat higher knowledge of biology would be beneficial for quality responses in some cases.

In choosing comparatively broad expert and non-expert scenarios our study favors those visualizations that are more flexible to interpretation. Even so, the visualizations for each topic naturally have different degrees of effectiveness based on the audience and the described scenario. Rather than identifying the single best visualization for a specific audience scenario, our overarching goal was instead to find general preferences and values for visualization selection.

D.7 Research Opportunities

This study opens a number of exciting opportunities for visualization research of biomedical processes. Gaps in biomedical illustration and visualization are readily apparent in all our topics. Illustration-driven works are currently filling in spaces in stories that cannot be easily told with data alone, e.g., aneurysm rupture, the cellular composition of blood, and the spread of tumors. These indicate that visualizing data is not always sufficient, and may in fact lead to a mismatch between audience and technique. However, data-driven visualization can offer a faster and realistic means to present phenomena that are laborious or impossible to create with current biomedical illustration workflows. While visualization research that applies illustrative techniques to patient data is relatively mature [288], illustrative techniques applied to represent a creator's mental model of a given phenomena or to represent a cohort are an open challenge [336].

Visualization research that intentionally considers layered messaging, e.g., one for communication targeted for a non-expert audience and one for analysis that targets an expert audience, may be interesting to consider. The overlapping preferences for assets between expert and non-expert audiences suggest that this may be amenable and more likely with increased demand for health communication. This layering may be achieved by superimposing visualization techniques in a manner similar to Pixar's storytelling approach: Pixar films are designed to entertain multiple levels of audiences, with numerous adult messages sprinkled throughout that do not affect the messages geared towards children. We imagine that this can be done with a thoughtful combination of data- and/or illustrative-driven assets. Linked juxtaposition may be another avenue to explore. For example, linking the process steps visualized in a highly abstracted asset, e.g., signal transduction with a basic glow sequence animation between primitive shapes, to a complex stochastic interaction visualization may help both experts and a non-expert audience to understand the sequence of a reaction framed in a realistic, complex environment.

D.8 Conclusion

The aim of our study was to better understand the development and evaluation process for visualizations of biomedical processes by different audiences. We particularly were interested in illuminating how visualization and biomedical illustration currently diverge and converge. Our findings show that both audience levels we surveyed place a high value on clarity and ability of a given asset to meet its stated communication objective. Moving forward, an optimal positioning for abstraction is likely in a middle space of both model and visual abstraction. We additionally found that some conventions are not as clear as we thought, e.g., glows can ambiguously indicate a call to attention, a pathological event, activation, etc., while other approaches were unexpectedly preferred, e.g., biomedical illustrations in place of data-driven visualizations. This latter preference occurred most often when the source data model was overly complex or did not capture the mechanism required to achieve the stated audience objective. Much of this study focused on communication. Future work that combines both biomedical il-

Illustration and visualization techniques in data analysis with domain experts also holds great potential.

Acknowledgements

This work is part of the project Visual Data Science for Large Scale Hypothesis Management in Imaging Biomarker Discovery (VIDI) funded by the University of Bergen and the Trond Mohn Foundation in Bergen (813558 and 811255). We are grateful to our study participants – without them, this work would not have been possible.

Paper E

Considering Best Practices in Color Palettes for Molecular Visualizations

Laura A. Garrison¹, and Stefan Bruckner¹

¹Dept. of Informatics, Univ. of Bergen, Norway

Abstract

Biomedical illustration and visualization techniques provide a window into complex molecular worlds that are difficult to capture through experimental means alone. Biomedical illustrators frequently employ color to help tell a molecular story, e.g., to identify key molecules in a signaling pathway. Currently, color use for molecules is largely arbitrary and often chosen based on the client, cultural factors, or personal taste. The study of molecular dynamics is relatively young, and some stakeholders argue that color use guidelines would throttle the growth of the field. Instead, content authors have ample creative freedom to choose an aesthetic that, e.g., supports the story they want to tell. However, such creative freedom comes at a price. The color design process is challenging, particularly for those without a background in color theory. The result is a semantically inconsistent color space that reduces the interpretability and effectiveness of molecular visualizations as a whole. Our contribution in this paper is threefold. We first discuss some of the factors that contribute to this array of color palettes. Second, we provide a brief sampling of color palettes used in both industry and research sectors. Lastly, we suggest considerations for developing best practices around color palettes applied to molecular visualization.




Figure E.1: Two color choices and effects to illustrate a molecular reaction between a ligand and its protein receptor. (a) Orange ligand with similar luminance and saturation to its receptor. (b) Ligand and receptor with higher luminance and saturation, with additional highly saturated glow effect.

E.1 Introduction

Suppose you have landed a project for the visualization of the mechanism of action of a novel cancer treatment drug that is about to come to market. One of the key elements of the brief is showing how the drug acts at a molecular level. The visualization you produce must be accurate, as well as beautiful, informative, and memorable: your client wants hospital administrators, doctors, and patients to be interested in, and to opt in to, this drug. In telling this story, color palette will be one of the main and most challenging creative decisions you will have to make in the production process. The brief may include a suggested color palette that aligns with the drug's branding, or you may be free to choose colors that you feel are most appropriate to tell the story that you want to tell. If you choose poorly, you risk a visualization that is unappealing, ineffective, or incomprehensible. If you are a biomedical illustrator, this is a common scenario that you are faced with. As a researcher who works with molecules, aspects of this scenario may also be quite familiar. Color selection can be overwhelming, particularly for novices, and outside of standard best practices for color there are no general guidelines in place for the coloring of molecules. This paper identifies some of the rationale for the broad use of color in molecular visualizations, provides a set of contemporary color palette examples, and discusses considerations towards best practices to lead to more interpretable, accurate, and consistent molecular visualizations without compromising on aesthetics or overly limiting creative freedom.

Visualizations can tell stories that aid in exploration, analysis, and communication of complex molecular phenomena to a range of audiences and user types [168, 169, 239]. For this paper, we generalize to three story features in a molecular visualization where color is an important consideration:

1. **Focus + context molecules.** Molecular visualizations are often structured in a visual hierarchy such that *focus* molecules are shown prominently and in full detail. *Context* molecules or structures are de-emphasized and provide an overview of and add visual interest to the scene. Both Fig. E.1a and E.1b use color to help develop this hierarchy. Color increases the prominence of the focus molecules (ligand and receptor) and allows the context molecules (lipid bilayer of the cell




Figure E.2: Simple biomedical illustration depicting key molecules in a pathway.

membrane) to recede into the background while still providing locational context for the scene.

2. **Molecular reaction(s).** Molecules can interact in reactions that fundamentally change their properties, synthesize new molecules, or destroy molecules. A specific and commonly–visualized scenario is *ligand binding*, where a *ligand* is defined as any substance that forms a complex with another molecule to serve a biological purpose [180]. The specific region of the molecule that the ligand binds to is known as a *binding site*. This step initiates (or blocks) a series of reactions that contribute to pathways integral to the life cycle and behavior of a cell, with natural implications in drug development and protein engineering research. Fig. E.1 shows two different color approaches for this event that experiment with saturation and luminance to draw attention and semantically connect to the concept of “binding and activation.”
3. **Molecular pathway.** A sequence of molecular reactions, often which are initiated by a ligand binding event, describe a molecular pathway. Understanding molecular pathways and their functions is critical to understanding the functioning of higher-order structures such as cells, tissues, organs, organ systems, and even the entire body. An example of this is shown in Fig. E.2 with three key molecules in a given intracellular pathway. Color helps provide functional semantics to the visualization: similar colors show that the three molecules are connected, and a color progression indicates the order of the molecules in the pathway.

Color plays a vital role in conveying each of these story features. While it can be represented in several formats, such as RGB (red, green, blue), CMYK (cyan, magenta, yellow, key: black), or HSL (hue, saturation, lightness) [417], for the purposes of this paper it is most intuitive and useful to think of color in the HSL color space. This encompasses the three color properties depicted on the 3D cylinder shown in Fig. E.3a. Hue specifies a base color, e.g., cyan, that is localized by angle around the color wheel illustrated in Fig. E.3b. Saturation defines the purity of a hue. Values span the inner to




Figure E.3: Color. (a) Source: https://en.wikipedia.org/wiki/HSL_and_HSV. (b) Adapted from: https://commons.wikimedia.org/wiki/File:RGB_color_wheel_360.svg

outer perimeter of the cylinder, from no saturation (grey) to full saturation, e.g., pure cyan. Lightness specifies color brightness, ranging from the bottom (black) to the top of the cylinder (white). Mixing black into a color produces a shade, while blending with white produces a tint.

A color palette is the combination of colors used to design a visualization. A number of color harmony rules aid in creating visually pleasing palettes. Derived from Itten's seven models of color contrast [222], harmony rules may be *monochromatic*, *analogous*, or *complementary*, among others. *Monochromatic* palettes are formed from tints and shades of a single color, as in Fig. E.4a. *Analogous* palettes comprise colors that are adjacent on the color wheel, as in Fig. E.4b. This type of palette is employed in Fig. E.2 to indicate that the molecules are part of the same pathway, and are therefore functionally connected. *Complementary* palettes are comprised of colors that are opposite each other on the color wheel, as in Fig. E.4c. Colors from these palettes can be used to draw attention to a particular element, to guide the eye through a narrative, or to establish a visual hierarchy of focus + context elements in a molecular visualization.

While the artist often has creative license to choose their palette, other factors come into play. Clients from different sectors have different aims. A pharmaceutical company has different requirements than an educational or research institution. The in-




Figure E.4: Three common color harmony rules with base color blue. Created in Adobe Color: <https://color.adobe.com/create/color-wheel>.

tended audience may have certain cultural sensitivities, visual or cognitive disabilities, or other factors to consider. While on an individual basis these requirements can provide guard rails that limit the color design space, color selection remains broad and inconsistent overall. This leads to diluted semantic meaning of molecular structures, which can impact a visualization's interpretability and effectiveness on a larger scale. For example, if COVID-19 spike proteins are colored blue, rather than the red used in the well-known version produced by Alissa Eckert and Dan Higgins for the CDC, can everyone still recognize it as the COVID-19 virus? What are the consequences if they can't?

Many education and research-oriented applications use the CPK coloring convention for atoms [92] when atomic-scale resolution is key to the visualization. At the cellular scale there are established colors for certain cell types, where the red blood cell is perhaps the most obvious example. This is always red, unless there is an express reason to show it otherwise, e.g., deoxygenated cells. Immune cells are often shown in cool colors that echo the soothing blue color often seen in the medical field. Molecules can be similarly classified, to some extent, into related groups according to structure or function. With a standard in place for semantically coloring atoms, and an informal semantic coloring practice for coloring cells, why not have something in place for molecules? Limited works in visualization have addressed color treatment in molecular visualizations. These consider the use of illumination models to cue features on molecular surfaces [195, 490] and coloring of multiscale molecular visualizations [256, 535]. While the application of high luminance colors to focus objects is a consistent recommendation of these works, color assignment on the whole is largely arbitrary and lacks consistent semantic meaning. Such works, alongside the broad community of professionals who craft them, can form the foundation for a set of best practices. This would enable easier creation of molecular visualizations that are more interpretable and effective, as well as aesthetically-pleasing.

E.2 Related Work

In this section we briefly discuss works that explore color associations, as well as the use of color in molecular visualizations within the areas of biomedical illustration and scientific visualization.

Color can elicit different emotional and psychological reactions and associations [4, 316]. According to Itten, in general, 'all tints (light colors) represent the brighter and better aspects of life, whereas shades (dark colors) symbolize the dark and negative forces [222].' However, different cultures have often different affective interpretations of color. This is well-summarized in the visualization "Colours in Culture" by David McCandless and AlwaysWithHonor.com. In this graphic, we see the color black associating with, e.g., *intelligence* for Asian cultures and *style* for Japanese and Hindu cultures. Even in the same culture, a color can take on different meaning in different contexts, suggesting a more subjective and nuanced interpretation of color. Considering black again for Native American cultures in the "Colours in Culture" graphic, we see it associates with both *balance* and *death*. As a Western culture

<https://phil.cdc.gov/Details.aspx?pid=23311>

<https://informationisbeautiful.net/visualizations/colours-in-cultures/>

example, Wexner's [553] study of color–mood associations with 94 psychology students at Purdue University found that participants strongly associated the color black to *despondent, dejected, unhappy, melancholy* as well as *powerful, strong, masterful*. Adams & Osgood [4] conducted a ground-breaking study on the affective meanings of color across 23 different cultural groups using bipolar adjectives. These adjectives were grouped into three factors: *Evaluation (E)*, *Potency (P)*, and *Activity (A)*. Examples for each factor include: good ↔ bad for *E*, strong ↔ weak for *P*, and active ↔ passive for *A*. They explored the perception of color in general, as well as seven distinct colors: white, grey, black, red, yellow, green, and blue. Among their findings, blue, white, and green were associated across nearly all 23 cultures as *good*, while black and grey more typically associated with *bad*. Black and grey furthermore were associated with *passive*, implying a degree of subjectivity in assessing the mood of these colors. Red is considered a strongly *active* color, although cultures disagree on its *evaluation*. Yellow exposes similar cultural disagreements on its *evaluation*. Filmmakers frequently take advantage of such color–mood associations to define the tone or mood of a film. Wei et al. [547] analyze the consistency of this on a global (entire film) and local (short shot sequences) scale. In this study, the authors collected and classified multidimensional feature vectors, including movie pace, movie dynamics, and dominant color ratio, to determine the mood of a film. Using the color–mood associations developed by Mahnke [316], an American psychologist, their approach found approximately 80% accuracy for mood/film genre classification according to color.

Color affect is also well-studied in visualization, with studies such as Bartram et al.'s work on the strength of associations between certain color palettes and affective response [37]. For example, *calm* often associates with cool colors with high lightness and low saturation. This strong association was found again in a related study by Kulahcioglu & de Melo [275] investigating affective word clouds. In contrast, *playful* or *exciting* do not exhibit such distinct color palette associations, and Bartram et al [37] note the need for more nuanced analyses of color harmony patterns. Color meaning can also be highly individual, particularly when associated with concepts or lesser-known objects [456]. In their crowd-sourced study of color palette selection for visualizations, Ahmad et al. [8] demonstrate this in quotes from two users who differ on the meaning of blue, red, and white in terms of *agree, disagree, and neutral*. However, color can also have strong semantic associations that positively impact performance. Lin et al. [303] found a faster response time for comparison–related tasks when data are assigned semantically meaningful colors for fruits (e.g., yellow ↔ banana), drinks, vegetables, and brands. For concepts or objects which are often lesser known and lack such strong semantic associations for the public, Schloss et al. [443] have shown that, with sufficient context, audiences can still infer meanings of colors. This holds promise for molecules, which are often unfamiliar to the public. Although there remain myriad reasons for subjective and variable interpretations of color semantics, e.g., culture, color blindness, ethnicity, we can use insights from these and similar works to appropriately leverage context and semantics to tell more consistent stories in molecular visualization.

Biomedical illustration is a field devoted to illustrating and animating biological and medical topics, often with a human focus. Biomedical illustrators follow perceptual principles in color design of a molecular visualization, but frequently take artistic license regarding the specific colors in a color palette. David Goodsell's watercolor

paintings of molecular machinery are foundational to the practice of illustrating and visualizing molecules, which use color strategically to encode the spatial organization of molecules [165, 169]. However, he notes that the majority of his colors are “*completely arbitrary and are chosen solely for aesthetic appeal* [168].” Biomedical illustrators also frequently employ perceptual color techniques to draw audience attention to the main narrative of the visualization [231]. For example, Jenkinson et al.’s [232] perceptual study on scene complexity versus learning outcomes uses desaturated, low contrast colors for context molecules/scene elements, while applying *complementary* and highly saturated/bright colors for focus molecules (ligand and protein receptor) in all treatments. Johnson & Hertig suggest the same such approach in their guide to the visual analysis and communication of biomolecular structural data [239]. Wong furthermore notes that small scene objects need increased hue, saturation, and/or brightness to stand out in a visualization [561], and suggests the simple trick of squinting at a visualization to assess for color visibility and evenness. While these approaches are useful for aesthetics and guiding the narrative, none suggest the use of specific colors to semantically highlight particular structural or functional features.

A wealth of publications address the use and efficacy of colormaps, such as the controversial rainbow colormap [57], in visualization. Our concern with color in this work is specific to scientific visualization, and more narrowly to the coloring of molecules and their environments. For an overview of color scales and guidelines for color use in a general visualization context we refer to surveys by Silva et al. [465] and Zhou et al. [579]. Biomedical illustrators and researchers who create molecular visualizations often rely on tools such as ColorBrewer [185], Colourmap Hospital [128], Colorgorical [173], or Adobe Color [6] to determine a color palette for their scene. Many of these tools are designed for chart visualization, as opposed to complex molecular structures in 2D or 3D. Palettes generated from such general-purpose tools can be ineffective or difficult to interpret when applied to molecular visualization. Adobe Color [6] is designed for artists and flexibly allows color selection according to defined color harmony rules, e.g., *complementary* colors. However, this still offers a staggering array of choices and requires a degree of color expertise to use. A similar tool with additional constraints for color selection could be more useful for content authors creating these assets who lack a background in color theory.

Color is used to provide structural cues on a molecular surface. These cues are aided by illumination models, such as Hermosilla et al.’s recent approach [195] that includes realistic diffuse color bleeding over a complete molecular scene. Ambient occlusion and directional lighting are shown by Szafer et al. [490] to help in interpreting molecular surface colors that are in shadow, while stylization can make interpretation more difficult. This study also found luminance-varying ramps to perform better than isoluminant ramps, since shadows reduce the luminance range. This suggests that focus molecules or regions may be better assigned high luminance values to direct attention and improve readability of the molecule’s surface. These works show the importance of selecting appropriate illumination models in a molecular visualization, and should be discussed in future best practices.

Limited works address color application across spatial scales in molecular visualization. Waldin et al. [535] present a technique that employs a systematically adjustable color scheme that mainly relies on hue shift across different levels of magnification, from atomic resolution to a complete virus. Their technique allows the viewer

to clearly distinguish between structures of interest at a given level of magnification between, e.g., atoms, domains in a single molecule, or structural compartments of a virus. Colors are generally saturated, and luminance is used as a focus device for particular features, e.g., the quantity of amino acids in a scene, or to show structural details, e.g., secondary structures in a protein domain. This use of luminance to drive a main narrative aligns with conventions in biomedical illustration. However, some instances may require greater contrast than an *analogous* color palette can achieve. Additionally, Waldin et al.'s coloring technique is limited to structure, rather than function. Lastly, the initial molecular color selections made by the user are arbitrary and lack semantic association. Klein et al. [256] apply a similar adaptive multi-scale coloring scheme in the context of microtubule dynamics, where molecules are mainly colored according to structure and have similar arbitrary initial coloring assignments. Developing guidelines from some of the basic rules established in these works, in conjunction with more specific structural or functional coloring rules, could lead to more effective molecular visualizations.

E.3 Color Choices in Molecular Visualization

The colors used to visualize molecules are dependent upon a number of considerations, such as corporate branding, personal taste, and cultural sensitivities. We demonstrate the breadth of color palette choices with a brief sample of color palettes drawn from the Association of Medical Illustrators 2021 Online Salon and from the 2021 VIZBI Poster Gallery in the Proteins category.

E.3.1 Color Considerations

Numerous factors influence the choice of a molecular color palette. Pharmaceutical marketing and research are major drivers for the production of molecular visualizations. Marketing videos for a new drug are contracted to specialty biomedical illustration studios every year. In many of these instances, the brief requests color palettes to follow the branding of the drug or the parent pharmaceutical company. In other instances the client may focus less on brand colors but instead request a palette reflective of a particular mood, e.g., comforting or dramatic. Such client-specific color requests can be helpful in constraining a design space that is at times overwhelming, but are partially responsible for the broad range of palettes and the lack of semantic consistency in the coloring of particular molecules. Beyond the pharmaceutical sector, different target areas often play a role in color palette choices. A molecular visualization aimed at academic/educational use often drives the author to make different color choices than a visualization for pharmaceutical use, as the main goal for this sector is often to engage and teach a broad, diverse audience. For example, Drew Berry's work mainly targets the educational and research sector, with visualizations for colorblind-friendliness and frequent use of yellows, blues, and purples applied to lambert shaders with ambient occlusion [546]. In contrast, XVIVO Scientific Animation Studio [570] tends towards more high-end, flashier lighting, shading, and rendering techniques, which contribute

<https://meetings.ami.org/2021/onlinesalon/>

<https://vizbi.org/Posters/>

to color palettes with greater drama and contrast for clients in the biotech and pharmaceutical sectors.

Culture is also a major driver of color selection, given its affective role in visualization. Many molecular visualizations incorporate blue, or a close analogous color, into their palette for the comforting, pleasant emotions that often associate with this color [4, 96]. Various works have shown that cool colors, e.g., blue, green, are more passive than warm colors, e.g., red, yellow, and orange [4, 37, 222, 332], and thus can be good context color choices. While the color red is an active, i.e., highly visually salient color [4, 332], depending on the culture it can indicate danger or, conversely, luck or happiness according to Chinese culture. Biomedical illustrators based in North America commonly use red to indicate aberrant molecular activity, e.g., constitutive activation, where a molecule is always turned “on.” For East Asian cultures this is not semantically intuitive, and may be taken to mean the opposite for audiences that do not have this exposure to North American conventions.

Likely the most significant element of color choice for molecular visualizations comes down to the author’s tastes and aesthetic preferences. Furthermore, biomedical illustrators and studios often wish to develop a house style that sets them apart from other studios. The decision process for color selection may often be guided by basic perceptual principles, e.g., saturated focus and desaturated context, and color harmony rules. However, the ultimate decision to color a ligand purple, orange, or another color is the author’s decision with little to no semantics attached. Exceptions include the case mentioned previously in a North American context, where red is often used to indicate aberrant activity of a molecule [148], and the use of red to color hemoglobin, which is the oxygen-bearing protein in red blood cells [165]. However, the majority of molecular visualization color palettes are guided by the author’s aesthetic sensibilities and their storytelling goals. The visualizations produced by Drew Berry, such as his set of DNA animations [46], are one example of this, where color application is highly aesthetic and tells a clear story, but the coloring of the individual molecules does not necessarily tie to their respective structural or functional properties.

E.3.2 Color Strategies: A Contemporary Sample

To illustrate the broad use of color in practice for molecular visualizations today, we conducted an informal study where we extracted the color palettes from 20 molecular visualizations that were produced in the last year. This ensures that our sampling captures recent trends in color design. Since such works may be created by biomedical illustrators, bioinformaticians, structural biologists, and visualization researchers [148], we sampled from two venues that attract these professions, and sampled ten palettes from each. The Association of Medical Illustrators (AMI) is a global, although primarily oriented to North America, society of biomedical illustrators. Every year the association hosts a juried salon where student and professional work can be submitted, and which is subsequently posted online. The ten palettes we sampled from the salon featured works where molecules were the main story element, and came from either static images or stills from a larger animation. Visualizing Biological Data (VIZBI) is an annual meeting that brings together diverse professions to discuss advances in research

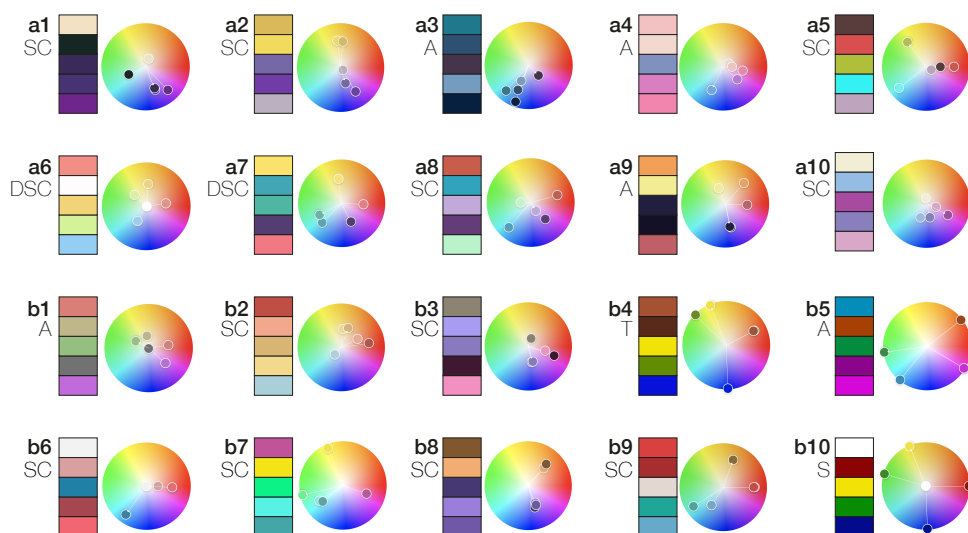


Figure E.5: 20 sampled molecular visualization color palettes. (a) 2021 AMI Online Salon color palettes. (b) 2021 VIZBI Protein poster color palettes. *Rule abbr:* SC: split complementary, DSC: double split complementary, A: analogous, T: triad, S: square.

in biological data, including molecular data. Each year includes a large poster session, which is divided into topical categories. The ten palettes we sampled came from the 2021 Proteins poster category. While VIZBI is advertised as an international venue, we acknowledge a North American bias in our sampling of biomedical illustration work from the AMI, since this is a primarily North American-oriented organization.

We generated color palettes from each molecular visualization using the Adobe Color tool [6], which has a feature to extract a color palette from an image. Palette extraction can be done manually by user selection of pixels of interest in the image, or semi-automatically according to *color mood*, e.g., *colorful* or *muted*. We used this semi-automatic feature as a way to minimize subjectivity bias for this phase. Experimentation with each of the five color mood options found that *colorful* consistently yielded palettes that captured the broadest color range in the image and most often included colors of molecules that may have proportionally occupied fewer pixels but were important to include in terms of the story. Palettes from this mode were the closest to what we, as experienced authors of molecular visualizations, would have generated by eye. However, in some instances the system did not pick up an important story feature. In these instances we manually adjusted the palette, replacing a redundant color, e.g., if the generated palette included two shades of a color, to include this color instead. This manual adjustment to the palette introduces a possible subjective bias in the study design.

After a palette is generated from an image, Adobe Color stores the resulting color palette as only a “custom” color harmony rule. Our intervention was required to match this custom palette to the closest-matching harmony rule, which we did by attempting to duplicate the palette by eye for each rule setting. This involved at times changing what the system had identified as the central color, or by making slight adjustments

to a palette color to identify the harmony rule that was the closest, if not exact, match for the generated palette. This human subjectivity is another possible source of bias, although we did our best to limit this through consistency of choices and monitor use. We further discuss this possible bias in Sec. E.5.

Fig. E.5 shows the set of 20 palettes that we generated from this process. Palettes a1–a10 are sourced from the AMI Salon, while b1–b10 are sourced from VIZBI Protein posters. For further details on the content authors, title and link to their original work, color harmony rules, and resulting palettes we refer to **supplementary material**. Although palettes are not always a precise match to a given harmony rule, we were able to match all 20 to a closest harmony rule. *Split complementary* is the most common rule, employed in 11 palettes in both groups, with two additional *double split complementary* palettes used in the AMI group. Palettes with an *analogous* rule are the second most common (five), and occur more frequently in the AMI group. The VIZBI group exhibits harmony rules that are closest to a *triad* (b4) and a *square* (b10) rule.

Generally, the palettes generated from the AMI group tend to be less saturated, show a greater contrast range, and favor *split/double split complementary* rules in their palettes. VIZBI group color rules for palettes are more variable, although *split complementary* is the dominant rule for this group as well. This latter point could reflect that some of the posters submitted were created in part by biomedical illustrators, e.g., b1 and b2. Cool base colors (purple, green, or blue) are more common (11) relative to warm colors, but this is fairly evenly-balanced between both groups.

Although we did not observe consistent color semantics across all 20 visualizations, some patterns emerge. We discuss these in the following, but note the need for a larger, international study before using such findings as the basis for formal and actionable guidelines. VIZBI posters that visualize COVID-19 more consistently use red for the spike proteins (b2, b6, b9), although this is not true for all (b4). The AMI group visualization of COVID-19 (a4), while using a warm color, uses pink instead. Purple is a popular color for receptor proteins for the AMI group (a2, a3, a5, a8), although also used for membrane molecules (a7, a8, a10), ligands (a4), or other elements (a9). VIZBI posters use purple less frequently, and more often for a ligand (b1, b5, b7) rather than for a receptor (b3) or other structural element (b8). In comparison, AMI works favor yellow/orange for a ligand (a1, a2, a3, a8, a10) or other focus molecule (a6, a7, a9). Yellow/orange is used more broadly, and less frequently, in the VIZBI group. These findings show that, while color selection on the whole is largely arbitrary, there is a clear aspect of peer influence in color selection, particularly in the AMI group. Best practices could be formalized from this, particularly if a larger survey with international groups and cultures uncovers a similar and stronger pattern of color application. In both groups on a per-visualization basis, individual molecular coloring is consistent for structural elements, e.g., molecules comprising a cell membrane are all colored the same or *analogously*. However, the semantics behind the color choices are unclear, except in the case of the COVID-19 spike protein's red coloring. Emphasizing a visual hierarchy or to focus attention on the main narrative appears to be the primary factor in color selection, with only sporadic consistency in coloring particular types of molecules, e.g., ligands.

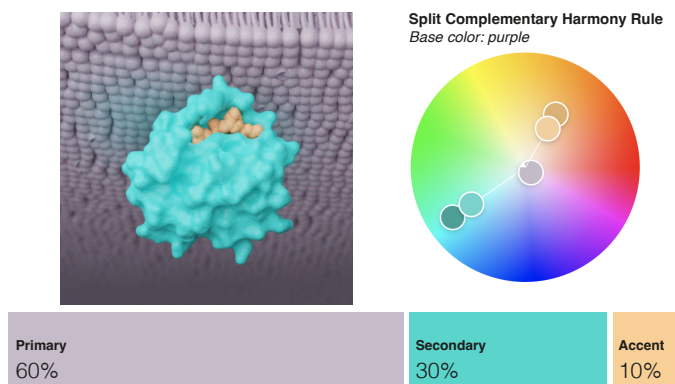


Figure E.6: Example of 60-30-10 rule used in a biomedical illustration. Explorable color palette at https://color.adobe.com/color-name_LG-color-theme-19646985/

E.4 Considerations for Molecular Coloring Best Practices

Molecular visualization remains a young and rapidly growing field. Allowing room for creativity is important for biomedical illustrators, structural biologists, and experts from related fields to innovate on the ways that we visualize molecules. However, framing coloring approaches within a set of best practices can provide a common ground that aids in the **aesthetics**, **interpretability**, and ultimate **effectiveness** of a molecular visualization while also simplifying the design process for content authors with limited training in color theory.

Aesthetics. Aesthetics are integral to drawing and guiding attention in a molecular visualization. Best practices can help content authors who lack formal training or intuition in color theory to easily craft more aesthetically-pleasing visualizations. The *60-30-10* rule from interior design is a useful rule of thumb to guide the composition and harmony of a molecular visualization. In this rule, 60% of the scene should be a dominant color, 30% a secondary color, and the last 10% an accent color, as demonstrated in Fig. E.6. The *split complementary* harmony rule that is popular with biomedical illustrators aligns with this practice. Molecular visualizations that follow established color harmony rules and basic perceptual principles can be more aesthetically-pleasing, and furthermore may be easier to interpret in their ability to guide user attention to salient story points [231, 239, 561].

Interpretability. Color is an important driving force in the interpretability, or readability, of a molecular visualization. It can help the audience focus on the intended parts of the story, which ultimately leads to a more effective visualization. For instance, readability guidelines for text-based presentations recommend roughly 80% contrast between focus and context elements [459]. Additionally, employing *active* versus *passive* color ranges can clearly define focus+context elements and establish a scene hierarchy. Although different cultures or other contexts may assign different semantic meanings, warm, saturated, or light value colors perceptually tend to draw attention, while cool, desaturated, dark value colors often recede visually when positioned with active colors [222]. This perceptual feature is already in broad use in biomedical illustration, although, as we observed in our small study, the specific colors themselves are incon-

sistent. Leveraging this natural perceptive feature can aid interpretability of a molecular visualization by using color salience to draw attention to the most important, i.e., focus elements, of a visualization. For example, ligands and their receptors can be assigned high contrast colors relative to context molecules. *Complementary* colors can then be employed to differentiate the ligand from its receptor. A *split complementary* palette of yellow (ligand) and purple (receptor), such as in Fig. E.5.a2, is a colorblind-friendly choice for clear differentiation of ligand and receptor. Molecules that comprise a pathway can similarly be colored for high contrast against context elements, and these pathway molecules could be *analogously*-colored to indicate functional relatedness, as in Fig. E.2. Importance functions could be useful for rule-based methods to aid in generating a molecular visualization and assigning appropriate hue, saturation, and lightness values to assets.

Effectiveness. Lastly, an effective molecular visualization is read *correctly* by the intended audience. For example, in a given visualization, can the audience correctly identify a ligand? Coloring best practices can help to create a semantic layer of communication that provides an intuition to a broad audience of certain structural or functional properties of a molecule. This is particularly valuable since molecules are themselves rather abstract, often looking like “partially-melted gummy bears” that are difficult to relate to macroscale-world structures. Just as how a lay audience may not necessarily know exactly what a red blood cell is or what it does at a technical level, through shape and color cues they can recognize its basic properties and relate it to blood on a larger scale. This strategy can be extended to the molecular level by coloring hemoglobin red. Consistent coloring across multiscale may facilitate understanding of properties of other molecules as well.

Coloring molecules according to the type of pathway that they are involved in is another consideration. Certain color families and harmony rules could be applied to, e.g., signal transduction pathways, while metabolic or gene expression pathways are assigned a different color family. Further color treatment in the form of different applications of color fresnel or glow effects could then be used if the pathway is showing aberrant activity, as is frequent in cancer. This guideline remains general enough to allow for creative license and follows principles for aesthetic and interpretable visualizations, while adding additional semantics for greater understanding.

An alternative consideration could be assigning color families and harmony rules according to structure. This could be related to the structure of a single molecule, e.g., the domains of a protein could be assigned particular shades or tints. This guideline may also be applied to the structure of entire molecular scenes, e.g., membrane structures that are comprised of thousands of phospholipid molecules is assigned to a particular color family, while native internal molecular structures are allocated to a different portion of the color space. Currently, coloring by structure is mostly consistent within a given visualization, but this is not the case when comparing between the majority of visualizations. Coloring according to structure has some natural overlap with coloring according to function. Guidelines for both perspectives could be applied together to give color greater meaning in a molecular visualization.

Lastly, cultural associations are important considerations in developing guidelines to ensure the semantics come across as intended. In our globalized world, adoption of best practices that generalize across cultures is preferable to, e.g., palettes biased

towards Western sensibilities. Works such as Adams et al.'s [4] cross-cultural study of the affective meanings of color are a useful starting point. The development of guidelines for coloring of molecular visualization that help achieve these three aims may also draw inspiration from works on color specification models for scientific data, such as by Nardini et al. [360].

E.5 Limitations

This work presents a discussion and limited study of the space of color choices used in industry and research, and additionally draws on the authors' own experience in crafting molecular visualizations. While our aim is to summarize and introduce considerations for developing guidelines and best practices for semantically meaningful molecular color palettes, we note some limitations in this paper.

The color palettes selected in our study and the rationales we discuss for authors' color palette selection draw primarily from North American or European venues and motivations. This bias is important to note, and acts as further motivation to conduct larger, international studies and focus groups for the formal development of best practices. We do not comprehensively summarize the various rationales for color palette selection, and rather discuss a subset of the most common rationales. Similarly, we do not aim to comprehensively summarize the space of color usage in molecular visualization. This is an interesting direction and may be useful in the establishment of robust guidelines, but is beyond the scope of this work. Although limited, our considerations are meant to provide a starting point for establishing and discussing the need for best practices in coloring molecular visualizations.

A second source of possible subjectivity and bias is the generation of color palettes from images that we sampled in our study. Although Adobe Color is a high-quality tool popular in both research and industry for color palette creation, use of another tool may lead to variations in the color palettes that we generated. Our decision to use a particular color mood and subsequent minor adjustments we made to determine color palettes and their closest harmony rules were based on the subjective opinion of the researcher generating the color palette. Additionally, our criteria for adjusting automatically-generated palettes to include a color important to the story is another source of subjectivity in palette generation. This could have had an effect on the resulting harmony rule. Bias could have been mitigated in these instances by having another researcher generate palettes and compare the results for consistency. However, we did not incorporate this into our pipeline, as this is a preliminary study to explore color use and to motivate future work in this space.

Despite the fact that we frame our discussion of color in this paper around the HSL color space, this may not be the optimal color space. Although spatially uniform and therefore easier to describe, it is not perceptually uniform. For the latter, a color space such as HCL (hue, *chroma*, and lightness) [213, 577] may be a more suitable alternative.

The spatial range that molecular visualization covers is large, spanning from tiny individual molecules up to large scenes of molecules that form even larger structures. An extensive discussion of the multiscale nature of the molecular world falls beyond the scope of this work, but this is a major challenge that best practices for coloring molecular visualizations must account for. Audience tasks may also be different at

each scale. For example, CPK coloring rules are useful to identify residues for binding sites, but a different rule is necessary when exploring protein secondary domains, or an entire cell at molecular resolution. Works such as Waldin et al. [535] provide a solid foundation for best practices in coloring across spatial scales.

E.6 Challenges & Outlook

Establishing a robust set of guidelines that retain the flexibility for creative expression and innovation is challenging and requires further research in the form of perceptual user studies. For example, Jenkinson et al. [232] present a perceptual study that assesses the effect of molecular scene complexity on learning outcomes for undergraduate biology students. Presenting four scenes of increasing complexity that depict protein-ligand binding conformational changes in 3D, the authors found that the most complex scene was the most effective in achieving the intended learning outcomes. A similar study design could be adopted to assess the effects of different color palettes on learning outcomes.

Focus groups on color use with experts in biomedical illustration, visualization, structural biology, and bioinformatics are also necessary when developing guidelines to ensure that all stakeholders have a voice in the process. Our prior work has explored stakeholder perspectives from biomedical illustration and animation studios and visualization researchers [148]. This was a limited study that centered on North American and European molecular visualization conventions, and obtained these stakeholder perspectives through 1:1 interviews and in small focus groups (3–4 people). Such qualitative research methods should be used for the development of more formal guidelines, which we envision could take place within a workshop format that invites stakeholders from all relevant disciplines to achieve adequate representation. The Creative Visualization-Opportunities (CVO) Workshop [249] is a promising model to follow to ensure diverse stakeholder participation and engagement.

Colors are not interpreted in the same way in all cultures. Guideline development requires numerous, inclusive discussions and perceptual studies with people from different backgrounds. Further, international studies into the use of color and its meaning beyond Western aesthetics are necessary. As previously discussed, CVO workshops [249] that recruit stakeholders not only from a diversity of professions but also a diversity of cultures, genders, and ethnicities are imperative to developing guidelines that are globally relevant. Avoiding pitfalls related to color blindness is another challenge. For example, red-green color blindness is well-known and accounted for in standard design practice. Deciding to color hemoglobin red to enable its semantic mapping to red blood cells and blood can be problematic if green is also a part of the visualization's palette. Further nuanced discussion of affective and perceptual conflicts is necessary.

The use of color in the context of glows, as demonstrated in Fig. E.1b, also requires discussion. Glows are notoriously ambiguous in their semantics: they can indicate generic activity, aberrant activity, metabolism, growth, pain, among others. Conducting studies exploring the use and meaning of glows in the context of molecular events is likely the best way forward to determine best practices for their use (or disuse).

E.7 Conclusion

Current color palettes for molecular visualizations are largely arbitrary and determined by a wide range of factors. Implementing coloring best practices can lead to more interpretable and effective molecular visualizations that are stylistically and semantically consistent, without overly compromising on creative freedom or aesthetics. Such guidelines can also simplify the color design process. These guidelines could be applicable not only for communication-oriented tasks, but for exploratory and analytically-oriented visualizations as well. Color can be a means to elevate scientific and health literacy at individual and population levels. With molecular visualization becoming more prevalent in mainstream culture, we feel that color best practices can, and should, be implemented to improve public understanding of molecules.

Acknowledgements

This work is part of the project Visual Data Science for Large Scale Hypothesis Management in Imaging Biomarker Discovery (VIDI) funded by the University of Bergen and the Trond Mohn Foundation in Bergen (813558). Parts of this work have been carried out in the context of the Mohn Medical Imaging and Visualization Centre (MMIV) and the Center for Data Science (CEDAS) at the University of Bergen.

Bibliography





































For Paper A literature:








Scales: Molecular Cellular Tissue Organ Multiscale Survey

Tasks: Exploration Analysis Communication













*Color fill amount corresponds to primary task(s) each work addresses










- [1] ABDI, H., AND WILLIAMS, L. J. Principal component analysis. *Wiley Interdiscip Rev Comput Stat* 2, 4 (2010), 433–459, doi: 10.1002/wics.101. 14
- [2] ABDUL-RAHMAN, A., CHEN, M., AND LAIDLAW, D. H. A survey of variables used in empirical studies for visualization. In *Foundations of Data Visualization*. 2020, pp. 161–79. 3.3.1, D.4.3
- [3] ACKERMAN, M. J. The visible human project. *Proc IEEE* 86, 3 (1998), 504–511, doi: 10.1109/5.662875. A.9
- [4] ADAMS, F. M., AND OSGOOD, C. E. A cross-cultural study of the affective meanings of color. *J Cross Cult Psychol* 4, 2 (1973), 135–156, doi: 10.1177/002202217300400201. 2.3, 3.3.2, E.2, E.3.1, E.4
- [5] ADOBE. Adobe Creative Cloud. <https://www.adobe.com/creativecloud.html>, 2021. Viewed 18.06.2021. 3.3.1, D.4.2
- [6] ADOBE. Adobe color. Adobe Color: create color wheel, 2022. <https://color.adobe.com/create>. 1.3, 2.1, 3.3.2, E.2, E.3.2
- [7] AERTS, J., GEHLENBORG, N., MARAI, G. E., AND NIESELT, K. K. Visualization of biological data - crossroads (dagstuhl seminar 18161). *Dagstuhl Reports* 8, 4 (2018), 32–71, doi: 10.4230/DagRep.8.4.32. 3.1, A.1
- [8] AHMAD, J., HUYNH, E., AND CHEVALIER, F. When red means good, bad, or canada: exploring people’s reasoning for choosing color palettes. In *IEEE Trans Vis Comput Graph* (2021), pp. 56–60. 2.3, E.2
- [9] AHRENS, J., GEVECI, B., AND LAW, C. Paraview: An end-user tool for large data visualization. *Visualization Handbook 717*, 8 (2005), doi: 10.1016/B978-012387582-2/50038-1. . 3.1.2, 3.3.1, A.10, A.13
- [10] AIGNER, W., MIKSCH, S., MÜLLER, W., SCHUMANN, H., AND TOMINSKI, C. Visual methods for analyzing time-oriented data. *IEEE Trans Vis Comput Graph* 14, 1 (2007), 47–60, doi: 10.1109/TVCG.2007.70415. . 1.2.2, 2.1, A.1, A.4.2, A.13
- [11] AIGNER, W., MIKSCH, S., SCHUMANN, H., AND TOMINSKI, C. *Visualization of Time-Oriented Data*. Springer London, 2011. . A.13
- [12] ALHARBI, N., ALHARBI, M., MARTINEZ, X., KRONE, M., ROSE, A. S., BAADEN, M., LARAMEE, R. S., AND CHAVENT, M. Molecular visualization of computational biology data: A survey of surveys. In *Proc EuroVis* (2017). . A.5, A.13
- [13] ALHARBI, N., KRONE, M., CHAVENT, M., AND LARAMEE, R. S. Lod pli: Level of detail for visualizing time-dependent, protein-lipid interaction. In *Proc VISIGRAPP* (2019), pp. 164–174. . A.5.2, A.13












- [14] AMIRKHANOV, A., FRÖHLER, B., KASTNER, J., GRÖLLER, E., AND HEINZL, C. InSpectr: Multimodal exploration, visualization, and analysis of spectral data. *Comput Graph Forum* 33, 3 (2014), 91–100, doi: 10.1111/cgf.12365. 2.2, B.2
- [15] ANANTHANARAYANAN, R., ESSER, S. K., SIMON, H. D., AND MODHA, D. S. The cat is out of the bag: Cortical simulations with 10^9 neurons, 10^{13} synapses. Tech. rep., 2009.    . A.9, A.13
- [16] ANDERSON, E. W., PRESTON, G. A., AND SILVA, C. T. Using python for signal processing and visualization. *Comput Sci Eng* 12, 4 (2010), 90–95, doi: 10.1109/MCSE.2010.91.    . A.13
- [17] ANDREWS, B. Introduction to perceptual principles in medical illustration. In *Proc ACM SIGGRAPH Courses: Illustrative Visualization for Medicine and Science* (2006). 2.3, D.2
- [18] ANDREWS, K., OSMIĆ, M., AND SCHAGERL, G. Aggregated parallel coordinates: integrating hierarchical dimensions into parallel coordinates visualisations. In *Proc i-KNOW* (2015), pp. 1–4. 2.2, 14, C.1, C.2.2
- [19] ANDRIDGE, R. R., AND LITTLE, R. J. A review of hot deck imputation for survey non-response. *Int Stat Rev* 78, 1 (2010), 40–64, doi: 10.1111/j.1751-5823.2010.00103.x. 14
- [20] ANGELELLI, P., AND HAUSER, H. Straightening tubular flow for side-by-side visualization. *IEEE Trans Vis Comput Graph* 17, 12 (2011), 2063–2070, doi: 10.1109/TVCG.2011.235.    . A.16, A.8.1, A.13
- [21] ANGELELLI, P., SNARE, S. R., NYRNES, S. A., BRUCKNER, S., HAUSER, H., AND LØVSTAKKEN, L. Live ultrasound-based particle visualization of blood flow in the heart. In *Proc SCCG* (2014), pp. 42–49.    . A.8.1, A.13
- [22] ARMITAGE, P., BEHRENBRUCH, C., BRADY, M., AND MOORE, N. Extracting and visualizing physiological parameters using dynamic contrast-enhanced magnetic resonance imaging of the breast. *Med Image Anal* 9, 4 (2005), 315–329, doi: 10.1016/j.media.2005.01.001.    . A.7.2, A.13
- [23] AROUSSI, A., HOWDEN, L., AND VLOEBERGHES, M. 3D visualisation of cerebrospinal fluid flow within the human central nervous system. In *Distributed Frameworks for Multimedia Applications* (2006), pp. 1–7.    . A.13
- [24] ARTERO, A. O., DE OLIVEIRA, M. C. F., AND LEVKOWITZ, H. Uncovering clusters in crowded parallel coordinates visualizations. In *Proc IEEE IV* (2004), pp. 81–88. 2.2, C.1
- [25] ASAKAWA, D. S., NAYAK, K. S., BLEMKER, S. S., DELP, S. L., PAULY, J. M., NISHIMURA, D. G., AND GOLD, G. E. Real-time imaging of skeletal muscle velocity. *J Magn Reson Imaging* 18, 6 (2003), 734–739, doi: 10.1002/jmri.10422.    . A.8.5, A.13
- [26] ASHWORTH, S. L., SANDOVAL, R. M., TANNER, G. A., AND MOLITORIS, B. A. Two-photon microscopy: Visualization of kidney dynamics. *Kidney Int* 72, 4 (2007), 416–421, doi: 10.1038/sj.ki.5002315.    . A.6.1, A.13
- [27] ASSOCIATION OF MEDICAL ILLUSTRATORS. 2020 Association of Medical Illustrators Salon Winners. <https://www.ami.org/medical-illustration/view-art-and-animations/2020-salon-winners>, 2020. Viewed 18.06.2021. D.4
- [28] AUDIGIER, V., HUSSON, F., AND JOSSE, J. A principal component method to impute missing values for mixed data. *Adv Data Anal Classif* 10, 1 (2016), 5–26, doi: 10.1007/s11634-014-0195-1. 14, C.2.4, C.4
- [29] AUGUSTIN, C. M., NEIC, A., LIEBMANN, M., PRASSL, A. J., NIEDERER, S. A., HAASE, G., AND PLANK, G. Anatomically accurate high resolution modeling of human whole heart electromechanics: a strongly scalable algebraic multigrid solver method for nonlinear deformation. *J Comput Phys* 305 (2016), 622–646, doi: 10.1016/j.jcp.2015.10.045.    . A.4, A.8.2, A.9, A.13
- [30] AYACHE, N., BOISSEL, J.-P., BRUNAK, S., AND CLAPWORTHY, G. Towards virtual physiological human: Multilevel modelling and simulation of the human anatomy and physiology. Tech. rep., VPH Institute, 2005. 1.1, 3.1, A.1

- [31] AYMOZ, D., SOLÉ, C., PIERRE, J.-J., SCHMITT, M., DE NADAL, E., POSAS, F., AND PELET, S. Timing of gene expression in a cell-fate decision system. *Mol Syst Biol* 14, 4 (2018), e8024, doi: 10.15252/msb.20178024. 1.2.1, A.2
- [32] BACA, J., CARRUTH, D. W., CALHOUN, A., STEPHENS, M., AND LEWIS, C. Challenges in evaluating efficacy of scientific visualization for usability and aesthetics. In *Design, User Experience, and Usability. Practice and Case Studies* (2019), pp. 397–407. 2.3, D.2
- [33] BACH, B., DRAGICEVIC, P., ARCHAMBAULT, D., HURTER, C., AND CARPENDALE, S. A descriptive framework for temporal data visualizations based on generalized space-time cubes. In *Comput Graph Forum* (2017), vol. 36, pp. 36–61. 2.1, A.1
- [34] BAER, A., GASTEIGER, R., CUNNINGHAM, D., AND PREIM, B. Perceptual evaluation of ghosted view techniques for the exploration of vascular structures and embedded flow. *Comput Graph Forum* 30, 3 (2011), 811–20. 2.3, D.2, D.4.2
- [35] BAIDAK, B., HUSSAIN, Y., KELMINSON, E., JONES, T. R., FRANKE, L., AND HAEHN, D. Cellprofiler analyst web (cpaw)-exploration, analysis, and classification of biological images on the web. In *IEEE Trans Vis Comput Graph* (2021), pp. 131–135. . A.6.1, A.13
- [36] BARTOCCI, E., CHERRY, E. M., GLIMM, J., GROSU, R., SMOLKA, S. A., AND FENTON, F. H. Toward real-time simulation of cardiac dynamics. In *Proc CMSB* (2011), pp. 103–112. . A.7.2, A.13
- [37] BARTRAM, L., PATRA, A., AND STONE, M. Affective color in visualization. In *Proc CHI* (2017), pp. 1364–1374. 2.3, 3.3.2, E.2, E.3.1
- [38] BASOLI, F., GIANNITELLI, S. M., GORI, M., MOZETIC, P., BONFANTI, A., TROMBETTA, M., AND RAINER, A. Biomechanical characterization at the cell scale: present and prospects. *Front Physiol* 9 (2018), 1449, doi: 10.3389/fphys.2018.01449. A.6, A.11
- [39] BAZILEVS, Y., GOHEAN, J. R., HUGHES, T. J. R., MOSER, R. D., AND ZHANG, Y. Patient-specific isogeometric fluid-structure interaction analysis of thoracic aortic blood flow due to implantation of the jarvik 2000 left ventricular assist device. *Comput Methods Appl Mech Eng* 198, 45–46 (2009), 3534–3550, doi: 10.1016/j.cma.2009.04.015s. . A.13
- [40] BEHRENDT, B., KÖHLER, B., GRÄFE, D., GROTHOFF, M., GUTBERLET, M., AND PREIM, B. Semi-automatic vessel boundary detection in cardiac 4D PC-MRI data using FTLE fields. In *Proc EG VCBM* (2016), pp. 41–45. . A.3, A.13
- [41] BELLE, M., GODEFROY, D., COULY, G., MALONE, S. A., COLLIER, F., GIACOBINI, P., AND CHÉDOTAL, A. Tridimensional visualization and analysis of early human development. *Cell* 169, 1 (2017), 161–173, doi: 10.1016/j.cell.2017.03.008. . A.7.1, A.13
- [42] BERG, P., ROLOFF, C., BEUNG, O., VOSS, S., SUGIYAMA, S.-I., ARISTOKLEOUS, N., ANAYIOTOS, A. S., ASHTON, N., REVELL, A., BRESSLOFF, N. W., ET AL. The computational fluid dynamics rupture challenge 2013 – phase II: variability of hemodynamic simulations in two intracranial aneurysms. *J Biomech Eng* 137, 12 (2015), 121008. D.4.2
- [43] BERMAN, H. M., WESTBROOK, J., FENG, Z., GILLILAND, G., BHAT, T. N., WEISSIG, H., SHINDYALOV, I. N., AND BOURNE, P. E. The protein data bank. *Nucleic Acids Res* 28, 1 (2000), 235–42. D.4.2
- [44] BERRANEN, Y., HAYASHIBE, M., GILLES, B., AND GUIRAUD, D. 3d volumetric muscle modeling for real-time deformation analysis with fem. In *Proc EMBC* (2012), pp. 4863–4866. . A.8, A.8.5, A.13
- [45] BERRY, D. DNA animations for science-art exhibition. Walter and Eliza Hall Institute “The Future Starts Here” Exhibit, 2018. , https://www.youtube.com/watch?v=7Hk9jct2ozY&ab_channel=WEHImovies. A.5.3, A.13







































- [46] BERRY, D. DNA animations for science-art exhibition. Walter and Eliza Hall Institute “The Future Starts Here” Exhibit, 2018. https://www.youtube.com/watch?v=7Hk9jct2ozY&ab_channel=WEHImovies. E.3.1
- [47] BERTIN, J., BERG, W. J., AND WAINER, H. *Semiology of graphics: diagrams, networks, maps*. University of Wisconsin Press, 1983. 2.2, B.2
- [48] BESANÇON, L., SEMMO, A., BIAU, D., FRACHET, B., PINEAU, V., SARIALI, E. H., SOUBEYRAND, M., TAOUACHI, R., ISENBERG, T., AND DRAGICEVIC, P. Reducing affective responses to surgical images and videos through stylization. *Comput Graph Forum* 39, 1 (2019), 462–83. 2.3, D.2
- [49] BIEGING, E. T., FRYDRYCHOWICZ, A., WENTLAND, A., LANDGRAF, B. R., JOHNSON, K. M., WIEBEN, O., AND FRANÇOIS, C. J. In vivo three-dimensional mr wall shear stress estimation in ascending aortic dilatation. *J Magn Reson Imaging* 33, 3 (2011), 589–597, doi: 10.1002/jmri.22485. ■ □ ■ . A.8.1, A.13
- [50] BIRKELAND, Å., ŠOLTÉSZOVÁ, V., HÖNIGMANN, D., GILJA, O. H., BREKKE, S., ROPINSKI, T., AND VIOLA, I. The ultrasound visualization pipeline. In *Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization*. 2014. ■ . A.8, A.13
- [51] BLENDER ONLINE COMMUNITY. *Blender - a 3D modelling and rendering package*. Blender Foundation, 2018. 3.3.1, D.1, D.4.2
- [52] BLUMENSCHNEIN, M., BEHRISCH, M., SCHMID, S., BUTSCHER, S., WAHL, D. R., VILLINGER, K., RENNER, B., REITERER, H., AND KEIM, D. A. Smartexplore: Simplifying high-dimensional data analysis through a table-based visual analytics approach. In *Proc IEEE VAST* (2018), pp. 36–47. 3.2.1, B.6
- [53] BOCK, J., FRYDRYCHOWICZ, A., STALDER, A. F., BLEY, T. A., BURKHARDT, H., HENNIG, J., AND MARKL, M. 4D phase contrast MRI at 3 T: Effect of standard and blood-pool contrast agents on SNR, PC-MRA, and blood flow visualization. *Magn Reson Med* 63, 2 (2009), 330–338, doi: 10.1002/mrm.22199. ■ □ ■ . A.8.1, A.13
- [54] BOLINSKY, D. The inner life of the cell. *Proc ACM SIGGRAPH*, 2006. ■ □ ■ , <http://multimedia.mcb.harvard.edu/>. A.6.1, A.13
- [55] BORKIN, M. A., BYLINSKII, Z., KIM, N. W., BAINBRIDGE, C. M., YEH, C. S., BORKIN, D., PFISTER, H., AND OLIVA, A. Beyond memorability: Visualization recognition and recall. *IEEE Trans Vis Comput Graph* 22, 1 (2015), 519–28. 2.3, D.2
- [56] BORKIN, M. A., VO, A. A., BYLINSKII, Z., ISOLA, P., SUNKAVALLI, S., OLIVA, A., AND PFISTER, H. What makes a visualization memorable? *IEEE Trans Vis Comput Graph* 19, 12 (2013), 2306–15. 2.3, D.2
- [57] BORLAND, D., AND II, R. M. T. Rainbow color map (still) considered harmful. *IEEE Comput Graph Appl* 27, 2 (2007), 14–17, doi: 10.1109/MCG.2007.323435. 2.1, E.2
- [58] BORLAND, D., YI, H., GRANT, G. D., KEDZIORA, K. M., CHAO, H. X., HAGGERTY, R. A., KUMAR, J., WOLFF, S. C., COOK, J. G., AND PURVIS, J. E. The cell cycle browser: an interactive tool for visualizing, simulating, and perturbing cell-cycle progression. *Cell Syst* 7, 2 (2018), 180–184, doi: 10.1016/j.cels.2018.06.004. ■ □ ■ . A.6.1, A.13
- [59] BORN, S., PFEIFLE, M., MARKL, M., GUTBERLET, M., AND SCHEUERMANN, G. Visual analysis of cardiac 4D MRI blood flow using line predicates. *IEEE Trans Vis Comput Graph* 19, 6 (2013), 900–912, doi: 10.1109/TVCG.2012.318. ■ □ ■ . A.13
- [60] BOSTOCK, M., OGIEVETSKY, V., AND HEER, J. D3 data-driven documents. *IEEE Trans Vis Comput Graph* 17, 12 (2011), 2301–2309, doi: 10.1109/TVCG.2011.185. B.7, C.3
- [61] BOTTARO, S., AND LINDORFF-LARSEN, K. Biophysical experiments and biomolecular simulations: A perfect match? *Science* 361, 6400 (2018), 355–60. D.4.2











- [62] BOURQUI, R., AUBER, D., LACROIX, V., AND JOURDAN, F. Metabolic network visualization using constraint planar graph drawing algorithm. In *Proc IEEE IV* (2006), pp. 489–496. . A.5.3, A.13
- [63] BRADLEY, C., BOWERY, A., BRITTEN, R., BUDELMANN, V., CAMARA, O., CHRISTIE, R., COOKSON, A., FRANGI, A. F., GAMAGE, T. B., HEIDLBAUF, T., ET AL. OpenCMISS: a multi-physics & multi-scale computational infrastructure for the VPH/Physiome project. *Prog Biophys Mol Biol* 107, 1 (2011), 32–47, doi: 10.1016/j.pbiomolbio.2011.06.015. . A.1, A.13
- [64] BRANDES, U., DWYER, T., AND SCHREIBER, F. Visual understanding of metabolic pathways across organisms using layout in two and a half dimensions. *J Integr Bioinform* 1, 1 (2004), 1–16, doi: 10.2390/biecoll-jib-2004-2. . A.13
- [65] BREEN, D. E., WIDMANN, T., BAI, L., JU, F., DAHMANN, C., ET AL. Epithelial cell reconstruction and visualization of the developing drosophila wing imaginal disc. In *Proc BioVis* (2012), pp. 77–84. . A.7.1, A.13
- [66] BREHMER, M., AND MUNZNER, T. A multi-level typology of abstract visualization tasks. *IEEE Trans Vis Comput Graph* 19, 12 (2013), 2376–2385, doi: 10.1109/TVCG.2013.124. 1.2.2, 2.1, 2.2, 3.1.1, 3.2.1, A.1, A.6, A.4.2, D.2
- [67] BREHMER, M., AND MUNZNER, T. A multi-level typology of abstract visualization tasks. *IEEE Trans Vis Comput Graph* 19, 12 (2013), 2376–2385, doi: 10.1109/TVCG.2013.124. B.4.1
- [68] BREITWIESER, L., HESAM, A., DE MONTIGNY, J., VAVOURAKIS, V., IOSIF, A., JENNINGS, J., KAISER, M., MANCA, M., DI MEGLIO, A., AL-ARS, Z., RADEMAKERS, F., MUTLU, O., AND BAUER, R. Biodynamo: A modular platform for high-performance agent-based simulation. *Bioinformatics* 38, 2 (09 2021), 453–460, doi: 10.1093/bioinformatics/btab649. . A.7.2, A.15, A.13
- [69] BROCKMEIER, A. J., KRIMINGER, E. G., SANCHEZ, J. C., AND PRÍNCIPES, J. C. Latent state visualization of neural firing rates. In *Proc IEEE NER* (2011), pp. 144–147. . A.7.2, A.13
- [70] BRODBECK, D., AND GIRARDIN, L. Visualization of large-scale customer satisfaction surveys using a parallel coordinate tree. In *Proc IEEE IV* (2003), pp. 197–201. 2.2, C.1
- [71] BROOS, A., DE HOON, N., DE KONING, P., VAN DER GEEST, R., VILANOVA, A., AND JALBA, A. A framework for fast initial exploration of pc-mri cardiac flow. In *Proc EG VCBM* (2016), pp. 69–78. . A.8.1, A.13
- [72] BRUCKNER, S., AND MÖLLER, T. Isosurface similarity maps. *Comput Graph Forum* 29, 3 (2010), 773–782, doi: 10.1111/j.1467-8659.2009.01689.x. 2.2, B.2
- [73] BRYDEN, A., JR., G. N. P., AND GLEICHER, M. Automated illustration of molecular flexibility. *IEEE Trans Vis Comput Graph* 18, 1 (2012), 132–145, doi: 10.1109/TVCG.2010.250. . A.5.1, A.13
- [74] BURGESS, A., VUONG, J., ROGERS, S., MALUMBRES, M., AND O’DONOGHUE, S. I. SnapShot: phosphoregulation of mitosis. *Cell* 169, 7 (2017), 1358–1358, doi: 10.1016/j.cell.2017.06.003. . A.5.3, A.13
- [75] BURROWES, K., DE BACKER, J., SMALLWOOD, R., STERK, P., GUT, I., WIRIX-SPEETJENS, R., SIDDIQUI, S., OWERS-BRADLEY, J., WILD, J., MAIER, D., ET AL. Multi-scale computational models of the airways to unravel the pathophysiological mechanisms in asthma and chronic obstructive pulmonary disease (AirPROM). *Interface Focus* 3, 2 (2013), 20120057, doi: 10.1098/rsfs.2012.0057. . A.9, A.13
- [76] BYŠKA, J., LE MUZIC, M., GRÖLLER, M. E., VIOLA, I., AND KOZLÍKOVÁ, B. AnimoAminoMiner: Exploration of protein tunnels and their properties in molecular dynamics. *IEEE Trans Vis Comput Graph* 22, 1 (2015), 747–756, doi: 10.1109/TVCG.2015.2467434. . A.5.1, A.13
- [77] BYŠKA, J., TRAUTNER, T., MARQUES, S. M., DAMBORSKÝ, J., KOZLÍKOVÁ, B., AND WALDNER, M. Analysis of long molecular dynamics simulations using interactive focus+ context visualization. *Comput Graph Forum* 38, 3 (2019), 441–453, doi: 10.1111/cgf.13701. . A.5.2, A.13



























- [78] BYŠKA, J., JURČÍK, A., GRÖLLER, M. E., VIOLA, I., AND KOZLÍKOVÁ, B. Molecollar and tunnel heat map visualizations for conveying spatio-temporo-chemical properties across and along protein voids. *Comput Graph Forum* 34, 3 (2015), 1–10, doi: 10.1111/cgf.12612. . A.7, A.5.1, A.13
- [79] CAKMAK, E., JACKLE, D., SCHRECK, T., KEIM, D., AND FUCHS, J. Multiscale visualization: A structured literature analysis. *IEEE Trans Vis Comput Graph* (2021), doi: 10.1109/TVCG.2021.3109387. . 2.1, A.1, A.9, A.13
- [80] CANDAN, K., CARO, L. D., AND SAPINO, M. L. PhC: Multiresolution visualization and exploration of text corpora with parallel hierarchical coordinates. *ACM Trans Intell Syst Technol* 3, 2 (2012), 22, doi: 10.1145/2089094.2089098. 2.2, C.1
- [81] CASPI, R., BILLINGTON, R., FERRER, L., FOERSTER, H., FULCHER, C. A., KESELER, I. M., KOTHARI, A., KRUMMENACKER, M., LATENDRESSE, M., MUELLER, L. A., ONG, Q., PALEY, S., SUBHRAVETI, P., WEAVER, D. S., AND KARP, P. D. The metacyc database of metabolic pathways and enzymes and the biocyc collection of pathway/genome databases. *Nucleic Acids Res* 44, D1 (2015), D471, doi: 10.1093/nar/gkv1164. . A.5.3, A.13
- [82] CAVASSILA, S., DEVAL, S., HUEGEN, C., VAN ORMONDT, D., AND GRAVERON-DEMILLY, D. Cramér–Rao bounds: an evaluation tool for quantitation. *NMR Biomed* 14, 4 (2001), 278–283, doi: 10.1002/nbm.701. B.6
- [83] CHAARI, L., VINCENT, T., FORBES, F., DOJAT, M., AND CIUCIU, P. Fast joint detection-estimation of evoked brain activity in event-related fMRI using a variational approach. *IEEE Trans Med Imaging* 32, 5 (2013), 821–837, doi: 10.1109/TMI.2012.2225636. . A.13
- [84] CHABINIOK, R., WANG, V. Y., HADJICHARALAMBOUS, M., ASNER, L., LEE, J., SERMESANT, M., KUHLM, E., YOUNG, A. A., MOIREAU, P., NASH, M. P., ET AL. Multiphysics and multiscale modelling, data–model fusion and integration of organ physiology in the clinic: ventricular cardiac mechanics. *Interface Focus* 6, 2 (2016), 20150083, doi: 10.1098/rsfs.2015.0083. . A.9, A.13
- [85] CHEN, C., JIN, Y., UNVERZAGT, F. W., CHENG, Y., HAKE, A. M., LIANG, C., MA, F., SU, L., LIU, J., BIAN, J., ET AL. The association between selenium and lipid levels: a longitudinal study in rural elderly chinese. *Arch Gerontol Geriatr* 60, 1 (2015), 147–152, doi: 10.1016/j.archger.2014.09.005. C.3.2
- [86] CHEN, H., GUO, L., NIE, J., ZHANG, T., HU, X., AND LIU, T. A dynamic skull model for simulation of cerebral cortex folding. In *Proc MICCAI* (2010). . A.7.1, A.13
- [87] CHEN, M., AND EDWARDS, D. J. “isms” in visualization. In *Foundations of Data Visualization*. 2020, pp. 225–41. 2.3, D.2
- [88] CICKOVSKI, T. M., HUANG, C., CHATURVEDI, R., GLIMM, T., HENTSCHEL, H. G. E., ALBER, M. S., GLAZIER, J. A., NEWMAN, S. A., AND IZAGUIRRE, J. A. A framework for three-dimensional simulation of morphogenesis. *IEEE/ACM Trans Comput Biol Bioinform* 2, 4 (2005), 273–288, doi: 10.1109/TCBB.2005.46. . A.7.1, A.13
- [89] CLARAFI. Molecular maya (mmaya). Clarafi Tools, 2022. . <https://clarafi.com/tools/mmaya/>. A.5.1, A.13
- [90] CLEVELAND, W. S., AND MCGILL, R. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *J Am Stat Assoc* 79, 387 (1984), 531–554, doi: 10.2307/2288400. 3.2.1, B.6, C.2.2
- [91] COPPIN, P., HARVEY, J., VALEN-SENDSTAD, K., STEINMAN, D., AND STEINMAN, D. Illustration-inspired visualization of blood flow dynamics. In *Proc IEEE IV* (2014), pp. 333–335. . A.8.1, A.13
- [92] COREY, R. B., AND PAULING, L. Molecular models of amino acids, peptides, and proteins. *Rev Sci Instrum* 24, 8 (1953), 621–627, doi: 10.1063/1.1770803. 3.3.2, B.5, E.1











- [93] CRANE, J. C., OLSON, M. P., AND NELSON, S. J. SIVIC: open-source, standards-based software for DICOM MR spectroscopy workflows. *J Biomed Imaging* 2013 (2013), 12, doi: 10.1155/2013/169526. B.3
- [94] CRUZ, A., ARRAIS, J. P., AND MACHADO, P. Interactive network visualization of gene expression time-series data. In *Proc IEEE IV* (2018), pp. 574–580. . A.5.3, A.13
- [95] CSERMELY, P., KORCSMÁROS, T., KISS, H. J., LONDON, G., AND NUSSINOV, R. Structure and dynamics of molecular networks: a novel paradigm of drug discovery: a comprehensive review. *Pharmacol Ther* 138, 3 (2013), 333–408, doi: 10.1016/j.pharmthera.2013.01.016. . A.5, A.13
- [96] CYR, D., HEAD, M., AND LARIOS, H. Colour appeal in website design within and across cultures: A multi-method evaluation. *Int J Hum Comput Stud* 68, 1-2 (2010), 1–21, doi: 10.1016/j.ijhcs.2009.08.005. 3.3.2, E.3.1
- [97] DABDOUB, S. M., RUMPF, R. W., SHINDHELM, A. D., AND RAY, W. C. Moflow: visualizing conformational changes in molecules as molecular flow improves understanding. *BMC Bioinform* 9, 6 (2015), 1–12, doi: 10.1186/1753-6561-9-S6-S5. . A.5.1, A.13
- [98] DAHL, A. C. E., CHAVENT, M., AND SANSOM, M. S. P. Bendix: intuitive helix geometry analysis and abstraction. *Bioinformatics* 28, 16 (2012), 2193–2194, doi: 10.1093/bioinformatics/bts357. . A.5.1, A.13
- [99] DAVISON, T., SAMAVATI, F., AND JACOB, C. Lifebrush: painting, simulating, and visualizing dense biomolecular environments. *Comput Graph* 82 (2019), 232–242, doi: 10.1016/j.cag.2019.05.006. . A.6.1, A.13
- [100] DAYAN, J. H., OLIKER, A., SHARONY, R., BAUMANN, F., GALLOWAY, A., COLVIN, S. B., MILLER, D., AND GROSSI, E. A. Computer-generated three-dimensional animation of the mitral valve. *J Thorac Cardiovasc Surg* 127, 3 (2004), 763–769, doi: 10.1016/S0022-5223(03)00959-0. . A.8.2, A.13
- [101] DE HERAS CIECHOMSKI, P., KLANN, M., MANGE, R., AND KOEPL, H. From biochemical reaction networks to 3d dynamics in the cell: The zigcell3d modeling, simulation and visualisation framework. In *Proc BioVis* (2013), pp. 41–48. . A.4.2, A.6.1, A.13
- [102] DE HOON, N., JALBA, A., EISEMANN, E., AND VILANOVA, A. Temporal interpolation of 4D PC-MRI blood-flow measurements using bidirectional physics-based fluid simulation. In *Proc EG VCBM* (2016), pp. 59–68. . A.8.1, A.13
- [103] DE HOON, N., VAN PELT, R., JALBA, A., AND VILANOVA, A. 4D MRI flow coupled to physics-based fluid simulation for blood-flow visualization. *Comput Graph Forum* 33, 3 (2014), 121–130, doi: 10.1111/cgf.12368. . A.13
- [104] DE HOON, N. H., LAWONN, K., JALBA, A. C., EISEMANN, E., AND VILANOVA, A. Inkvis: A high-particle-count approach for visualization of phase-contrast magnetic resonance imaging data. In *Proc EG VCBM* (2019), pp. 177–188. . A.8.1, A.13
- [105] DERMIT, M., DODEL, M., AND MARDAKHEH, F. K. Methods for monitoring and measurement of protein translation in time and space. *Mol Biosyst* 13, 12 (2017), 2477–2488, doi: 10.1039/c7mb00476a. 3.1.2, 4, A.10, A.11
- [106] DEWAN, S., MCCABE, K. J., REGNIER, M., AND MCCULLOCH, A. D. Insights and challenges of multi-scale modeling of sarcomere mechanics in ctn and tm dcm mutants—genotype to cellular phenotype. *Front Physiol* 8 (2017), 151, doi: 10.3389/fphys.2017.00151. 3.1.1, A.4.1
- [107] DIXHOORN, A. F., VISSERS, B. H., FERRARINI, L., MILLES, J., AND BOTHA, C. P. Visual analysis of integrated resting state functional brain connectivity and anatomy. In *Proc EG VCBM* (2010), pp. 57–64. . A.8.4, A.13
- [108] DONDERS, A. R. T., VAN DER HEIJDEN, G. J., STIJNEN, T., AND MOONS, K. G. A gentle introduction to imputation of missing values. *J Clin Epidemiol* 59, 10 (2006), 1087–1091, doi: 10.1016/j.jclinepi.2006.01.014. 14, C.2.4, C.4

- [109] DONG, Y., AND PENG, C.-Y. J. Principled missing data methods for researchers. *Springerplus* 2, 1 (2013), 222, doi: 10.1186/2193-1801-2-222. C.4
- [110] DOWLING, M., WENSKOVITCH, J., FRY, J., HOUSE, L., AND NORTH, C. Sirius: Dual, symmetric, interactive dimension reductions. *IEEE Trans Vis Comput Graph* 25, 1 (2018), 172–182, doi: 10.1109/TVCG.2018.2865047. 2.2, C.1
- [111] DOWSON, N., BOURGEAT, P., ROSE, S., DAGLISH, M., SMITH, J., FAY, M., COULTHARD, A., WINTER, C., MACFARLANE, D., THOMAS, P., CROZIER, S., AND SALVADO, O. Joint factor and kinetic analysis of dynamic fdopa pet scans of brain cancer patients. In *Proc MICCAI* (2010). ■ ■ ■ . A.13
- [112] DUNIA, R., EDGAR, T. F., AND NIXON, M. Process monitoring using principal components in parallel coordinates. *AIChE J* 59, 2 (2013), 445–456, doi: 10.1002/aic.13846. 2.2, C.1
- [113] DUNIN-BARKOWSKI, W., LOVERING, A., OREM, J., BAEKEY, D., DICK, T., RYBAK, I., MORRIS, K., O’CONNOR, R., NUDING, S., SHANNON, R., AND LINDSEY, B. L-plotting—a method for visual analysis of physiological experimental and modeling multi-component data. *Neurocomputing* 74, 1–3 (2010), 328–336, doi: 10.1016/j.neucom.2010.03.015. ■ ■ ■ . A.7.2, A.13
- [114] DURAN, D., HERMOSILLA, P., ROPINSKI, T., KOZLIKOVA, B., VINACUA, A., AND VÁZQUEZ, P.-P. Visualization of large molecular trajectories. *IEEE Trans Vis Comput Graph* 25, 1 (2018), 987–996, doi: 10.1109/TVCG.2018.2864851. ■ ■ ■ . A.5.2, A.13
- [115] DURIKOVIC, R., KANEDA, K., AND YAMASHITA, H. Animation of biological organ growth based on l-systems. *Comput Graph Forum* 17, 3 (1998), 1–13, doi: 10.1111/1467-8659.00248. ■ ■ ■ . A.7.1, A.13
- [116] EDELSTEIN, S. Modeling of bimolecular reactions. In *Encyclopedia of Computational Neuroscience*, D. Jaeger and R. Jung, Eds. 2013, pp. 1–3. A.5.2
- [117] EDWARDS, M. R., GILROY, F. V., JIMENEZ, B. M., AND O’SULLIVAN, W. J. Alanine is a major end product of metabolism by giardia lamblia: a proton nuclear magnetic resonance study. *Mol Biochem Parasitol* 37, 1 (1989), 19–26, doi: 10.1016/0166-6851(89)90098-4. B.8
- [118] EGEBLAD, M., EWALD, A. J., ASKAUTRUD, H. A., TRUITT, M. L., WELM, B. E., BAINBRIDGE, E., PEETERS, G., KRUMMEL, M. F., AND WERB, Z. Visualizing stromal cell dynamics in different tumor microenvironments by spinning disk confocal microscopy. *Dis Model Mech* 1, 2–3 (2008), 155–167, doi: 10.1242/dmm.000596. ■ ■ ■ . A.6.2, A.13
- [119] EHRHARDT, J., WERNER, R., SCHMIDT-RICHBERG, A., AND HANDELS, H. Statistical modeling of 4d respiratory lung motion using diffeomorphic image registration. *IEEE Trans Med Imaging* 30, 2 (2011), 251–265, doi: 10.1109/TMI.2010.2076299. ■ ■ ■ . A.8.3, A.13
- [120] ELIASMITH, C., STEWART, T. C., CHOO, X., BEKOLAY, T., DEWOLF, T., TANG, Y., AND RASMUSSEN, D. A large-scale model of the functioning brain. *Science* 338, 6111 (2012), 1202–1205, doi: 10.1126/science.1225266. A.9
- [121] ELMQVIST, N., AND FEKETE, J. Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines. *IEEE Trans Vis Comput Graph* 16, 3 (2009), 439–454. C.2, C.2.3, C.2.3
- [122] ENGLUND, R., AND ROPINSKI, T. Evaluating the perception of semi-transparent structures in direct volume rendering techniques. In *Proc ACM SIGGRAPH* (2016), pp. 9:1–9:8. 2.3, D.2
- [123] ENGLUND, R., ROPINSKI, T., AND HOTZ, I. Coherence maps for blood flow exploration. In *Proc EG VCBM* (2016), pp. 79–88. ■ ■ ■ . A.8.1, A.13
- [124] FALK, M., DAUB, M., SCHNEIDER, G., AND ERTL, T. Modeling and visualization of receptor clustering on the cellular membrane. In *Proc BioVis* (2011), pp. 9–15. ■ ■ ■ . A.6.1, A.13

- [125] FALK, M., KLANN, M., REUSS, M., AND ERTL, T. Visualization of signal transduction processes in the crowded environment of the cell. In *Proc IEEE PacificVis* (2009), pp. 169–176.    . A.6.1, A.13, D.4.2
- [126] FALK, M., KLANN, M., REUSS, M., AND ERTL, T. 3D visualization of concentrations from stochastic agent-based signal transduction simulations. In *Proc ISBI* (2010), pp. 1301–1304.    . A.6.1, A.13
- [127] FALK, M., KRONE, M., AND ERTL, T. Atomistic visualization of mesoscopic whole-cell simulations using ray-casted instancing. *Comput Graph Forum* 32, 8 (2013), 195–206, doi: 10.1111/cgf.12197.    . A.6.1, A.13
- [128] FANG, H., WALTON, S., DELAHAYE, E., HARRIS, J., STORCHAK, D., AND CHEN, M. Categorical colormap optimization with visualization case studies. *IEEE Trans Vis Comput Graph* 23, 1 (2016), 871–880, doi: 10.1109/TVCG.2016.2599214. 2.1, E.2
- [129] FANG, Z., MÖLLER, T., HAMARNEH, G., AND CELLER, A. Visualization and exploration of time-varying medical image data sets. In *Proc Graphics Interface* (2007), pp. 281–288.    . A.13
- [130] FANGERAU, J., HÖCKENDORF, B., WITTBRODT, J., AND LEITTE, H. Similarity analysis of cell movements in video microscopy. In *Proc BioVis* (2012), pp. 69–76.    . A.6.2, A.13
- [131] FEDOROV, A., BEICHEL, R., KALPATHY-CRAMER, J., FINET, J., FILLION-ROBIN, J.-C., PUJOL, S., BAUER, C., JENNINGS, D., FENNESSY, F., SONKA, M., ET AL. 3D slicer as an image computing platform for the quantitative imaging network. *Magn Reson Imaging* 30, 9 (2012), 1323–1341, doi: 10.1016/j.mri.2012.05.001.    . 3.1.2, A.10, A.13
- [132] FEIG, M., AND SUGITA, Y. Reaching new levels of realism in modeling biological macromolecules in cellular environments. *J Mol Graph Model* 45 (2013), 144–156, doi: 10.1016/j.jmngm.2013.08.017. . A.6, A.13
- [133] FEIG, M., AND SUGITA, Y. Whole-cell models and simulations in molecular detail. *Annu Rev Cell Dev Biol* 35 (2019), 191–211, doi: 10.1146/annurev-cellbio-100617-062542. . A.6, A.13
- [134] FENG, D., AND CAI, W. Visualization of biomedical processes: local quantitative physiological functions in living human body. In *Proc CGI* (2000), pp. 319–323. A.8.4, A.13
- [135] FENG, D., KWOCK, L., LEE, Y., AND TAYLOR, R. M. Linked exploratory visualizations for uncertain MR spectroscopy data. In *Proc Visualization and Data Analysis* (2010), vol. 7530, p. 753004. B.2
- [136] FENG, D., LEE, Y., KWOCK, L., AND TAYLOR, R. M. Evaluation of glyph-based multivariate scalar volume visualization techniques. In *Proc APGV* (2009), pp. 61–68. 2.2, B.2
- [137] FENNER, J. W., BROOK, B., CLAPWORTHY, G., COVENEY, P. V., FEIPEL, V., GREGERSEN, H., HOSE, D. R., KOHL, P., LAWFORD, P., MCCORMACK, K. M., PINNEY, D., THOMAS, S. R., JAN, S. V. S., WATERS, S., AND VICECONTI, M. The europhysiome, step and a roadmap for the virtual physiological human. *Philos Trans R Soc A* 366, 1878 (2008), 2979–2999, doi: 10.1098/rsta.2008.0089. 1.1, 3.1, A.1
- [138] FIORAVANTE, M., SHOOK, A., THORPE, I., AND RHEINGANS, P. Visualizing motional correlations in molecular dynamics using geometric deformations. In *Proc EuroVis* (2013), pp. 311–320.    . A.5.1, A.13
- [139] FORKERT, N. D., FIEHLER, J., ILLIES, T., MÖLLER, D. P. F., HANDELS, H., AND SÄRING, D. 4D blood flow visualization fusing 3D and 4D MRA image sequences. *J Magn Reson Imaging* 36, 2 (2012), 443–453, doi: 10.1002/jmri.23652.    . A.8.1, A.13
- [140] FRIMAN, O., HENNEMUTH, A., HARLOFF, A., BOCK, J., MARKL, M., AND PEITGEN, H.-O. Probabilistic 4d blood flow mapping. In *Proc MICCAI* (2010).    . A.13
- [141] FRÖHLER, B., AMIRKHANOV, A., KASTNER, J., GRÖLLER, E., AND HEINZL, C. Multimodal visualization and analysis of spectral and XCT data. In *Proc Research Forum of the Austrian Universities of Applied Sciences* (2015). 2.2, B.2

- [142] FRYDRYCHOWICZ, A., BLEY, T. A., DITTRICH, S., HENNIG, J., LANGER, M., AND MARKL, M. Visualization of vascular hemodynamics in a case of a large patent ductus arteriosus using flow sensitive 3D CMR at 3T. *J Cardiovasc Magn Reson* 9, 3 (2007), 585–587, doi: 10.1080/10976640601015219. ■
 . A.13
- [143] FUA, Y., WARD, M., AND RUNDENSTEINER, E. Hierarchical parallel coordinates for exploration of large datasets. In *Proc IEEE VIS* (1999), pp. 43–508. 2.2, C.1
- [144] FUA, Y. H., WARD, M., AND RUNDENSTEINER, E. Navigating hierarchies with structure-based brushes. In *Proc IEEE InfoVis* (1999), pp. 58–64. 2.2, C.1
- [145] FUNAHASHI, A., MOROHASHI, M., KITANO, H., AND TANIMURA, N. Celldesigner: a process diagram editor for gene-regulatory and biochemical networks. *BIOSSILICO* 1, 5 (2003), 159–162, doi: 10.1016/S1478-5382(03)02370-9.  . A.13
- [146] FURMANOVA, K., JARESOVA, M., BYSKA, J., JURCIK, A., PARULEK, J., HAUSER, H., AND KOZLIKOVA, B. Interactive exploration of ligand transportation through protein tunnels. *BMC Bioinform* 18, 2 (2016), 1–16, doi: 10.1186/s12859-016-1448-0.  . A.5.2, A.13
- [147] FURMANOVÁ, K., JURČÍK, A., KOZLÁKOVÁ, B., HAUSER, H., AND BYŠKA, J. Multiscale visual drilldown for the analysis of large ensembles of multi-body protein complexes. *IEEE Trans Vis Comput Graph* 26, 1 (2019), 843–852, doi: 10.1109/TVCG.2019.2934333.  . A.5.2, A.13
- [148] GARRISON, L., MEUSCHKE, M., FAIRMAN, J., SMIT, N. N., PREIM, B., AND BRUCKNER, S. An exploration of practice and preferences for the visual communication of biomedical processes. In *Proc EG VCBM* (2021), pp. 1–12. 1.1, 3.3.2, A.5.3, A.8.1, E.3.1, E.3.2, E.6
- [149] GARRISON, L., VAŠÍČEK, J., CRAVEN, A. R., GRÜNER, R., SMIT, N., AND BRUCKNER, S. Interactive visual exploration of metabolite ratios in mr spectroscopy studies. *Comput Graph* 92 (2020), 1–12, doi: 10.1016/j.cag.2020.08.001.  . A.5.3, A.13
- [150] GARRISON, L., VAŠÍČEK, J., GRÜNER, R., SMIT, N., AND BRUCKNER, S. Spectramosaic: An exploratory tool for the interactive visual analysis of magnetic resonance spectroscopy data. In *Proc EG VCBM* (sep 2019). 2.2, B.1, B.2
- [151] GASTEIGER, R., NEUGEBAUER, M., BEUNG, O., AND PREIM, B. The flowlens: A focus-and-context visualization approach for exploration of blood flow in cerebral aneurysms. *IEEE Trans Vis Comput Graph* 17, 12 (2011), 2183–2192, doi: 10.1109/TVCG.2011.243. ■  . A.8.1, A.13
- [152] GASTEIGER, R., NEUGEBAUER, M., KUBISCH, C., AND PREIM, B. Adapted surface visualization of cerebral aneurysms with embedded blood flow information. In *Proc EG VCBM* (2010). ■  . A.8.1, A.13
- [153] GE, B., GUO, L., ZHANG, T., HU, X., HAN, J., AND LIU, T. Resting state fMRI-guided fiber clustering: Methods and applications. *Neuroinformatics* 11, 1 (2013), 119–133, doi: 10.1007/s12021-012-9169-7. ■
 . A.8.4, A.13
- [154] GEHLENBORG, N., O'DONOGHUE, S. I., BALIGA, N. S., GOESMANN, A., HIBBS, M. A., KITANO, H., KOHLBACHER, O., NEUEWEGER, H., SCHNEIDER, R., TENENBAUM, D., ET AL. Visualization of omics data for systems biology. *Nat Methods* 7, 3 (2010), S56–S68, doi: 10.1038/nmeth.1436. ■ . A.5, A.10, A.13
- [155] GERHARD, S., DADUCCI, A., LEMKADDEM, A., MEULI, R., THIRAN, J.-P., AND HAGMANN, P. The connectome viewer toolkit: An open source framework to manage, analyze, and visualize connectomes. *Front Neuroinform* 5, 3 (2011), 1–15, doi: 10.3389/fninf.2011.00003. ■  . A.13
- [156] GHAFARIZADEH, A., HEILAND, R., FRIEDMAN, S. H., MUMENTHALER, S. M., AND MACKLIN, P. PhysiCell: An open source physics-based cell simulator for 3-D multicellular systems. *PLoS Comput Biol* 14, 2 (2018), e1005991, doi: 10.1371/journal.pcbi.1005991.  . A.6.2, A.13










- [157] GIL, D., ARIS, R., BORRAS, A., RAMÍREZ, E., SEBASTIAN, R., AND VAZQUEZ, M. Influence of fiber connectivity in simulations of cardiac biomechanics. *Int J Comput Assist Radiol Surg* 14, 1 (2019), 63–72, doi: 10.1007/s11548-018-1849-9.   . A.9, A.13
- [158] GILJA, O. H., HAUSKEN, T., ØDEGAARD, S., AND BERSTAD, A. Three-dimensional ultrasonography of the gastric antrum in patients with functional dyspepsia. *Scand J Gastroenterol* 31, 9 (1996), 847–855, doi: 10.3109/00365529609051991.   . A.8.5, A.13
- [159] GILLMANN, C., SMIT, N. N., GRÖLLER, E., PREIM, B., VILANOVA, A., AND WISCHGOLL, T. Ten open challenges in medical visualization. *IEEE Comput Graph Appl* 41, 5 (2021), 7–15, doi: 10.1109/MCG.2021.3094858. 3.1, A.1
- [160] GLANCY, B. Visualizing mitochondrial form and function within the cell. *Trends Mol Med* 26, 1 (2020), 58–70, doi: 10.1016/j.molmed.2019.09.009.   . A.6, A.6.1, A.13
- [161] GLASSER, S., PREIM, U., TÖNNIES, K., AND PREIM, B. A visual analytics approach to diagnosis of breast dce-mri data. *Comput Graph* 34, 5 (2010), 602–611, doi: 10.1016/j.cag.2010.05.016.   . A.7.2, A.13
- [162] GLAUCHE, I., LORENZ, R., HASENCLEVER, D., AND ROEDER, I. A novel view on stem cell development: analysing the shape of cellular genealogies. *Cell Prolif* 42, 2 (2009), 248–263, doi: 10.1111/j.1365-2184.2009.00586.x.   . A.6.2, A.13
- [163] GLEICHER, M., ALBERS, D., WALKER, R., JUSUFI, I., HANSEN, C. D., AND ROBERTS, J. C. Visual comparison for information visualization. *Inf Vis* 10, 4 (2011), 289–309, doi: 10.1177/1473871611416549. 2.2, B.2
- [164] GOLDSTEIN, D. S. How does homeostasis happen? integrative physiological, systems biological, and evolutionary perspectives. *Am J Physiol Regul Integr Comp Physiol* 316, 4 (2019), R301–R317, doi: 10.1152/ajpregu.00396.2018. A.1
- [165] GOODSSELL, D. Molecular landscapes by David S. Goodsell. RCSB PDB-101, 2022. <https://pdb101.rcsb.org/sci-art/goodsell-gallery>. 2.3, E.2, E.3.1
- [166] GOODSSELL, D. S. Neuromuscular synapse. *Biochem Mol Biol Educ* 37, 4 (2009), 204–210, doi: 10.1002/bmb.20297.   . A.6.2, A.13
- [167] GOODSSELL, D. S., AND JENKINSON, J. Molecular illustration in research and education: past, present, and future. *J Mol Bio* 430, 21 (2018), 3969–3981, doi: 10.1016/j.jmb.2018.04.043. 2.3, 3.1
- [168] GOODSSELL, D. S., AND JOHNSON, G. T. Filling in the gaps: artistic license in education and outreach. *PLoS Comput Biol* 5, 12 (2007), e308. 2.3, 3.1, A.1, D.2, E.1, E.2
- [169] GOODSSELL, D. S., OLSON, A. J., AND FORLI, S. Art and science of the cellular mesoscale. *Trends Biochem Sci* 45, 6 (2020), 472–483, doi: 10.1016/j.tibs.2020.02.010. . 2.1, 2.3, A.1, A.6, A.13, E.1, E.2
- [170] GOODSTADT, M., AND MARTI-RENOM, M. A. Challenges for visualizing three-dimensional data in genomic browsers. *FEBS Lett*. 591, 17 (2017), 2505–2519, doi: 10.1002/1873-3468.12778. . 3.1.1, A.4.1, A.5, A.13
- [171] GRADL, S., WIRTH, M., ZILLIG, T., AND ESKOFIER, B. M. Visualization of heart activity in virtual reality: A biofeedback application using wearable sensors. In *Proc BSN* (2018), pp. 152–155.   . A.8.2, A.13
- [172] GRALKA, P., BECHER, M., BRAUN, M., FRIESS, F., MÜLLER, C., RAU, T., SCHATZ, K., SCHULZ, C., KRONE, M., REINA, G., ET AL. MegaMol—a comprehensive prototyping framework for visualizations. *Eur Phys J Spec Top* 227, 14 (2019), 1817–1829, doi: 10.1140/epjst/e2019-800167-5.   . A.5.1, A.13

- [173] GRAMAZIO, C. C., LAIDLAW, D. H., AND SCHLOSS, K. B. Colorlogical: creating discriminable and preferable color palettes for information visualization. *IEEE Trans Vis Comput Graph* 23, 1 (2017), 521–530, doi: 10.1109/TVCG.2016.2598918. 2.1, E.2
- [174] GREEN, S. B., AND SALKIND, N. J. *Using SPSS for Windows and Macintosh, books a la carte*. Pearson, 2016. C.3
- [175] GRESH, D. L., ROGOWITZ, B. E., WINSLOW, R. L., SCOLLAN, D. F., AND YUNG, C. K. Weave: A system for visually linking 3-d and statistical visualizations applied to cardiac simulation and measurement data. In *Proc IEEE VIS* (2000), pp. 489–492. . A.9, A.13
- [176] GRISS, J., VITERI, G., SIDIROPOULOS, K., NGUYEN, V., FABREGAT, A., AND HERMIAKOB, H. Reactomegsa-efficient multi-omics comparative pathway analysis. *Mol Cell Proteomics* 19, 12 (2020), 2115–2125, doi: 10.1074/mcp.TIR120.002155. . A.5, A.5.3, A.13
- [177] GÜNTHER, D., BOTO, R. A., CONTRERAS-GARCIA, J., PIQUEMAL, J.-P., AND TIERNY, J. Characterizing molecular interactions in chemical systems. *IEEE Trans Vis Comput Graph* 20, 14 (2014), 2476–2485, doi: 10.1109/TVCG.2014.2346403. . A.5.1, A.13
- [178] GUO, D., HAN, D., XU, X., YE, K., AND NIE, J. Spatiotemporal multiscale molecular cavity visualization and visual analysis. *J Vis (Tokyo)* 23 (2020), 661–676, doi: 10.1007/s12650-020-00646-x. . A.5.1, A.13
- [179] GUPTA, S., CZECH, J., KUCZEWSKI, R., BARTOL, T. M., SEJNOWSKI, T. J., LEE, R. E., AND FAEDER, J. R. Spatial stochastic modeling with mcell and cellblender. *arXiv* (2018). . A.5.2, A.13
- [180] HALL, J. E., AND GUYTON, A. C. *Guyton and Hall textbook of medical physiology*, 12th ed. ed. Saunders/Elsevier, 2011. 1.2.1, 3.1, 3.1.1, A.1, A.1, A.3, A.4.1, A.5.1, A.5.2, A.5.3, A.8.2, A.8.3, 2, E.7
- [181] HANDLEY, J. W., BRODLIE, K. W., AND CLAYTON, R. H. Multi-variate visualization of cardiac virtual tissue. In *Proc CBMS* (2006), pp. 665–672. . A.7.2, A.13
- [182] HANSEGÅRD, J., ORDERUD, F., AND RABBen, S. I. Real-time active shape models for segmentation of 3D cardiac ultrasound. In *Proc CAIP* (2007), pp. 157–164. . A.8.2, A.13
- [183] HARBIG, T. A., FRATTE, J., KRONE, M., AND NIESELT, K. K. OmicsTIDE: Interactive exploration of trends in multi-omics data. *bioRxiv* (2021), doi: 10.1101/2021.02.01.428836. . A.5.3, A.13
- [184] HARLOFF, A., ALBRECHT, F., SPREER, J., STALDER, A. F., BOCK, J., FRYDRYCHOWICZ, A., SCHÖLLHORN, J., HETZEL, A., SCHUMACHER, M., HENNIG, J., AND MARKL, M. 3D blood flow characteristics in the carotid artery bifurcation assessed by flow-sensitive 4D MRI at 3T. *Magn Reson Med* 61, 1 (2009), 65–74, doi: 10.1002/mrm.21774. . A.13
- [185] HARROWER, M., AND BREWER, C. A. Colorbrewer.org: an online tool for selecting colour schemes for maps. *Cartogr J* 40, 1 (2003), 27–37, doi: 10.1179/000870403235002042. 2.1, 3.2.1, B.6, E.2
- [186] HAUSER, H. Generalizing focus+ context visualization. In *Scientific visualization: The visual extraction of knowledge from data*. 2006, pp. 305–327. D.3
- [187] HAUSER, H., LEDERMANN, F., AND DOLEISCH, H. Angular brushing of extended parallel coordinates. In *Proc IEEE InfoVis* (2002), pp. 127–130. 14, C.2.2, C.2.2
- [188] HEIMDAL, A., STØYLEN, A., TORP, H., AND SKJÆRPE, T. Real-time strain rate imaging of the left ventricle by ultrasound. *J Am Soc Echocardiogr* 11, 11 (1998), 1013–1019, doi: 10.1016/S0894-7317(98)70151-8. . A.8.2, A.13
- [189] HEINRICH, J., AND WEISKOPF, D. State of the art of parallel coordinates. In *Proc EuroGraphics* (2013), pp. 96–116. 2.2, 14, C.1, C.2.2

- [190] HELLSTRÖM, J., ZAPATA, R. R., LIBARD, S., WIKSTRÖM, J., ORTIZ-NIETO, F., ALAFUZOFF, I., AND RAININKO, R. The value of magnetic resonance spectroscopy as a supplement to MRI of the brain in a clinical setting. *PLoS One* 13, 11 (2018), doi: 10.1371/journal.pone.0207336. 3.2.1, B.3
- [191] HEMPEL, S., GÖTZENBERGER, L., KÜHN, I., MICHALSKI, S. G., RILLIG, M. C., ZOBEL, M., AND MOORA, M. Mycorrhizas in the central european flora: relationships with plant life history traits and ecology. *Ecology* 94, 6 (2013), 1389–1399, doi: 10.1890/12-1700.1. C.9, C.3.1, C.3.1
- [192] HENNEMUTH, A., MAHNKEN, A., KÜHNEL, C., OELTZE, S., AND PEITGEN, H.-O. Ct late enhancement segmentation for the combined analysis of coronary arteries and myocardial viability. In *Proc EG VCBM* (2008), pp. 1–9. ■ ■ ■ ■ ■. A.13
- [193] HENZLER-WILDMAN, K., AND KERN, D. Dynamic personalities of proteins. *Nature* 450, 7172 (2007), 964–972, doi: 10.1038/nature06522. 1.2.1, A.2, A.5.1
- [194] HERMOSILLA, P., ESTRADA, J., GUALLAR, V., ROPINSKI, T., ÀLVAR VINACUA, AND VÁZQUEZ, P.-P. Physics-based visual characterization of molecular interaction forces. *IEEE Trans Vis Comput Graph* 23, 1 (2017), 731–740, doi: 10.1109/TVCG.2016.2598825. ■ ■ ■ ■ ■. A.5.2, A.13
- [195] HERMOSILLA, P., VÁZQUEZ, P., VINACUA, A., AND ROPINSKI, T. A general illumination model for molecular visualization. In *Comput Graph Forum* (2018), vol. 37, pp. 367–378. 2.1, 3.3.2, E.1, E.2
- [196] HESSE, M., RAULF, A., PILZ, G.-A., HABERLANDT, C., KLEIN, A. M., JABS, R., ZAEHRES, H., FÜGEMANN, C. J., ZIMMERMANN, K., TREBICKA, J., ET AL. Direct visualization of cell division using high-resolution imaging of M-phase of the cell cycle. *Nat Commun* 3, 1 (2012), 1–12, doi: 10.1038/ncomms2089. ■ ■ ■ ■ ■. A.6.1, A.13
- [197] HESTER, R., BROWN, A., HUSBAND, L., ILIESCU, R., PRUETT, W. A., SUMMERS, R. L., AND COLEMAN, T. HumMod: a modeling environment for the simulation of integrative human physiology. *Front Physiol* 2 (2011), 12, doi: 10.3389/fphys.2011.00012. ■ ■ ■ ■ ■. A.9, A.13
- [198] HINTZE, J. L., AND NELSON, R. D. Violin plots: a box plot-density trace synergism. *Am Stat* 52, 2 (1998), 181–184, doi: 10.1080/00031305.1998.10480559. B.6
- [199] HIRST, J. D., GLOWACKI, D. R., AND BAADEN, M. Molecular simulations and visualization: introduction and overview. *Faraday Discuss.* 169 (2014), 9–22, doi: 10.1039/C4FD90024C. ■ ■ ■ ■ ■. A.5, A.13
- [200] HOBSON, C. M., FALVO, M. R., AND SUPERFINE, R. A survey of physical methods for studying nuclear mechanics and mechanobiology. *APL Bioeng.* 5, 4 (2021), 041508, doi: 10.1063/5.0068126. ■ ■ ■ ■ ■. A.6, A.13
- [201] HOEHME, S., BRULPORT, M., BAUER, A., BEDAWY, E., SCHORMANN, W., HERMES, M., PUPPE, V., GEBHARDT, R., ZELLMER, S., SCHWARZ, M., BOCKAMP, E., TIMMEL, T., HENGSTLER, J. G., AND DRASDO, D. Prediction and validation of cell alignment along microvessels as order principle to restore tissue architecture in liver regeneration. *Proc Natl Acad Sci USA* 107, 23 (2010), 10371–10376, doi: 10.1073/pnas.0909374107. ■ ■ ■ ■ ■. A.7.1, A.13
- [202] HOEHME, S., AND DRASDO, D. A cell-based simulation software for multi-cellular systems. *Bioinformatics* 26, 20 (2010), 2641–2642, doi: 10.1093/bioinformatics/btq437. ■ ■ ■ ■ ■. A.13
- [203] HÖLLT, T., PEZZOTTI, N., VAN UNEN, V., KONING, F., EISEMANN, E., LELIEVELDT, B., AND VILANOVA, A. Cytosplore: Interactive immune cell phenotyping for large single-cell datasets. *Comput Graph Forum* 35, 3 (2016), 171–180, doi: 10.1111/cgf.12893. ■ ■ ■ ■ ■. A.6.1, A.13
- [204] HÖLLT, T., PEZZOTTI, N., VAN UNEN, V., KONING, F., LELIEVELDT, B. P., AND VILANOVA, A. Cytguide: Visual guidance for hierarchical single-cell analysis. *IEEE Trans Vis Comput Graph* 24, 1 (2017), 739–748, doi: 10.1109/TVCG.2017.2744318. ■ ■ ■ ■ ■. A.6.1, A.13
- [205] HOLTEN, D., AND VAN WIJK, J. J. Force-directed edge bundling for graph visualization. In *Comput Graph Forum* (2009), vol. 28, pp. 983–990. C.5

- [206] HONG, G.-R., PEDRIZZETTI, G., TONTI, G., LI, P., WEI, Z., KIM, J. K., BAWEJA, A., LIU, S., CHUNG, N., HOULE, H., NARULA, J., AND VANNAN, M. A. Characterization and quantification of vortex flow in the human left ventricle by contrast echocardiography using vector particle image velocimetry. *JACC Cardiovasc Imaging* 1, 6 (2008), 705–717, doi: 10.1016/j.jcmg.2008.06.008. ■ ■ ■ . A.8.1, A.13
- [207] HUANG, M. L., HUANG, T.-H., AND ZHANG, X. A novel virtual node approach for interactive visual analytics of big datasets in parallel coordinates. *Future Gener Comput Syst* 55 (2016), 510–523, doi: 10.1016/j.future.2015.02.003. 2.2, 14, C, C.1, C.2.2
- [208] HUBMAP CONSORTIUM. The human body at cellular resolution: the NIH human biomolecular atlas program. *Nature* 574, 7777 (2019), 187, doi: 10.1038/s41586-019-1629-x. 3.1, A.1
- [209] HUITEMA, H., AND VAN LIERE, R. Interactive visualization of protein dynamics. In *Proc IEEE VIS* (2000), pp. 465–468. ■ ■ ■ . A.5.1, A.13
- [210] HUMPHREY, W., DALKE, A., AND SCHULTEN, K. VMD: Visual molecular dynamics. *J Mol Graph Model* 14, 1 (1996), 33–38, doi: 10.1016/0263-7855(96)00018-5. ■ ■ ■ . A.5.1, A.10, A.13
- [211] HUNTER, P. J. Modeling human physiology: The iups/embs physiome project. *Proc IEEE* 94, 4 (2006), 678–691, doi: 10.1109/jProc2006.871767. 1.1, 3.1, A.1
- [212] IGLESIAS-GUITIAN, J. A., ALIAGA, C., JARABO, A., AND GUTIERREZ, D. A biophysically-based model of the optical properties of skin aging. In *Comput Graph Forum* (2015), vol. 34, pp. 45–55. ■ ■ ■ . A.7.1, A.13
- [213] IHAKA, R. Colour for presentation graphics. In *Proc DSC* (2003), p. 2. E.5
- [214] ING-SIMMONS, E., AND VAQUERIZAS, J. M. Visualising three-dimensional genome organisation in two dimensions. *Development* 146, 19 (2019), dev177162, doi: 10.1242/dev.177162. ■ ■ ■ . A.5, A.13
- [215] INGRAM, S., MUNZNER, T., IRVINE, V., TORY, M., BERGNER, S., AND MÖLLER, T. Dimstiller: Workflows for dimensional analysis and reduction. In *Proc IEEE VAST* (2010), pp. 3–10. 2.2, C.1
- [216] INSLEY, J. A., GRINBERG, L., AND PAPKA, M. E. Visualizing multiscale, multiphysics simulation data: Brain blood flow. In *Proc LDAV* (2011), pp. 3–7. ■ ■ ■ . A.9, A.13
- [217] INSTITUTE, N. N. H. G. R. Biological pathways fact sheet, 2020. <https://www.genome.gov/about-genomics/fact-sheets/Biological-Pathways-Fact-Sheet>. A.5.3
- [218] IONASEC, R. I., GRBIC, S., VITANOVSKI, D., VOIGT, I., GEORGESCU, B., VEGA-HIGUERA, F., AND COMANICIU, D. Complete valvular heart apparatus model from 4D cardiac CT. In *Proc MICCAI* (2010). ■ ■ ■ . A.8.2, A.13
- [219] ISENBERG, P., ISENBERG, T., SEDLMAYER, M., CHEN, J., AND MÖLLER, T. Visualization as seen through its research paper keywords. *IEEE Trans Vis Comput Graph* 23, 1 (2016), 771–80. 2.3, D.2
- [220] ISENBERG, T. Visual abstraction and stylisation of maps. *Cartogr J* 50, 1 (2013), 8–18. D.4
- [221] ISENBERG, T., NEUMANN, P., CARPENDALE, S., SOUSA, M. C., AND JORGE, J. A. Non-photorealistic rendering in context: an observational study. In *Proc NPAR* (2006), pp. 115–26. 2.3, D.2, D.4
- [222] ITTEN, J. *The art of color*. Reinhold Pub Corp, 1974. 1.2.3, 2.3, 3.3.2, E.1, E.2, E.3.1, E.4
- [223] IVANOV, P. C. The new field of network physiology: Building the human physiome. *Front Physiol* (2021), 1, doi: 10.3389/fnnetp.2021.711778. A.11
- [224] IWASA, J. Visualizing the origins of life: Molecular animation for scientific research and education. In *ACM SIGGRAPH 2007: Educators Program*. 2007. ■ ■ ■ . A.13
- [225] JAINEK, W. M., BORN, S., BARTZ, D., STRASSER, W., AND FISCHER, J. Illustrative hybrid visualization and exploration of anatomical and functional brain data. In *Comput Graph Forum* (2008), vol. 27, pp. 855–862. ■ ■ ■ . A.8.4, A.13

- [226] JANETZKO, H., STEIN, M., SACHA, D., AND SCHRECK, T. Enhancing parallel coordinates: Statistical visualizations for analyzing soccer data. In *Proc VDA* (2016). 14, C.2.2
- [227] JANOOS, F., NOUANESENGSY, B., MACHIRAJU, R., SHEN, H. W., SAMMET, S., KNOPP, M., AND MÓROCZ, I. A. Visual analysis of brain activity from fMRI data. *Comput Graph Forum* 28, 3 (2009), 903–910, doi: 10.1111/j.1467-8659.2009.01458.x. ■ ■ ■ . A.8.4, A.13
- [228] JAWAD, M., EVERS, M., GERWING, A., HERICK, M., SEIBERT, D., BAUER, J., KUGEL, H., OHRMANN, P., AND LINSEN, L. A visual analytics approach for comparing cohorts in single-voxel magnetic resonance spectroscopy data. In *Biomedical Visualisation*. 2019, pp. 115–136. □ ■ ■ . 2.2, A.5.3, A.13, B.2
- [229] JAWAD, M., MOLCHANOV, V., AND LINSEN, L. Coordinated image-and feature-space visualization for interactive magnetic resonance spectroscopy imaging data analysis. In *Proc VISGRAPP* (2019), pp. 118–128. □ ■ ■ . 2.2, A.5.3, A.13, B.2
- [230] JAWAD, M., MOLCHANOV, V., AND LINSEN, L. A reproducibility study for visual mrsi data analytics. In *Proc VISGRAPP* (2019), vol. 1182, pp. 362–388. 2.2
- [231] JENKINSON, J. The role of craft-based knowledge in the design of dynamic visualizations. In *Learning from Dynamic Visualization: Innovations in Research and Application*. 2017, pp. 93–117. 2.3, 3.3.2, D.2, D.4, E.2, E.4
- [232] JENKINSON, J., AND MCGILL, G. Visualizing protein interactions and dynamics: evolving a visual language for molecular animation. *CBE Life Sci Educ* 11, 1 (2012), 103–110, doi: 10.1187/cbe.11-08-0071. □ ■ ■ . 2.3, 3.3.1, A.5.2, A.13, D.2, D.6, E.2, E.6
- [233] JENNINGS, D., RAGHUNAND, N., AND GILLIES, R. J. Imaging hemodynamics. *Cancer Metastasis Rev* 27, 4 (2008), 589–613, doi: 10.1007/s10555-008-9157-4. A.8
- [234] JENSEN, E. C. Overview of live-cell imaging: requirements and methods used. *Anat Rec* 296, 1 (2013), 1–8, doi: 10.1002/ar.22554. A.6
- [235] JEONG, D. H., ZIEMKIEWICZ, C., FISHER, B., RIBARSKY, W., AND CHANG, R. iPCA: An interactive system for PCA-based visual analytics. *Comput Graph Forum* 28, 3 (2009), 767–774, doi: 10.1111/j.1467-8659.2009.01475.x. 2.2, 14, C.1, C.2.2
- [236] JIANG, Y.-F., LIN, S.-S., CHEN, J.-M., TSAI, H.-Z., HSIEH, T.-S., AND FU, C.-Y. Electron tomographic analysis reveals ultrastructural features of mitochondrial cristae architecture which reflect energetic state and aging. *Sci Rep* 7, 1 (2017), 1–11, doi: 10.1038/srep45474. ■ ■ ■ . A.6.1, A.13
- [237] JIANU, R., KEBING, Y., LULU, C., VINH, N., SALOMON, A. R., AND LAIDLAW, D. H. Visual integration of quantitative proteomic data, pathways, and protein interactions. *IEEE Trans Vis Comput Graph* 16, 4 (2010), 609–620, doi: 10.1109/TVCG.2009.106. □ ■ ■ . A.5.3, A.13
- [238] JIQUAN, L., JINGYI, F., DUAN, H., AND SIPING, C. Simulation data mapping in virtual cardiac model. In *Proc Eng Med Biol Soc* (2004), pp. 1937–1940. ■ ■ ■ . A.8.2, A.13
- [239] JOHNSON, G. T., AND HERTIG, S. A guide to the visual analysis and communication of biomolecular structural data. *Nat Rev Mol Cell Biol* 15, 10 (2014), 690–698, doi: 10.1038/nrm3874. ■ . 2.3, 3.1, 3.3.2, A.1, A.5, A.13, D.2, E.1, E.2, E.4
- [240] JOHNSON, G. T., NOSKE, A., AND MARSH, B. Rapid visual inventory & comparison of complex 3D structures. NSF Scientific Visualization Challenge 2011, 2011. ■ ■ ■ , <http://www.wired.com/2012/02/science-visualizations-2011/>. A.13
- [241] JÖNSSON, D., BERGSTRÖM, A., FORSELL, C., SIMON, R., ENGSTRÖM, M., YNNERMAN, A., AND HOTZ, I. A visual environment for hypothesis formation and reasoning in studies with fMRI and multi-variate clinical data. In *Proc EG VCBM* (2019). ■ ■ ■ . A.8.4, A.13

- [242] JURČÍK, A., FURMANOVÁ, K., BYŠKA, J., VONÁSEK, V., VÁVRA, O., ULBRICH, P., HAUSER, H., AND KOZLÍKOVÁ, B. Visual analysis of ligand trajectories in molecular dynamics. In *Proc IEEE PacificVis* (2019), pp. 212–221. . A.5.2, A.13
- [243] KANEHISA, M., AND SATO, Y. KEGG mapper for inferring cellular functions from protein sequences. *Protein Sci.* 29, 1 (2020), 28–35, doi: 10.1002/pro.3711. . A.5.3, A.13
- [244] KAUR, G., AND KARKI, B. B. Bifocal parallel coordinates plot for multivariate data visualization. In *VISGRAPP (3: IVAPP)* (2018), pp. 176–183. C.2.2
- [245] KEHRER, J., AND HAUSER, H. Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Trans Vis Comput Graph* 19, 3 (2012), 495–513, doi: 10.1109/TVCG.2012.110. 2.1, A.1
- [246] KERPPOLA, T. K. Visualization of molecular interactions by fluorescence complementation. *Nat Rev Mol Cell Biol* 7, 6 (2006), 449–456, doi: 10.1038/nrm1929. . A.5.2, A.13
- [247] KERR, R. A., BARTOL, T. M., KAMINSKY, B., DITTRICH, M., CHANG, J.-C. J., BADEN, S. B., SEJNOWSKI, T. J., AND STILES, J. R. Fast monte carlo simulation methods for biological reaction-diffusion systems in solution and on surfaces. *SIAM J Sci Comput* 30, 6 (2008), 3126–49. D.1, D.4.2
- [248] KERREN, A., KUCHER, K., LI, Y.-F., AND SCHREIBER, F. BioVis Explorer: A visual guide for biological data visualization techniques. *PLoS One* 12, 11 (2017), e0187341, doi: 10.1371/journal.pone.0187341. . A.5, A.13
- [249] KERZNER, E., GOODWIN, S., DYKES, J., JONES, S., AND MEYER, M. A framework for creative visualization-opportunities workshops. *IEEE Trans Vis Comput Graph* 25, 1 (2019), 748–758, doi: 10.1109/TVCG.2018.2865241. 3.3.1, D.4.2, E.6
- [250] KIKINIS, R., PIEPER, S. D., AND VOSBURGH, K. G. 3D slicer: a platform for subject-specific image analysis, visualization, and clinical support. In *Intraoperative imaging and image-guided therapy*. 2014, pp. 277–89. 3.3.1, D.4.2
- [251] KIM, M., BORDAS, R., VOS, W., HARTLEY, R. A., BRIGHTLING, C. E., KAY, D., GRAU, V., AND BURROWES, K. S. Dynamic flow characteristics in normal and asthmatic lungs. *Int J Numer Method Biomed Eng* 31, 12 (2015), doi: 10.1002/cnm.2730. . A.4, A.18, A.8.3, A.9, A.13
- [252] KINDLMANN, G. Superquadric tensor glyphs. In *Proc IEEE VIS* (2004), pp. 147–154. A.8.4
- [253] KINDLMANN, G., AND SCHEIDEGGER, C. An algebraic process for visualization design. *IEEE Trans Vis Comput Graph* 20, 12 (2014), 2181–2190. C.2.2
- [254] KING, Z. A., DRÄGER, A., EBRAHIM, A., SONNENSCHN, N., LEWIS, N. E., AND PALSSON, B. O. Escher: A web application for building, sharing, and embedding data-rich visualizations of biological pathways. *PLoS Comput Biol* 11, 8 (2015), 1–13, doi: 10.1371/journal.pcbi.1004321. . A.5.3, A.13
- [255] KINGMAN, P. Our resilient genome. IFF St. Tropez 2015, 2014. . A.13
- [256] KLEIN, T., VIOLA, I., GRÖLLER, E., AND MINDEK, P. Multi-scale procedural animations of microtubule dynamics based on measured data. *IEEE Trans Vis Comput Graph* 26, 1 (2019), 622–632, doi: 10.1109/TVCG.2019.2934612. . 2.1, 3.3.2, A.6.1, A.13, E.1, E.2
- [257] KLEMM, P., OELTZE-JAFRA, S., LAWONN, K., HEGENSCHN, K., VÖLZKE, H., AND PREIM, B. Interactive visual analysis of image-centric cohort study data. *IEEE Trans Vis Comput Graph* 20, 12 (2014), 1673–1682, doi: 10.1109/tvcg.2014.2346591. 2.2, B.2
- [258] KLINDER, T., LORENZ, C., AND OSTERMANN, J. Prediction framework for statistical respiratory motion modeling. In *Proc MICCAI* (2010). . A.8.3, A.13
- [259] KOHL, P., AND NOBLE, D. Systems biology and the virtual physiological human. *Mol Syst Biol* 5, 1 (2009), 292, doi: 10.1038/msb.2009.51. A.1


























- [260] KÖHLER, B., BORN, S., VAN PELT, R. F., HENNEMUTH, A., PREIM, U., AND PREIM, B. A survey of cardiac 4D PC-MRI data processing. In *Comput Graph Forum* (2017), vol. 36, pp. 5–35. ■ . 2.1, A.1, A.8, A.8.1, A.13
- [261] KÖHLER, B., BORN, S., VAN PELT, R. F. P., PREIM, U., AND PREIM, B. A survey of cardiac 4D PC-MRI data processing. In *Proc EG VCBM* (2015), pp. 1–30. ■ . A.13
- [262] KÖHLER, B., GASTEIGER, R., PREIM, U., THEISEL, H., GUTBERLET, M., , AND PREIM, B. Semi-automatic vortex extraction in 4D PC-MRI cardiac blood flow data using line predicates. *IEEE Trans Vis Comput Graph* 19, 12 (2013), 2773–2782, doi: 10.1109/TVCG.2013.189. ■ ■ ■ ■ . A.13
- [263] KOHOUT, J., AND ČERVENKA, M. Muscle deformation using position based dynamics. In *Proc BIOSTEC* (2020), pp. 486–509. ■ ■ ■ ■ . A.8.5, A.13
- [264] KOHOUT, J., AND KUKAČKA, M. Real-time modelling of fibrous muscle. In *Comput Graph Forum* (2014), vol. 33, pp. 1–15. ■ ■ ■ ■ . A.8.5, A.13
- [265] KOLESÁR, I., BYŠKA, J., PARULEK, J., HAUSER, H., AND KOZLÍKOVÁ, B. Unfolding and interactive exploration of protein tunnels and their dynamics. In *Proc EG VCBM* (2016), pp. 1–10. ■ ■ ■ ■ . A.5.1, A.13
- [266] KOLESÁR, I., PARULEK, J., VIOLA, I., BRUCKNER, S., STAVRUM, A.-K., AND HAUSER, H. Interactively illustrating polymerization using three-level model fusion. *BMC Bioinform* 15, 345 (2014), 1–16, doi: 10.1186/1471-2105-15-345. ■ ■ ■ ■ . A.5.2, A.13
- [267] KOZLÍKOVÁ, B., KRONE, M., FALK, M., LINDOW, N., BAADEN, M., BAUM, D., VIOLA, I., PARULEK, J., AND HEGE, H.-C. Visualization of biomolecular structures: State of the art revisited. In *Comput Graph Forum* (2017), vol. 36, pp. 178–204. ■ . 2.1, A.1, A.5, A.13, D.4.2
- [268] KOZLÍKOVÁ, B., KRONE, M., LINDOW, N., FALK, M., BAADEN, M., BAUM, D., VIOLA, I., PARULEK, J., AND HEGE, H.-C. Visualization of biomolecular structures: State of the art. In *Proc EuroVis* (2015), pp. 61–81. ■ . A.13
- [269] KRIEGEL, H., KRÖGER, P., AND ZIMEK, A. Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering. *ACM Trans Knowl Discov Data* 3, 1 (2009), 1–58, doi: 10.1145/1497577.1497578. C
- [270] KRISHNAN, H., GARTH, C., GÜHRING, J., GÜLSÜN, M. A., GREISER, A., AND JOY, K. I. Analysis of time-dependent flow-sensitive pc-mri data. *IEEE Trans Vis Comput Graph* 18, 6 (2012), 966–977, doi: 10.1109/TVCG.2011.80. ■ ■ ■ ■ . A.13
- [271] KRONE, M., FRIESS, F., SCHARNOWSKI, K., REINA, G., FADEMRECHT, S., KULSCHEWSKI, T., PLEISS, J., AND ERTL, T. Molecular surface maps. *IEEE Trans Vis Comput Graph* 23, 1 (2016), 701–710, doi: 10.1109/TVCG.2016.2598824. ■ ■ ■ ■ . A.5.1, A.13
- [272] KRONE, M., KOZLÍKOVÁ, B., LINDOW, N., BAADEN, M., BAUM, D., PARULEK, J., HEGE, H.-C., AND VIOLA, I. Visual analysis of biomolecular cavities: State of the art. In *Comput Graph Forum* (2016), vol. 35, pp. 527–551. ■ . A.5, A.13
- [273] KUANG, X., ZHANG, H., ZHAO, S., AND MCGUFFIN, M. Tracing tuples across dimensions: A comparison of scatterplots and parallel coordinate plots. *Comput Graph Forum* 31, 3pt4 (2012), 1365–1374, doi: 10.1111/j.1467-8659.2012.03129.x. C.2.2
- [274] KÜHN, I., DURKA, W., AND KLOTZ, S. BioFlor: a new plant-trait database as a tool for plant invasion ecology. *Diversity and Distributions* 10, 5/6 (2004), 363–365, doi: 10.1111/j.1366-9516.2004.00106.x. C.3.1
- [275] KULAHCIOGLU, T., AND DE MELO, G. Paralinguistic recommendations for affective word clouds. In *Proc IUI* (2019), pp. 132–143. 2.3, E.2

- [276] KULP, S., METAXAS, D., QIAN, Z., VOROS, S., AXEL, L., AND MIHALEF, V. Patient-specific modeling and visualization of blood flow through the heart. In *Proc ISBI* (2011), pp. 1692–1697. ■ ■ ■ ■ . A.8.2, A.13
- [277] LAIDLAW, D. H., AHRENS, E. T., KREMERS, D., AVALOS, M. J., JACOBS, R. E., AND READHEAD, C. Visualizing diffusion tensor images of the mouse spinal cord. In *Proc IEEE VIS* (1998), pp. 127–134. A.8.4
- [278] LAM, H., BERTINI, E., ISENBERG, P., PLAISANT, C., AND CARPENDALE, S. Empirical studies in information visualization: Seven scenarios. *IEEE Trans Vis Comput Graph* 18, 9 (2012), 1520–36. 3.3.1, D.4.1
- [279] LAMPE, O. D., VIOLA, I., REUTER, N., AND HAUSER, H. Two-level approach to efficient visualization of protein dynamics. *IEEE Trans Vis Comput Graph* 13, 6 (2007), 1616–1623, doi: 10.1109/TVCG.2007.70517. □ ■ ■ ■ . A.5.1, A.13
- [280] LANCIANO, T., BONCHI, F., AND GIONIS, A. Explainable classification of brain networks via contrast subgraphs. In *Proc ACM KDD* (2020), pp. 3308–3318. ■ ■ ■ ■ . A.19, A.8.4, A.13
- [281] LANGE, D., POLANCO, E., JUDSON-TORRES, R., ZANGLE, T., AND LEX, A. Loon: Using exemplars to visualize large-scale microscopy data. *IEEE Trans Vis Comput Graph* 28, 01 (2022), 248–258, doi: 10.1109/TVCG.2021.3114766. ■ ■ ■ ■ . A.6.1, A.13
- [282] LANTZ, J., GUPTA, V., HENRIKSSON, L., KARLSSON, M., PERSSON, A., CARLHÄLL, C.-J., AND EBBERS, T. Impact of pulmonary venous inflow on cardiac flow simulations: comparison with in vivo 4D flow MRI. *Ann Biomed Eng* 47, 2 (2019), 413–424, doi: 10.1007/s10439-018-02153-5. ■ ■ ■ ■ . A.8.1, A.13
- [283] LASSERRE, S., HERNANDO, J., HILL, S., SCHÜRMAN, F., DE MIGUEL ANASAGASTI, P., JAOUDE, G. A., AND MARKRAM, H. A neuron membrane mesh representation for visualization of electrophysiological simulations. *IEEE Trans Vis Comput Graph* 18, 2 (2012), 214–227, doi: 10.1109/TVCG.2011.55. ■ ■ ■ ■ . A.7.2, A.13
- [284] LAWONN, K., GASTEIGER, R., AND PREIM, B. Adaptive surface visualization of vessels with animated blood flow. *Comput Graph Forum* 33, 8 (2014), 16–27, doi: 10.1111/cgf.12355. ■ ■ ■ ■ . A.13
- [285] LAWONN, K., GLASSER, S., VILANOVA, A., PREIM, B., AND ISENBERG, T. Occlusion-free blood flow animation with wall thickness visualization. *IEEE Trans Vis Comput Graph* 22, 1 (2016), 728–737, doi: 10.1109/TVCG.2015.2467961. ■ ■ ■ ■ . A.8.1, A.13
- [286] LAWONN, K., GÜNTHER, T., AND PREIM, B. Coherent view-dependent streamlines for understanding blood flow. In *Proc EuroVis* (2014), pp. 19–23. ■ ■ ■ ■ . A.8.1, A.13
- [287] LAWONN, K., SMIT, N. N., BÜHLER, K., AND PREIM, B. A survey on multimodal medical data visualization. In *Comput Graph Forum* (2018), vol. 37, pp. 413–438. ■ ■ . 2.1, A.1, A.8, A.13, D.4.2
- [288] LAWONN, K., VIOLA, I., PREIM, B., AND ISENBERG, T. A survey of surface-based illustrative rendering for visualization. *Comput Graph Forum* 37, 6 (2018), 205–34. D.4, D.7
- [289] LÊ, S., JOSSE, J., HUSSON, F., ET AL. Factominer: an R package for multivariate analysis. *J Stat Softw* 25, 1 (2008), 1–18, doi: 10.18637/jss.v025.i01. 14, C, C.3
- [290] LE NOVERE, N., HUCKA, M., MI, H., MOODIE, S., SCHREIBER, F., SOROKIN, A., DEMIR, E., WEGNER, K., ALADJEM, M. I., WIMALARATNE, S. M., ET AL. The systems biology graphical notation. *Nat Biotechnol* 27, 8 (2009), 735–741, doi: 10.1038/nbt.1558. □ ■ ■ ■ . A.5.3, A.13
- [291] LEE, D., GLUECK, M., KHAN, A., FIUME, E., AND JACKSON, K. Modeling and simulation of skeletal muscle for computer graphics: A survey. *Found Trends Comput Graph Vis* 7, 4 (2012), 229–276, doi: 10.1561/06000000036. ■ ■ . A.8.5, A.13
- [292] LEE, R., KARR, J. R., AND COVERT, M. W. Wholecellviz: Data visualization for whole-cell models. *BMC Bioinform* 14, 1 (2013), 1–9, doi: 10.1186/1471-2105-14-253. ■ ■ ■ ■ . A.11, A.6.1, A.13

- [293] LEE, T.-Y., CHAUDHURI, A., PORIKLI, F., AND SHEN, H.-W. Cyclestack: Inferring periodic behavior via temporal sequence visualization in ultrasound video. In *Proc IEEE PacificVis* (2010), pp. 89–96. ■ ■ ■ . A.8.3, A.13
- [294] LEGGIO, B., LAUSSU, J., CARLIER, A., GODIN, C., LEMAIRE, P., AND FAURE, E. Morphonet: an interactive online morphological browser to explore complex multi-scale data. *Nat Commun* 10, 1 (2019), 1–8, doi: 10.1038/s41467-019-10668-1. ■ ■ ■ . A.9, A.20, A.13
- [295] LEI, J., AKHTER, N., QIAO, W., AND SHEHU, A. Reconstruction and decomposition of high-dimensional landscapes via unsupervised learning. In *Proc ACM KDD* (2020), pp. 2505–2513. □ ■ ■ . A.5.1, A.13
- [296] LEISTIKOW, S., NAHARDANI, A., HOERR, V., AND LINSEN, L. Interactive visual similarity analysis of measured and simulated multi-field tubular flow ensembles. In *Proc EG VCBM* (2020), pp. 139–150. ■ ■ ■ . A.8.1, A.13
- [297] LETUNIC, I., YAMADA, T., KANEHISA, M., AND BORK, P. ipath: interactive exploration of biochemical pathways and networks. *Trends Biochem Sci* 33, 3 (2008), 101–103, doi: 10.1016/j.tibs.2008.01.001. □ ■ ■ . A.5.3, A.13
- [298] LEWIS, C. *Using the "thinking-aloud" method in cognitive interface design*. IBM TJ Watson Research Center, 1982. B.8
- [299] LEX, A., PARTL, C., KALKOFEN, D., STREIT, M., GRATZL, S., WASSERMANN, A. M., SCHMALSTIEG, D., AND PFISTER, H. Entourage: Visualizing relationships between biological pathways using contextual subsets. *IEEE Trans Vis Comput Graph* 19, 12 (2013), 2536–2545, doi: 10.1109/TVCG.2013.154. □ ■ ■ . A.5.3, A.13
- [300] LEX, A., STREIT, M., KRUIJFF, E., AND SCHMALSTIEG, D. Caleydo: Design and evaluation of a visual analysis framework for gene expression data in its biological context. In *Proc IEEE PacificVis* (2010), pp. 57–64. □ ■ ■ . A.5.3, A.13
- [301] LI, R., AND CHEN, J. Toward a deep understanding of what makes a scientific visualization memorable. In *Proc IEEE SciVis* (2018), pp. 26–31. 2.3, D.2
- [302] LIN, D. W., JOHNSON, S., AND HUNT, C. A. Modeling liver physiology: combining fractals, imaging and animation. In *Proc IEMBS* (2004), pp. 3120–3123. ■ ■ ■ . A.8.5, A.13
- [303] LIN, S., FORTUNA, J., KULKARNI, C., STONE, M., AND HEER, J. Selecting semantically-resonant colors for data visualization. In *Comput Graph Forum* (2013), vol. 32, pp. 401–410. 2.3, E.2
- [304] LINDOW, N., BAUM, D., PROHASKA, S., AND HEGE, H.-C. Accelerated visualization of dynamic molecular surfaces. In *Comput Graph Forum* (2010), vol. 29, pp. 943–952. □ ■ ■ . A.5.1, A.13
- [305] LIPŠA, D. R., LARAMEE, R. S., COX, S. J., ROBERTS, J. C., WALKER, R., BORKIN, M. A., AND PFISTER, H. Visualization for the physical sciences. In *Comput Graph Forum* (2012), vol. 31, pp. 2317–2347. A.1
- [306] LIU, J.-Y., LV, W.-J., JIAN, J.-B., XIN, X.-H., ZHAO, X.-Y., AND HU, C.-H. High-resolution three-dimensional visualization of hepatic sinusoids in cirrhotic rats via serial histological sections. *Histol Histopathol* 36, 5 (2021), 577–586, doi: 10.14670/HH-18-339. ■ ■ ■ . A.7.1, A.13
- [307] LIU, S., MALJOVEC, D., WANG, B., BREMER, P., AND PASCUCCI, V. Visualizing high-dimensional data: Advances in the past decade. *IEEE Trans Vis Comput Graph* 23, 3 (2017), 1249–1268, doi: 10.1109/TVCG.2016.2640960. 2.2, C.1
- [308] LIU, Z., AND STASKO, J. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE Trans Vis Comput Graph* 16, 6 (2010), 999–1008, doi: 10.1109/TVCG.2010.177. C.2.2
- [309] LLOYD, C. M., LAWSON, J. R., HUNTER, P. J., AND NIELSEN, P. F. The cellml model repository. *Bioinformatics* 24, 18 (2008), 2122–2123, doi: 10.1093/bioinformatics/btn390. A.6

- [310] LUNDERVOLD, A. On consciousness, resting state fMRI, and neurodynamics. *Nonlinear Biomed Phys* 4, 1 (2010), 1–18, doi: 10.1186/1753-4631-4-S1-S9. ■ ■ ■ . A.8.4, A.13
- [311] L'YI, S., WANG, Q., LEKSCHAS, F., AND GEHLENBORG, N. Gosling: A grammar-based toolkit for scalable and interactive genomics data visualization. *IEEE Trans Vis Comput Graph* (2021), 1–1, doi: 10.1109/TVCG.2021.3114876. □ ■ ■ . A.9, A.5.3, A.13
- [312] MAC GABHANN, F., JI, J. W., AND POPEL, A. S. Vegf gradients, receptor activation, and sprout guidance in resting and exercising skeletal muscle. *J Appl Physiol* 102, 2 (2007), 722–734, doi: 10.1152/jap-physiol.00800.2006. ■ ■ ■ . A.9, A.13
- [313] MACHIRAJU, G. B., MALLICK, P., AND FRIEBOES, H. B. Multicompartment modeling of protein shedding kinetics during vascularized tumor growth. *Sci Rep* 10, 1 (2020), 1–16, doi: 10.1038/s41598-020-73866-8. ■ □ ■ . A.13, A.7.1, A.13
- [314] MADEHKHAKSAR, F., LUO, Z., PRONOST, N., AND EGGES, A. Modeling and simulating virtual anatomical humans. In *3D Multiscale Physiological Human*. 2014, pp. 137–164. ■ ■ ■ . A.9, A.13
- [315] MADHU, B., JAUHAINEN, A., MCGUIRE, S., AND GRIFFITHS, J. R. Exploration of human brain tumour metabolism using pairwise metabolite-metabolite correlation analysis (MMCA) of HR-MAS 1H NMR spectra. *PLoS One* 12, 10 (2017), e0185980. B.4.1
- [316] MAHNKE, F. H. *Color, environment, and human response: an interdisciplinary understanding of color and its use as a beneficial element in the design of the architectural environment*. John Wiley & Sons, 1996. 2.3, E.2
- [317] MANSI, T., VOIGT, I., MENGUE, E. A., IONASEC, R., GEORGESCU, B., NOACK, T., SEEBURGER, J., AND COMANICIU, D. Towards patient-specific finite-element simulation of mitralclip procedure. In *Proc MICCAI* (2011), pp. 452–459. ■ ■ ■ . A.8, A.8.2, A.13
- [318] MARGULIES, D. S., BÖTTGER, J., WATANABE, A., AND GORGOLEWSKI, K. J. Visualizing the human connectome. *NeuroImage* 80 (2013), 445–461, doi: 10.1016/j.neuroimage.2013.04.111. ■ . A.8.4, A.13
- [319] MARIAS, K., NIKIFORAKI, K., MANIKIS, G. C., KONTOPODIS, E., AND PAPANIKOLAOU, N. Visualizing tumor environment with perfusion and diffusion mri: Computational challenges. In *Proc CGI*. 2016, pp. 113–116. ■ ■ ■ . A.7.2, A.13
- [320] MARIÉS, A., MAYS, N., HUNT, M., WONG, K. F., LAYTON, W., BOUDREAU, R., ROSANO, C., AND MARAI, G. E. Grace: A visual comparison framework for integrated spatial and non-spatial geriatric data. *IEEE Trans Vis Comput Graph* 19, 12 (2013), 2916–2925, doi: 10.1109/TVCG.2013.161. ■ ■ ■ . A.8.4, A.13
- [321] MARINO, J., AND KAUFMAN, A. Prostate cancer visualization from MR imagery and MR spectroscopy. *Comput Graph Forum* 30, 3 (2011), 1051–1060, doi: 10.1111/j.1467-8659.2011.01954.x. 2.2, B.2
- [322] MARITAN, M., AUTIN, L., KARR, J., COVERT, M. W., OLSON, A. J., AND GOODSSELL, D. S. Building structural models of a whole mycoplasma cell. *J Mol Biol* 434, 2 (2022), 167351, doi: 10.1016/j.jmb.2021.167351. ■ ■ ■ . A.6.1, A.13
- [323] MARKL, M., FRYDRYCHOWICZ, A., KOZERKE, S., HOPE, M., AND WIEBEN, O. 4D flow MRI. *J Magn Reson Imaging* 36, 5 (2012), 1015–1036, doi: 10.1002/jmri.23632. ■ . A.8, A.13
- [324] MARKL, M., KILNER, P. J., AND EBBERS, T. Comprehensive 4d velocity mapping of the heart and great vessels by cardiovascular magnetic resonance. *J Cardiovasc Magn Reson* 13, 1 (2011), 1–22, doi: 10.1186/1532-429X-13-7. ■ . A.8.1, A.13
- [325] MARKLUND, E. G., AND BENESCH, J. L. Weighing-up protein dynamics: the combination of native mass spectrometry and molecular dynamics simulations. *Curr Opin Struct Biol* 54 (2019), 50–58, doi: 10.1016/j.sbi.2018.12.011. □ ■ ■ . A.5, A.5.2, A.13









- [326] MARTINEZ, X., KRONE, M., ALHARBI, N., ROSE, A. S., LARAMEE, R. S., O'DONOGHUE, S., BAADEN, M., AND CHAVENT, M. Molecular graphics: Bridging structural biologists and computer scientists. *Structure* 27, 11 (2019), 1617–1623, doi: 10.1016/j.str.2019.09.001. ■ . A.5, A.13
- [327] MASULLO, L. A., STEINER, F., ZÄHRINGER, J., LOPEZ, L. F., BOHLEN, J., RICHTER, L., COLE, F., TINNEFELD, P., AND STEFANI, F. D. Pulsed interleaved miniflux. *Nano Letters* 21, 1 (2020), 840–846, doi: 10.1021/acs.nanolett.0c04600. A.11
- [328] MATKOVIC, K., FREILER, W., GRACANIN, D., AND HAUSER, H. ComVis: A coordinated multiple views system for prototyping new visualization technology. In *Proc IV* (2008), pp. 215–220. 2.2, B.2
- [329] MATLAB. 9.7.0.1190202 (R2019b). The MathWorks Inc, 2018. B.5, B.7
- [330] MCFARLANE, N. J., MA, X., CLAPWORTHY, G. J., BESSIS, N., AND TESTI, D. A survey and classification of visualisation in multiscale biomedical applications. In *Proc IEEE IV* (2012), pp. 561–566. ■ . 2.1, A.1, A.13
- [331] MCGEE, F., GHONIEM, M., MELANÇON, G., OTJACQUES, B., AND PINAUD, B. The state of the art in multilayer network visualization. In *Comput Graph Forum* (2019), vol. 38, pp. 125–149. ■ . 2.1, A.1, A.5, A.13
- [332] MEHTA, R., AND ZHU, R. Blue or red? exploring the effect of color on cognitive task performances. *Science* 323, 5918 (2009), 1226–1229, doi: 10.1126/science.1169144. E.3.1
- [333] MEIER, S., HENNEMUTH, A., FRIMAN, O., BOCK, J., MARKL, M., AND PREUSSER, T. Non-invasive 4d blood flow and pressure quantification in central blood vessels via pc-mri. In *Computing in Cardiology* (2010), pp. 903–906. ■ □ ■ □ . A.13
- [334] METZ, C., BAKA, N., KIRISLI, H., SCHAAP, M., VAN WALSUM, T., KLEIN, S., NEEFJES, L., MOLLET, N., LELIEVELDT, B., DE BRUIJNE, M., AND NIESSEN, W. Conditional shape models for cardiac motion estimation. In *Proc MICCAI* (2010), vol. 13, pp. 452–459. ■ □ ■ □ . A.8, A.8.2, A.13
- [335] MEUSCHKE, M., B.KÖHLER, PREIM, U., PREIM, B., AND LAWONN, K. Semi-automatic vortex flow classification in 4D PC-MRI data of the aorta. *Comput Graph Forum* 35, 3 (2016), 351–360, doi: 10.1111/cgf.12911. ■ □ ■ □ . A.8.1, A.13
- [336] MEUSCHKE, M., GARRISON, L., SMIT, N., BRUCKNER, S., LAWONN, K., AND PREIM, B. Towards narrative medical visualization. *arXiv* (2021), doi: 10.48550/arXiv.2108.05462. D.7
- [337] MEUSCHKE, M., VOSS, S., GAIDZIK, F., PREIM, B., AND LAWONN, K. Skyscraper visualization of multiple time-dependent scalar fields on surfaces. *Comput Graph* (2021), doi: 10.1103/PhysRevFluids.6.110505. ■ □ ■ □ . A.8.1, A.13
- [338] MEUSCHKE, M., WICKENHÖFER, R., PREIM, B., AND LAWONN, K. Aneulysis—a system for aneurysm data analysis. *Comput Graph* 98 (2020), 197–209, doi: 10.1016/j.cag.2021.06.001. ■ □ ■ □ . A.8.1, A.13
- [339] MEYER, M., MUNZNER, T., DEPACE, A., AND PFISTER, H. Multeesum: A tool for comparative spatial and temporal gene expression data. *IEEE Trans Vis Comput Graph* 16, 6 (2010), 908–917, doi: 10.1109/TVCG.2010.137. ■ □ ■ □ . A.7.1, A.13
- [340] MEYER, M., MUNZNER, T., AND PFISTER, H. Mizbee: a multiscale synteny browser. *IEEE Trans Vis Comput Graph* 15, 6 (2009), 897–904, doi: 10.1109/TVCG.2009.167. □ □ ■ □ . A.5.3, A.13
- [341] MEYER, M., WONG, B., STYCZYNSKI, M., MUNZNER, T., AND PFISTER, H. Pathline: A tool for comparative functional genomics. In *Comput Graph Forum* (2010), vol. 29, pp. 1043–1052. □ □ ■ □ . 3.2.1, A.5.3, A.13, B.6
- [342] MEYER-SPRADOW, J., STEGGER, L., DÖRING, C., ROPINSKI, T., AND HINRICHS, K. Glyph-based spect visualization for the diagnosis of coronary artery disease. *IEEE Trans Vis Comput Graph* 14, 6 (2008), 1499–1506, doi: 10.1109/TVCG.2008.136. ■ □ ■ □ . A.8.2, A.13

- [343] MILLER, M. I., TWARD, D. J., AND TROUVE, A. Hierarchical computational anatomy: Unifying the molecular to tissue continuum via measure representations of the brain. *bioRxiv* (2021), doi: 10.1101/2021.04.19.440540v1.full.    . A.9, A.13
- [344] MILO, R., JORGENSEN, P., MORAN, U., WEBER, G., AND SPRINGER, M. Bionumbers—the database of key numbers in molecular and cell biology. *Nucleic Acids Res* 38, suppl_1 (2010), D750–D753, doi: 10.1093/nar/gkp889. 3.1.2, A.4.1, A.10
- [345] MOHAMMED, H., AL-AWAMI, A. K., BEYER, J., CALI, C., MAGISTRETTI, P., PFISTER, H., AND HADWIGER, M. Abstractocyte: A visual tool for exploring nanoscale astroglial cells. *IEEE Trans Vis Comput Graph* 24, 1 (2017), 853–861, doi: 10.1109/TVCG.2017.2744278.    . A.7.2, A.13
- [346] MORRIS, C. J., YANG, J. N., AND SCHEER, F. A. The impact of the circadian timing system on cardiovascular and metabolic function. *Prog Brain Res* 199 (2012), 337–358, doi: 10.1016/B978-0-444-59427-3.00019-8. A.10
- [347] MÖRTH, E., HALDORSEN, I. S., BRUCKNER, S., AND SMIT, N. N. Paraglyder: Probe-driven interactive visual analysis for multiparametric medical imaging data. In *Proc CGI* (2020), pp. 351–363.    . A.7.2, A.14, A.13
- [348] MÖRTH, E., WAGNER-LARSEN, K., HODNELAND, E., KRAKSTAD, C., HALDORSEN, I. S., BRUCKNER, S., AND SMIT, N. N. Radex: Integrated visual exploration of multiparametric studies for radiomic tumor profiling. In *Comput Graph Forum* (2020), vol. 39, pp. 611–622.    . A.7.2, A.13
- [349] MUÑOZ, R., AND OKUNIEV, D. TypeForm. <https://www.typeform.com/>, 2012. Viewed 18.06.2021. 3.3.1, D.4.3
- [350] MULDER, A. M., YOSHIOKA, C., BECK, A. H., BUNNER, A. E., MILLIGAN, R. A., POTTER, C. S., CARRAGHER, B., AND WILLIAMSON, J. R. Visualizing ribosome biogenesis: parallel assembly pathways for the 30s subunit. *Science* 330, 6004 (2010), 673–677, doi: 10.1126/science.1193220.    . A.13
- [351] MÜLLER, J., GARRISON, L., SCHREIBER, S., BRUCKNER, S., HAUSER, H., AND OELTZE-JAFFRA, S. Integrated dual analysis of quantitative and qualitative high-dimensional data. *IEEE Trans Vis Comput Graph* 27, 6 (2021), 2953—2966, doi: 10.1109/TVCG.2021.3056424. 2.2, 14, C.1, C.2.3
- [352] MULLINS, P. G., MCGONIGLE, D. J., O’GORMAN, R. L., PUTS, N. A., VIDYASAGAR, R., EVANS, C. J., EDDEN, R. A. E., ET AL. Current practice in the use of MEGA-PRESS spectroscopy for the detection of GABA. *Neuroimage* 86, 1 (2014), 43–52, doi: 10.1016/j.neuroimage.2012.12.004. B.3
- [353] MUNZNER, T. A nested model for visualization design and validation. *IEEE Trans Vis Comput Graph* 15, 6 (2009), 921–28. 2.1, D.2
- [354] MUNZNER, T. *Visualization analysis and design*. CRC press, 2014. 1.2.2, 1.2.3, 2.1, D.2
- [355] MUZIC, M. L., PARULEK, J., STAVRUM, A. K., AND VIOLA, I. Illustrative visualization of molecular reactions using omniscient intelligence and passive agents. *Comput Graph Forum* 33, 3 (2014), 141–150, doi: 10.1111/cgf.12370.    . A.5.2, A.13
- [356] MUZIC, M. L., WALDNER, M., PARULEK, J., AND VIOLA, I. Illustrative timelapse: a technique for illustrative visualization of particle simulations on the mesoscale level. In *Proc IEEE PacificVis* (2015), pp. 247–254.    . A.5.3, A.13
- [357] MÜLLER, W., NOCKE, T., AND SCHUMANN, H. Enhancing the visualization process with principal component analysis to support the exploration of trends. In *Proc APVis* (2006), pp. 121–130. 2.2, C.1
- [358] NAGLE, M. P., TAM, G. S., MALTZ, E., HEMMINGER, Z., AND WOLLMAN, R. Bridging scales: From cell biology to physiology using in situ single-cell technologies. *Cell Syst* 12, 5 (2021), 388–400, doi: 10.1016/j.cels.2021.03.002. A.6, A.7

- [359] NAKAO, M., TOKUNO, J., CHEN-YOSHIKAWA, T., DATE, H., AND MATSUDA, T. Surface deformation analysis of collapsed lungs using model-based shape matching. *Int J Comput Assist Radiol Surg* 14, 10 (2019), 1763–1774, doi: 10.1007/s11548-019-02013-0. ■ ■ ■ . A.8.3, A.13
- [360] NARDINI, P., CHEN, M., SAMSEL, F., BUJACK, R., BÖTTINGER, M., AND SCHEUERMANN, G. The making of continuous colormaps. *IEEE Trans Vis Comput Graph* 27, 6 (2019), 3048–3063, doi: 10.1109/TVCG.2019.2961674. E.4
- [361] NARECHANIA, A., KARDUNI, A., WESSLEN, R., AND WALL, E. Vitality: Promoting serendipitous discovery of academic literature with transformers & visual analytics. *IEEE Trans Vis Comput Graph* 28, 1 (2021), 486–496, doi: 10.1109/TVCG.2021.3114820. 3.2, 3.1.1, 3.1.2, A.3, A.10
- [362] NAYAK, S., AND IWASA, J. H. Preparing scientists for a visual future. *EMBO Reports* 20, 11 (2019), e49347. D.1
- [363] NAYAK, S., LIU, H., HSU, G. I., AND IWASA, J. H. Using 3D animation to visualize hypotheses. *Trends Biochem Sci* 45, 7 (2020), 633–634, doi: 10.1016/j.tibs.2020.04.009. □ ■ ■ . A.5.2, A.13
- [364] NAZIR, S., KHAN, M. N., ANWAR, S., ADNAN, A., ASADI, S., SHAHZAD, S., AND ALI, S. Big data visualization in cardiology—a systematic review and future directions. *IEEE Access* 7 (2019), 115945–115958, doi: 10.1109/ACCESS.2019.2936133. ■ . A.8.2, A.13
- [365] NEDEL, L. P., AND THALMANN, D. Real time muscle deformations using mass-spring systems. In *Proc CGI* (1998), pp. 156–165. ■ ■ ■ . A.8.5, A.13
- [366] NETZEL, R., VUONG, J., ENGELKE, U., O'DONOGHUE, S., WEISKOPF, D., AND HEINRICH, J. Comparative eye-tracking evaluation of scatterplots and parallel coordinates. *Vis Inform* 1, 2 (2017), 118–131, doi: 10.1016/j.visinf.2017.11.001. C.2.2
- [367] NEUGEBAUER, M., GASTEIGER, R., BEUING, O., DIEHL, V., SKALEJ, M., AND PREIM, B. Map displays for the analysis of scalar data on cerebral aneurysm surfaces. *Comput Graph Forum* 28, 3 (2009), 895–902, doi: 10.1111/j.1467-8659.2009.01459.x. ■ ■ ■ . A.13
- [368] NGUYEN, N., STRNAD, O., KLEIN, T., LUO, D., ALHARBI, R., WONKA, P., MARITAN, M., MINDEK, P., AUTIN, L., GOODSSELL, D. S., ET AL. Modeling in the time of COVID-19: Statistical and rule-based mesoscale models. *IEEE Trans Vis Comput Graph* 27, 2 (2021), 722, doi: 10.1109/TVCG.2020.3030415. □ ■ ■ . A.5.2, A.13
- [369] NGUYEN, T. K., EKLUND, A., OHLSSON, H., HERNELL, F., LJUNG, P., FORSELL, C., ANDERSSON, M., KNUTSSON, H., AND YNNERMAN, A. Concurrent volume visualization of real-time fMRI. In *Proc Volume Graphics* (2010), pp. 53–60. A.8.5
- [370] NIELSEN, C. B., CANTOR, M., DUBCHAK, I., GORDON, D., AND WANG, T. Visualizing genomes: Techniques and challenges. *Nat Methods* 7, 3 (2010), S5–S15, doi: 10.1038/nmeth.1422. ■ . A.5, A.10, A.13
- [371] NOBREGA, T. H. C., CARVALHO, D. D. B., AND WANGENHEIM, A. V. Simplified simulation and visualization of tubular flows with approximate centerline generation. In *Proc CBMS* (2009), pp. 1–7. ■ ■ ■ . A.8.1, A.13
- [372] NUNES, M., LARUELO, A., KEN, S., LAPRIE, A., AND BÜHLER, K. A survey on visualizing magnetic resonance spectroscopy data. In *Proc EG VCBM* (2014), pp. 21–30. ■ . 2.2, A.5.3, A.13
- [373] NUNES, M., ROWLAND, B., SCHLACHTER, M., KEN, S., MATKOVIC, K., LAPRIE, A., AND BÜHLER, K. An integrated visual analysis system for fusing MR spectroscopy and multi-modal radiology imaging. In *Proc IEEE VAST* (2014), pp. 53–62. ■ ■ ■ . 2.2, A.13, B.2
- [374] NUSRAT, S., HARBIG, T., AND GEHLENBORG, N. Tasks, techniques, and tools for genomic data visualization. In *Comput Graph Forum* (2019), vol. 38, pp. 781–805. ■ . A.5, A.13
- [375] O'DONOGHUE, S. I. Grand challenges in bioinformatics data visualization. *Frontiers in Bioinformatics* 1 (2021), 13, doi: 10.3389/fbinf.2021.669186. 3.1, A.1
































- [376] O'DONOGHUE, S. I., BALDI, B. F., CLARK, S. J., DARLING, A. E., HOGAN, J. M., KAUR, S., MAIER-HEIN, L., MCCARTHY, D. J., MOORE, W. J., STENAU, E., ET AL. Visualization of biomedical data. *Annu Rev Biomed Data Sci* 1 (2018), 275–304, doi: 10.1146/annurev-biodatasci-080917-013424. ■ . 2.1, A.1, A.13
- [377] O'DONOGHUE, S. I., GOODSSELL, D. S., FRANGAKIS, A. S., JOSSINET, F., LASKOWSKI, R. A., NILGES, M., SAIBIL, H. R., SCHAFFERHANS, A., WADE, R. C., WESTHOF, E., AND OLSON, A. J. Visualization of macromolecular structures. *Nat Methods* 7, 3s (2010), S42–S55, doi: 10.1038/nmeth.1427. ■ . A.5, A.10, A.13
- [378] O'DONOGHUE, S. I., GOODSSELL, D. S., FRANGAKIS, A. S., JOSSINET, F., LASKOWSKI, R. A., NILGES, M., SAIBIL, H. R., SCHAFFERHANS, A., WADE, R. C., WESTHOF, E., AND OLSON, A. J. Visualizing biological data — now and in the future. *Nat Methods* 7, 3s (2010), S2–S4, doi: 10.1038/nmeth.f.301. ■ . 3.1, A.1, A.10, A.13
- [379] OELTZE, S., DOLEISCH, H., HAUSER, H., MUIGG, P., AND PREIM, B. Interactive visual analysis of perfusion data. *IEEE Trans Vis Comput Graph* 13, 6 (2007), 1392–1399, doi: 10.1109/TVCG.2007.70569. ■ ■ ■ ■ . A.8.2, A.13
- [380] OELTZE, S., GROTHUES, F., HENNEMUTH, A., AND PREIM, B. Integrated visualization of morphologic and perfusion data for the analysis of coronary artery disease. In *Proc VisSym* (2006), pp. 131–138. ■ ■ ■ ■ . A.8.2, A.13
- [381] OELTZE, S., LEHMANN, D. J., KUHN, A., JANIGA, G., THEISEL, H., AND PREIM, B. Blood flow clustering and applications in virtual stenting of intracranial aneurysms. *IEEE Trans Vis Comput Graph* 20, 5 (2014), 686–701, doi: 10.1109/TVCG.2013.2297914. ■ ■ ■ ■ . A.8.1, A.13
- [382] OELTZE, S., PREIM, B., HAUSER, H., RØRVIK, J., AND LUNDERVOLD, A. Visual analysis of cerebral perfusion data four interactive approaches and a comparison. In *Proc ISPA* (2009), pp. 582–589. ■ ■ ■ ■ . A.7.2, A.13
- [383] OELTZE-JAFRA, S., CEBRAL, J. R., JANIGA, G., AND PREIM, B. Cluster analysis of vortical flow in simulations of cerebral aneurysm hemodynamics. *IEEE Trans Vis Comput Graph* 22, 1 (2016), 757–766, doi: 10.1109/TVCG.2015.2467203. ■ ■ ■ ■ . A.8.1, A.13
- [384] OELTZE-JAFRA, S., MEUSCHKE, M., NEUGEBAUER, M., SAALFELD, S., LAWONN, K., JANIGA, G., HEGE, H.-C., ZACHOW, S., AND PREIM, B. Generation and visual exploration of medical flow data: Survey, research trends and future challenges. In *Comput Graph Forum* (2019), vol. 38, pp. 87–125. ■ . A.8.1, A.13, D.4.2, D.4.2
- [385] OELTZE-JAFRA, S., YNNERMAN, A., BRUCKNER, S., AND HAUSER, H. Rejuvenated medical visualization: Large-scale, whole-body visualization, visualizing physiology, non-standard imaging and simulations, and cohort studies. In *Proc IEEE VIS (Tutorials)* (2015). 3.1, A.1
- [386] OU, H. D., PHAN, S., DEERINCK, T. J., THOR, A., ELLISMAN, M. H., AND O'SHEA, C. C. ChromEMT: Visualizing 3D chromatin structure and compaction in interphase and mitotic cells. *Science* 357, 6349 (2017), eaag0025, doi: 10.1126/science.aag0025. ■ ■ ■ ■ . A.6.1, A.13
- [387] PAGES, J. Analyse factorielle de données mixtes. *Revue de Statistique Appliquée* 52, 4 (2004), 93–111. 14, 14
- [388] PÁLENIK, J., BYŠKA, J., BRUCKNER, S., AND HAUSER, H. Scale-space splatting: Reforming spacetime for cross-scale exploration of integral measures in molecular dynamics. *IEEE Trans Vis Comput Graph* 26, 1 (2019), 643–653, doi: 10.1109/TVCG.2019.2934258. ■ ■ ■ ■ . A.5.2, A.13
- [389] PALSSON, E. A 3-d model used to explore how cell adhesion and stiffness affect cell sorting and movement in multicellular systems. *J Theor Biol* 254, 1 (2008), 1–13, doi: 10.1016/j.jtbi.2008.05.004. ■ ■ ■ ■ . A.6.2, A.13
- [390] PARK, D., DRUCKER, S., FERNANDEZ, R., AND ELMQVIST, N. Atom: A grammar for unit visualizations. *IEEE Trans Vis Comput Graph* 24, 12 (2017), 3032 – 3043, doi: 10.1109/TVCG.2017.2785807. 2.2, 3.2.1, B.2, B.6

- [391] PARK, J., AND PARK, S. I. Strain analysis and visualization: left ventricle of a heart. *Comput Graph* 24, 5 (2000), 701–714, doi: 10.1016/S0097-8493(00)00073-X. ■ ■ ■. A.8.2, A.13
- [392] PASI, M., BOULOUIS, G., FOTIADIS, P., AURIEL, E., CHARIDIMOU, A., HALEY, K., AYRES, A., SCHWAB, K. M., GOLDSTEIN, J. N., ROSAND, J., ET AL. Distribution of lacunes in cerebral amyloid angiopathy and hypertensive small vessel disease. *Neurology* 88, 23 (2017), 2162–2168, doi: 10.1212/WNL.0000000000004007. C.3.3
- [393] PAVLOPOULOS, G. A., MALLIARAKIS, D., PAPANIKOLAOU, N., THEODOSIOU, T., ENRIGHT, A. J., AND ILIOPOULOS, I. Visualizing genome and systems biology: technologies, tools, implementation techniques and trends, past, present and future. *Gigascience* 4, 1 (2015), s13742–015, doi: 10.1186/s13742-015-0077-2. ■. A.5, A.13
- [394] PEARSON, K. Liii. on lines and planes of closest fit to systems of points in space. *Lond Edinb Dublin Philos Mag J Sci* 2, 11 (1901), 559–572, doi: 10.1080/14786440109462. 2.2, 14, C.1, 14
- [395] PERIN, C., DRAGICEVIC, P., AND FEKETE, J. Bertifier: New interactions for crafting tabular visualizations. *IEEE Trans Vis Comput Graph* 20, 12 (2014), 2082 – 2091, doi: 10.1109/TVCG.2014.2346279. B.5
- [396] PERNOD, E., SERMESANT, M., RELAN, J., AND DELINGETTE, H. Interactive real time simulation of cardiac radio-frequency ablation. In *Proc EG VCBM* (2010), pp. 91–98. ■ ■ ■. A.8.2, A.13
- [397] PETERSCH, B., AND HÖNIGMANN, D. Blood flow in its context: Combining 3D B-mode and color doppler ultrasonic data. *IEEE Trans Vis Comput Graph* 13, 4 (2007), 748–757, doi: 10.1109/TVCG.2007.1018. ■ ■ ■. A.8.1, A.13
- [398] PEZESHKIAN, W., KÖNIG, M., WASSENAAR, T. A., AND MARRINK, S. J. Backmapping triangulated surfaces to coarse-grained membrane models. *Nat Commun* 11, 1 (2020), 1–9, doi: 10.1038/s41467-020-16094-y. ■ ■ ■. A.10, A.6.1, A.13
- [399] PFISTER, H., KAYNIG, V., BOTHA, C. P., BRUCKNER, S., DERCKSEN, V. J., HEGE, H.-C., AND ROERDINK, J. B. Visualization in connectomics. In *Scientific Visualization*. 2014, pp. 221–245. ■. A.8, A.8.4, A.13
- [400] POTTER, K., KNISS, J., RIESENFELD, R., AND JOHNSON, C. R. Visualizing summary statistics and uncertainty. vol. 29, pp. 823–832. B.6
- [401] PREIM, B., AND MEUSCHKE, M. A survey of medical animations. *Comput Graph* 90 (2020), 145–168, doi: 10.1016/j.cag.2020.06.003. ■. 2.1, A.1, A.8, A.13
- [402] PREIM, B., AND MEUSCHKE, M. A survey of medical animations. *Comput Graph* 90 (2020), 145–68. 3.3.1, D.6
- [403] PREIM, B., OELTZE, S., MLEJNEK, M., GROLLER, E., HENNEMUTH, A., AND BEHRENS, S. Survey of the visual exploration and analysis of perfusion data. *IEEE Trans Vis Comput Graph* 15, 2 (2009), 205–220, doi: 10.1109/TVCG.2008.95. ■. A.7, A.13
- [404] PREIM, B., ROPINSKI, T., AND ISENBERG, P. A critical analysis of the evaluation practice in medical visualization. In *Proc VCBM* (2018), pp. 45–56. D.4
- [405] PRETORIUS, A. J., KHAN, I. A., AND ERRINGTON, R. J. Cell lineage visualisation. *Comput Graph Forum* 34, 3 (2015), 21–30, doi: 10.1111/cgf.12614. ■ ■ ■. A.6.2, A.13
- [406] PRETORIUS, A. J., KHAN, I. A., AND ERRINGTON, R. J. A survey of visualization for live cell imaging. In *Comput Graph Forum* (2017), vol. 36, pp. 46–63. ■. A.6, A.13
- [407] PROKOP, Z., GORA, A., BREZOVSKY, J., CHALOUPOKOVA, R., STEPANKOVA, V., AND DAMBORSKY, J. Engineering of protein tunnels: Keyhole-lock-key model for catalysis by the enzymes with buried active sites. In *Protein Engineering Handbook*. 2012, pp. 421–64. D.4.2

- [408] PROVENCHER, S. W. Automatic quantitation of localized in vivo ¹H spectra with LCMoDel. *NMR in Biomedicine* 14, 4 (2001), 260–264, doi: 10.1002/nbm.698. 3.6, B.1, B.3
- [409] PURVIS, L. A. B., CLARKE, W. T., BIASIOLLI, L., VALKOVIČ, L., ROBSON, M. D., AND RODGERS, C. T. OXSA: An open-source magnetic resonance spectroscopy analysis toolbox in MATLAB. *PLoS One* 12, 9 (2017), e0185356, doi: 10.1371/journal.pone.0185356. B.3
- [410] QIN, Y., HUTTLIN, E. L., WINSNES, C. F., GOSZTYLA, M. L., WACHEUL, L., KELLY, M. R., BLUE, S. M., ZHENG, F., CHEN, M., SCHAFFER, L. V., ET AL. A multi-scale map of cell structure fusing protein images and interactions. *Nature* 600, 7889 (2021), 536–542, doi: 10.1038/s41586-021-04115-9. . A.6.1, A.13
- [411] QUTUB, A. A., MAC GABHANN, F., KARAGIANNIS, E. D., VEMPATI, P., AND POPEL, A. S. Multiscale models of angiogenesis. *IEEE Eng Med Biol Mag* 28, 2 (2009), 14–31, doi: 10.1109/MEMB.2009.931791. . 3.1.2, A.7, A.7.1, A.9, A.13
- [412] RAUTEK, P., BRUCKNER, S., GRÖLLER, E., AND VIOLA, I. Illustrative visualization: New technology or useless tautology? *ACM SIGGRAPH Comput Graph* 42, 3 (2008). 1.1, 2.3, D.1, D.2, D.3
- [413] RAYMAN, M. P. Selenium and human health. *Lancet* 379, 9822 (2012), 1256–1268, doi: 10.1016/S0140-6736(11)61452-9. C.3.2
- [414] RECLUSA, P., VERSTRAELEN, P., TAVERNA, S., GUNASEKARAN, M., PUCCI, M., PINTELO, I., CLAES, N., DE MIGUEL-PÉREZ, D., ALESSANDRO, R., BALS, S., ET AL. Improving extracellular vesicles visualization: From static to motion. *Sci Rep* 10, 1 (2020), 1–9, doi: 10.1038/s41598-020-62920-0. . A.6.1, A.13
- [415] REISACHER, M., VIOLA, I., AND MUZIC, M. L. Cellpathway: A simulation tool for illustrative visualization of biochemical networks. In *Proc WSCG* (2016), pp. 99–105. D.4.2
- [416] RHODES, O., PERES, L., ROWLEY, A. G., GAIT, A., PLANA, L. A., BRENNINKMEIJER, C., AND FURBER, S. B. Real-time cortical simulation on neuromorphic hardware. *Philos Trans R Soc A* 378, 2164 (2020), 20190160, doi: 10.1098/rsta.2019.0160. . A.7.2, A.13
- [417] RHYNE, T.-M. Applying artistic color theories to visualization. In *Expanding the Frontiers of Visual Analytics and Visualization*. 2012, pp. 263–283. 1.2.3, E.1
- [418] RICHER, G., SANSEN, J., LALANNE, F., AUBER, D., AND BOURQUI, R. Enabling hierarchical exploration for large-scale multidimensional data with abstract parallel coordinates. In *Proc BigVis* (2018), vol. 2083, pp. 76–83. 2.2, C.1
- [419] RINIKER, S. Fixed-charge atomistic force fields for molecular dynamics simulations in the condensed phase: An overview. *J Chem Inf Model* 58, 3 (2018), 565–578, doi: 10.1021/acs.jcim.8b00042. A.5.1
- [420] RISPOLI, V. C., NIELSEN, J. F., NAYAK, K. S., AND CARVALHO, J. L. Computational fluid dynamics simulations of blood flow regularized by 3D phase contrast MRI. *Biomed Eng Online* 14, 1 (2015), 1–23, doi: 10.1186/s12938-015-0104-7. . A.8, A.8.1, A.13
- [421] ROBERTSON, G. G., MACKINLAY, J. D., AND CARD, S. The perspective wall: Detail and context smoothly integrated. In *Proc ACM CHI* (1991), vol. 91, pp. 173–179. C.2.2
- [422] ROHRSCHEIDER, M., HEINE, C., REICHENBACH, A., KERREN, A., AND SCHEUERMANN, G. A novel grid-based visualization approach for metabolic networks with advanced focus & context view. In *Proc Graph Drawing* (2010), pp. 268–279. . A.13
- [423] ROMEO, M., MONTEAGUDO, C., AND SÁNCHEZ-QUIRÓS, D. Muscle and fascia simulation with extended position based dynamics. In *Comput Graph Forum* (2020), vol. 39, pp. 134–146. . A.8.5, A.9, A.13
- [424] ROPINSKI, T., VIOLA, I., BIERMANN, M., HAUSER, H., AND HINRICHS, K. Multimodal visualization with interactive closeups. In *Proc TPCG* (2009), pp. 17–24. . A.8.5, A.13

- [425] ROSEN, P., BURTON, B., POTTER, K., AND JOHNSON, C. R. muview: A visual analysis system for exploring uncertainty in myocardial ischemia simulations. In *Visualization in Medicine and Life Sciences III*. 2016, pp. 49–69. ■ ■□□. A.8.2, A.13
- [426] RÖSSLER, F., TEJADA, E., FANGMEIER, T., ERTL, T., AND KNAUFF, M. Gpu-based multi-volume rendering for the visualization of functional brain images. In *Proc SimVis* (2006), pp. 305–18. ■ ■□□. A.8.4, A.13
- [427] RYAN, M., MCGILL, G., AND WILSON, E. O. *E.O.Wilson's life on earth*. Wilson Biodiversity Foundation, 2014. ■ □□■. 3.1, A.1, A.13
- [428] RZEPECKI, J., VAQUERO, R. M. M., VAIS, A., FRIESE, K.-I., AND WOLTER, F.-E. Multimodal approach for natural biomedical multi-scale exploration. In *Proc ISVC* (2014), pp. 620–631. ■ ■□□. A.9, A.13
- [429] SACHA, D., ZHANG, L., SEDLMAIR, M., LEE, J. A., PELTONEN, J., WEISKOPF, D., NORTH, S. C., AND KEIM, D. A. Visual interaction with dimensionality reduction: A structured literature analysis. *IEEE Trans Vis Comput Graph* 23, 1 (2017), 241–250, doi: 10.1109/TVCG.2016.2598495. 2.2, C.1, C.2
- [430] SADRIEH, A., DOMANSKI, L., PITT-FRANCIS, J., MANN, S. A., HODKINSON, E. C., NG, C.-A., PERRY, M. D., TAYLOR, J. A., GAVAGHAN, D., SUBBIAH, R. N., ET AL. Multiscale cardiac modelling reveals the origins of notched t waves in long qt syndrome type 2. *Nat Commun* 5, 1 (2014), 1–11, doi: 10.1038/ncomms6069. ■ ■□□. A.9, A.13
- [431] SAKAUE-SAWANO, A., KUROKAWA, H., MORIMURA, T., HANYU, A., HAMA, H., OSAWA, H., KASHIWAGI, S., FUKAMI, K., MIYATA, T., MIYOSHI, H., IMAMURA, T., OGAWA, M., MASAI, H., AND MIYAWAKI, A. Visualizing spatiotemporal dynamics of multicellular cell-cycle progression. *Cell* 132, 3 (2008), 487–498, doi: 10.1016/j.cell.2007.12.033. ■ ■□□. A.6.1, A.13
- [432] SALVADOR-MARTÍNEZ, I., GRILLO, M., AVEROF, M., AND TELFORD, M. J. Celavi: an interactive cell lineage visualization tool. *Nucleic Acids Res* (2021), doi: 10.1093/nar/gkab325. ■ ■□□. A.6.2, A.13
- [433] SAMATOVA, N. F., BREIMYER, P., HENDRIX, W., SCHMIDT, M. C., AND RHYNE, T. An outlook into ultra-scale visualization of large-scale biological data. In *Proc UltraVis* (2008), pp. 29–39. ■. A.5, A.13
- [434] SAND, O., SJAASTAD, Ø. V., HAUG, E., AND TOVERUD, K. C. *Menneskets fysiologi*, 3 ed. Gyldendal Akademisk, 2014. ■ □□■. 3.1, A.1, A.13
- [435] SANTHANAM, A. P., IMIELINSKA, C., DAVENPORT, P., KUPELIAN, P., AND ROLLAND, J. P. Modeling real-time 3-d lung deformations for medical visualization. *IEEE Trans Inf Technol Biomed* 12, 2 (2008), 257–270, doi: 10.1109/TITB.2007.899489. ■ ■□□. A.8.3, A.13
- [436] SAPORTA, G. Simultaneous analysis of qualitative and quantitative data. In *Proc Societa Italiana di Statistica* (1990), vol. 1, pp. 62–72. C.2.4
- [437] SARMA, G. P., LEE, C. W., PORTEGYS, T., GHAYOOMIE, V., JACOBS, T., ALICEA, B., CANTARELLI, M., CURRIE, M., GERKIN, R. C., GINGELL, S., ET AL. Openworm: Overview and recent advances in integrative biological simulation of *caenorhabditis elegans*. *Philos Trans R Soc B* 373, 1758 (2018), 20170382, doi: 10.1098/rstb.2017.0382. ■ ■□□. A.9, A.13
- [438] SCHAFER, J. L., AND GRAHAM, J. W. Missing data: our view of the state of the art. *Psychol Methods* 7, 2 (2002), 147, doi: 10.1037/1082-989X.7.2.147. 14
- [439] SCHATZ, K., FRANCO-MORENO, J. J., SCHÄFER, M., ROSE, A. S., FERRARIO, V., PLEISS, J., VÁZQUEZ, P.-P., ERTL, T., AND KRONE, M. Visual analysis of large-scale protein-ligand interaction data. In *Comput Graph Forum* (2021), vol. 40, pp. 394–408. ■ ■□□. A.5.2, A.8, A.13
- [440] SCHATZ, K., KRONE, M., PLEISS, J., AND ERTL, T. Interactive visualization of biomolecules' dynamic and complex properties. *Eur Phys J Spec Top* 227, 14 (2019), 1725–1739, doi: 10.1140/EPJST/E2019-800162-Y. ■. A.5, A.13

- [441] SCHLACHTER, M., RAIDOU, R. G., MUREN, L. P., PREIM, B., PUTORA, P. M., AND BÜHLER, K. State-of-the-art report: Visual computing in radiation therapy planning. In *Comput Graph Forum* (2019), vol. 38, pp. 753–779. ■. A.7, A.13
- [442] SCHLEICH, J.-M., ALMANGE, C., DILLENSEGER, J.-L., AND COATRIEUX, J.-L. Understanding normal cardiac development using animated models. *IEEE Comput Graph Appl* 22, 1 (2002), 14–19, doi: 10.1109/38.974513. ■ □ □ ■. A.7.1, A.13
- [443] SCHLOSS, K. B., LESSARD, L., WALMSLEY, C. S., AND FOLEY, K. Color inference in visual communication: the meaning of colors in recycling. *Cogn Res: Princ Implic* 3, 1 (2018), 1–17, doi: 10.1186/s41235-018-0090-y. 2.3, E.2
- [444] SCHMIDT-EHRENBERG, J., BAUM, D., AND HEGE, H.-C. Visualizing dynamic molecular conformations. In *Proc IEEE VIS* (2002), pp. 235–242. □ □ □ ■. A.5.1, A.13
- [445] SCHNEIDER, R. K., AND COSTA, I. G. Crosstalk: analysis and visualization of ligand–receptor networks. *Bioinformatics* 1 (2021), 3, doi: 10.1093/bioinformatics/btab370. □ □ □ ■. A.5.2, A.13
- [446] SCHÖLL, A., BRAUN, C., DAUB, M., SCHNEIDER, G., AND WUNDERLICH, H.-J. Adaptive parallel simulation of a two-timescale model for apoptotic receptor-clustering on gpus. In *Proc BIBM* (2014), pp. 424–431. □ □ □ ■. A.6.1, A.13
- [447] SCHOLZ, M., AND SELBIG, J. Visualization and analysis of molecular data. In *Metabolomics*. 2007, pp. 87–104. ■. A.5, A.13
- [448] SCHREIBER, F. Comparison of metabolic pathways using constraint graph drawing. In *Proc Asia-Pacific Bioinformatics Conf* (2003), pp. 105–110. □ □ □ ■. A.5.3, A.13
- [449] SCHREINER, S. M., KOO, P. K., ZHAO, Y., MOCHRIE, S. G., AND KING, M. C. The tethering of chromatin to the nuclear envelope supports nuclear mechanics. *Nat Commun* 6, 1 (2015), 1–13, doi: 10.1038/ncomms8159. □ □ □ ■. A.6.1, A.13
- [450] SCHRÖDINGER, LLC. The PyMOL molecular graphics system. PyMOL, The PyMOL Molecular Graphics System, Schrödinger, LLC., 2015. A.5.1
- [451] SCHWENKE, M., HENNEMUTH, A., FISCHER, B., AND FRIMAN, O. Blood flow computation in phase-contrast mri by minimal paths in anisotropic media. In *Proc MICCAI* (2011). ■ □ □ ■. A.13
- [452] SECRIER, M., AND SCHNEIDER, R. Visualizing time-related data in biology, a review. *Brief Bioinformatics* 15, 5 (2014), 771–782, doi: 10.1093/bib/bbt021. ■. 2.1, 3.1.1, A.1, A.4.1, A.10, A.13
- [453] SEDLMAIR, M., BREHMER, M., INGRAM, S., AND MUNZNER, T. Dimensionality reduction in the wild : Gaps and guidance. *Dept Comput Sci, Univ British Columbia, Vancouver, BC, Canada, Tech Rep TR-2012-03* (2012). C.2
- [454] SEGAL, E., YELENSKY, R., KAUSHAL, A., PHAM, T., REGEV, A., KOLLER, D., AND FRIEDMAN, N. Genexpress: a visualization and statistical analysis tool for gene expression and sequence data. In *Proc ISMB* (2004), vol. 18. □ □ □ ■. A.13
- [455] SENGUPTA, P. P., PEDRIZZETTI, G., KILNER, P. J., KHERADVAR, A., EBBERS, T., TONTI, G., FRASER, A. G., AND NARULA, J. Emerging trends in CV flow visualization. *JACC Cardiovasc Imaging* 5, 3 (2012), 305–316, doi: 10.1016/j.jcmg.2012.01.003. A.8
- [456] SETLUR, V., AND STONE, M. C. A linguistic approach to categorical color assignment for data visualization. *IEEE Trans Vis Comput Graph* 22, 1 (2015), 698–707, doi: 10.1109/TVCG.2015.2467471. 2.3, E.2
- [457] SFORZA, D. M., PUTMAN, C. M., AND CEBRAL, J. R. Hemodynamics of cerebral aneurysms. *Annu Rev Fluid Mech* 41 (2009), 91–107. D.4.2

- [458] SHAIKH, B., MARUPILLA, G., WILSON, M., BLINOV, M. L., MORARU, I. I., AND KARR, J. R. runbiosimulations: an extensible web application that simulates a wide range of computational modeling frameworks, algorithms, and formats. *bioRxiv* (2021), doi: 10.1093/nar/gkab411.    . A.5.3, A.13
- [459] SHALLBETTER, J. L. F. Color and contrast. WORQX: Design, Resources, and Tutorials, 2021. https://worqx.com/color/color_contrast.htm. E.4
- [460] SHANNON, P., MARKIEL, A., OZIER, O., BALIGA, N. S., WANG, J. T., RAMAGE, D., AMIN, N., SCHWIKOWSKI, B., AND IDEKER, T. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res* 13, 11 (2003), 2498–2504, doi: 10.1101/gr.1239303.    . A.5.3, A.13
- [461] SHAW, D. E., MARAGAKIS, P., LINDORFF-LARSEN, K., PIANA, S., DROR, R. O., EASTWOOD, M. P., BANK, J. A., JUMPER, J. M., SALMON, J. K., SHAN, Y., ET AL. Atomic-level characterization of the structural dynamics of proteins. *Science* 330, 6002 (2010), 341–346, doi: 10.1126/science.1187409. A.5
- [462] SHEHARYAR, A., CHITIBOI, T., KELLER, E., RAHMAN, O., SCHNELL, S., MARKL, M., BOUHALI, O., AND LINSEN, L. Spatio-temporal visualization of regional myocardial velocities. In *Proc EG VCBM* (2016), pp. 89–98.    . A.8.2, A.13
- [463] SHERBONDY, A., AKERS, D., MACKENZIE, R., DOUGHERTY, R., AND WANDELL, B. Exploring connectivity of the brain’s white matter with dynamic queries. *IEEE Trans Vis Comput Graph* 11, 4 (2005), 419–430, doi: 10.1109/TVCG.2005.59.    . A.8.4, A.13
- [464] SHINOHARA, M., SABRA, K., GENNISSON, J.-L., FINK, M., AND TANTER, M. Real-time visualization of muscle stiffness distribution with ultrasound shear wave imaging during muscle contraction. *Muscle Nerve* 42, 3 (2010), 438–441, doi: 10.1002/mus.21723.    . A.8.5, A.13
- [465] SILVA, S., SANTOS, B. S., AND MADEIRA, J. Using color in visualization: A survey. *Comput Graph* 35, 2 (2011), 320–333, doi: 10.1016/j.cag.2010.11.015. 2.1, E.2
- [466] SIMÕES, T., LOPES, D., DIAS, S., FERNANDES, F., PEREIRA, J., JORGE, J., BAJAJ, C., AND GOMES, A. Geometric detection algorithms for cavities on protein surfaces in molecular graphics: a survey. In *Comput Graph Forum* (2017), vol. 36, pp. 643–683.  . A.5, A.13
- [467] SOLOVEY, M., AND SCIALDONE, A. COMUNET: A tool to explore and visualize intercellular communication. *Bioinformatics* 36, 15 (2020), 4296–4300, doi: 10.1093/bioinformatics/btaa482.    . A.6.2, A.13
- [468] SOMARAKIS, A., ISSSELSTEIJN, M. E., LUK, S. J., KENKHUIS, B., DE MIRANDA, N. F., LELIEVELDT, B. P., AND HÖLLT, T. Visual cohort comparison for spatial single-cell omics-data. *arXiv* (2020), doi: 10.1109/TVCG.2020.3030336.    . A.7.1, A.13
- [469] SONI, B., THOMPSON, D., AND MACHIRAJU, R. Visualizing particle/flow structure interactions in the small bronchial tubes. *IEEE Trans Vis Comput Graph* 14, 6 (2008), 1412–1427, doi: 10.1109/TVCG.2008.183.    . A.8.3, A.13
- [470] SORGER, J., MINDEK, P., KLEIN, T., JOHNSON, G., AND VIOLA, I. Illustrative transitions in molecular visualization via forward and inverse abstraction transform. In *Proc EG VCBM* (2016), pp. 21–30.    . A.9, A.13
- [471] SOUSA, M. C., GOOCH, A. A., AND GOOCH, B. Illustrative scientific visualization framework. In *Proc Computational Aesthetics* (2005), p. 57–65. 2.3, D.2
- [472] SOUTHERN, J., PITT-FRANCIS, J., WHITELEY, J., STOKELEY, D., KOBASHI, H., NOBES, R., KADOOKA, Y., AND GAVAGHAN, D. Multi-scale computational modelling in biology and physiology. *Prog Biophys Mol Biol* 96, 1-3 (2008), 60–89, doi: 10.1016/j.pbiomolbio.2007.07.019.    . 3.1.1, A.4.1, A.6, A.7, A.8, A.13
- [473] SPORNS, O., TONONI, G., AND KÖTTER, R. The human connectome: a structural description of the human brain. *PLoS Comput Biol* 1, 4 (2005), e42, doi: 10.1371/journal.pcbi.0010042. A.8.4

- [474] SRIVATSAN, S. R., REGIER, M. C., BARKAN, E., FRANKS, J. M., PACKER, J. S., GROSJEAN, P., DURAN, M., SAXTON, S., LADD, J. J., SPIELMANN, M., ET AL. Embryo-scale, single-cell spatial transcriptomics. *Science* 373, 6550 (2021), 111–117, doi: 10.1126/science.abb9536. ■ ■ ■ . 1, A.12, A.6.2, A.13
- [475] STAGG, C., AND ROTHMAN, D. L. *Magnetic resonance spectroscopy: tools for neuroscience research and emerging clinical applications*. Academic Press, 2013. 3.2.1, B.3, B.4.1, B.6, B.6
- [476] STALLING, D., WESTERHOFF, M., HEGE, H.-C., ET AL. Amira: A highly interactive system for visual data analysis. *Visualization Handbook 38* (2005), 749–67, doi: 10.1016/B978-012387582-2/50040-X. ■ ■ ■ . 3.1.2, A.10, A.13
- [477] STANKOVIC, Z., ALLEN, B. D., GARCIA, J., JARVIS, K. B., AND MARKL, M. 4D flow imaging with MRI. *Cardiovasc Diagn Ther* 4, 2 (2014), 173–192, doi: 10.3978/j.issn.2223-3652.2014.01.02. ■ . A.8.1, A.13
- [478] STARY, H. C., CHANDLER, A. B., GLAGOV, S., GUYTON, J. R., INSULL JR, W., ROSENFELD, M. E., SCHAFFER, S. A., SCHWARTZ, C. J., WAGNER, W. D., AND WISSLER, R. W. A definition of initial, fatty streak, and intermediate lesions of atherosclerosis. a report from the committee on vascular lesions of the council on arteriosclerosis, american heart association. *Circulation* 89, 5 (1994), 2462–2478, doi: 10.1161/01.CIR.89.5.2462. C.3.2
- [479] STEFAN, D., DI CESARE, F., ANDRASESCU, A., POPA, E., LAZARIEV, A., VESCOVO, E., STRBAK, O., WILLIAMS, S., STARCUK, Z., CABANAS, M., ET AL. Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package. *Meas Sci Technol* 20, 10 (2009), 104035, doi: 10.1088/0957-0233/20/10/104035. B.3
- [480] STILES, J. R., BARTOL, T. M., ET AL. Monte carlo methods for simulating realistic synaptic microphysiology using MCell. In *Computational Neuroscience: Realistic Modeling for Experimentalists*. 2001, pp. 87–127. D.1, D.4.2
- [481] STILES, J. R., VAN HELDEN, D., BARTOL, T. M., SALPETER, E. E., AND SALPETER, M. M. Miniature endplate current rise times less than 100 microseconds from improved dual recordings can be modeled with passive acetylcholine diffusion from a synaptic vesicle. *Proc National Academy of Sciences of the United States of America* 93, 12 (1996), 5747–52. D.1, D.4.2
- [482] STOLTE, C., TANG, D., AND HANRAHAN, P. Polaris: A system for query, analysis, and visualization of multidimensional relational databases. *IEEE Trans Vis Comput Graph* 8, 1 (2002), 52–65, doi: 10.1145/1400214.1400234. 2.2, B.2
- [483] STOPPEL, S., HODNELAND, E., HAUSER, H., AND BRUCKNER, S. Graxels: Information rich primitives for the visualization of time-dependent spatial data. In *Proc EG VCBM* (2016), pp. 183–192. ■ ■ ■ . A.8.3, A.13
- [484] STURM, W., SCHRECK, T., HOLZINGER, A., AND ULLRICH, T. Discovering medical knowledge using visual analytics. In *Proc EG VCBM* (2015), pp. 71–81. ■ . A.5, A.13
- [485] SUN, Y., AND PERIASAMY, A. Localizing protein–protein interactions in living cells using fluorescence lifetime imaging microscopy. In *Advanced Fluorescence Microscopy*. 2015, pp. 83–107. A.5
- [486] SUNDAR, H., LITT, H., AND SHEN, D. Estimating myocardial motion by 4d image warping. *Pattern Recognit* 42, 11 (2009), 2514–2526, doi: 10.1016/j.patcog.2009.04.022. ■ ■ ■ . A.8.2, A.13
- [487] SUNDARAM, T. A., AND GEE, J. C. Towards a model of lung biomechanics: pulmonary kinematics via registration of serial lung images. *Med Image Anal* 9, 6 (2005), 524–537, doi: 10.1016/j.media.2005.04.002. ■ ■ ■ . A.8.3, A.13
- [488] SUNDÉN, E., BOCK, A., JÖNSSON, D., YNNERMAN, A., AND ROPINSKI, T. Interaction techniques as a communication channel when presenting 3d visualizations. In *Proc 3DVis* (2014), pp. 61–65. D.4

- [489] SWANBERG, K. M., LANDHEER, K., PITT, D., AND JUCHEM, C. Quantifying the metabolic signature of multiple sclerosis by in vivo proton magnetic resonance spectroscopy: Current challenges and future outlook in the translation from proton signal to diagnostic biomarker. *Front Neurol* 10 (2019), doi: <https://doi.org/10.3389/fneur.2019.01173>. 3.2.1, B.1
- [490] SZAFIR, D. A., SARIKAYA, A., AND GLEICHER, M. Lightness constancy in surface visualization. *IEEE Trans Vis Comput Graph* 22, 9 (2015), 2107–2121, doi: 10.1109/TVCG.2015.2500240. 2.1, 3.3.2, E.1, E.2
- [491] TAHERI-ARAGHI, S., BROWN, S. D., SAULS, J. T., MCINTOSH, D. B., AND JUN, S. Single-cell physiology. *Annu Rev Biophys* 44 (2015), 123–142, doi: 10.1146/annurev-biophys-060414-034236. A.6
- [492] TANAKA, M., SAKAMOTO, T., SUGAWARA, S., NAKAJIMA, H., KATAHIRA, Y., OHTSUKI, S., AND KANAI, H. Blood flow structure and dynamics, and ejection mechanism in the left ventricle: Analysis using echo-dynamography. *J Cardiol* 52, 2 (2008), 86–101, doi: 10.1016/j.jjcc.2008.05.005. ■ ■ ■ . A.8.1, A.13
- [493] TATU, A., MAASS, F., FÄRBER, I., BERTINI, E., SCHRECK, T., SEIDL, T., AND KEIM, D. Subspace search and visualization to make sense of alternative clusterings in high-dimensional data. In *Proc IEEE VAST* (2012), pp. 63–72. 2.2, C.1
- [494] TAUTZ, L., HÜLLEBRAND, M., STEINMETZ, M., VOIT, D., FRAHM, J., AND HENNEMUTH, A. Exploration of interventricular septum motion in multi-cycle cardiac mri. In *Proc EG VCBM* (2017), pp. 169–178. ■ ■ ■ . A.17, A.8.2, A.13
- [495] TERAN, J., SIFAKIS, E., BLEMKER, S. S., NG-THOW-HING, V., LAU, C., AND FEDKIW, R. Creating and simulating skeletal muscle from the visible human data set. *IEEE Trans Vis Comput Graph* 11, 3 (2005), 317–328, doi: 10.1109/TVCG.2005.42. ■ ■ ■ . A.13
- [496] TERMEER, M., BESCÓS, J. O., BREEUWER, M., VILANOVA, A., GERRITSEN, F., AND GRÖLLER, M. E. Covicad: Comprehensive visualization of coronary artery disease. *IEEE Trans Vis Comput Graph* 13, 6 (2007), 1632–1639, doi: 10.1109/TVCG.2007.70550. ■ ■ ■ . A.8.2, A.13
- [497] TERMEER, M., BESCÓS, J. O., BREEUWER, M., VILANOVA, A., GERRITSEN, F., GRÖLLER, M. E., AND NAGEL, E. Visualization of myocardial perfusion derived from coronary anatomy. *IEEE Trans Vis Comput Graph* 14, 6 (2008), 1595–1602, doi: 10.1109/TVCG.2008.180. ■ ■ ■ . A.8.2, A.13
- [498] TESTI, D., CLAPWORTHY, G., AYLWARD, S., FRANGI, A., AND CHRISTIE, R. Interactive visualization of multiscale biomedical data: an integrated approach. In *Proc BioVis* (2011), pp. 3–4. 1.1, 3.1, A.1
- [499] THOMASS, B., WALTER, J., KRONE, M., HASSE, H., AND ERTL, T. Interactive exploration of polymer-solvent interactions. In *Proc VMV* (2011), pp. 301–308. ■ ■ ■ . A.5.2, A.13
- [500] TODD, S., TODD, P., LEYMARIE, F. F., LATHAM, W., KELLEY, L. A., STERNBERG, M., HUGUES, J., AND TAYLOR, S. FoldSynth: interactive 2D/3D visualisation platform for molecular strands. In *Proc EG VCBM* (2015). ■ ■ ■ . A.5.1, A.13
- [501] TOMA, A., MANG, A., SCHÜTZ, T. A., BECKER, S., AND BUZUG, T. M. Is it necessary to model the matrix degrading enzymes for simulating tumour growth? In *Proc VMV* (2011), pp. 361–368. ■ ■ ■ . A.7.1, A.13
- [502] TOMA, A., RÉGNIER-VIGOUROUX, A., MANG, A., BECKER, S., SCHUETZ, T. A., AND BUZUG, T. M. In-silico modelling of tumour-immune system interactions for glioblastomas. *IFAC Proc Volumes* 45, 2 (2012), 1237–1242, doi: 10.3182/20120215-3-AT-3016.00219. ■ ■ ■ . A.7.1, A.13
- [503] TOMINSKI, C., AND SCHUMANN, H. *Interactive Visual Data Analysis*. AK Peters/CRC Press, 2020. 2.1, D.2
- [504] TORY, M. User studies in visualization: A reflection on methods. In *Handbook of Human Centric Visualization*. 2014, pp. 411–26. 3.3.1, D.4.3

- [505] TORY, M., AND MÖLLER, T. Evaluating visualizations: do expert reviews work? *IEEE Comput Graph Appl* 25, 5 (2005), 8–11. 2.3, D.2
- [506] TORY, M., ROBER, N., MÖLLER, T., CELLER, A., AND ATKINS, M. S. 4d space-time techniques: A medical imaging case study. In *Proc IEEE VIS* (2001), pp. 473–592. ■ . A.8, A.13
- [507] TRÖGER, J., HOISCHEN, C., PERNER, B., MONAJEMBASHI, S., BARBOTIN, A., LÖSCHBERGER, A., EGGELING, C., KESSELS, M. M., QUALMANN, B., AND HEMMERICH, P. Comparison of multiscale imaging methods for brain research. *Cells* 9, 6 (2020), 1377, doi: 10.3390/cells9061377. A.7
- [508] TRUONG, V., AND DUNCAN, N. W. Suggestions for improving the visualization of magnetic resonance spectroscopy voxels and spectra. *Royal Soc Open Sci* 7, 8 (2020), 200600, doi: 10.1098/rsos.200600. 2.2
- [509] TUFTE, E. R. *The visual display of quantitative information*. Graphics Press, 1986. 2.3, D.2
- [510] TURKAY, C., FILZMOSER, P., AND HAUSER, H. Brushing dimensions—a dual visual analysis model for high-dimensional data. *IEEE Trans Vis Comput Graph* 17, 12 (2011), 2591–2599, doi: 10.1109/TVCG.2011.178. 2.2, C.1, C.2.3
- [511] TURKAY, C., LUNDERVOLD, A., LUNDERVOLD, A., AND HAUSER, H. Hypothesis generation by interactive visual exploration of heterogeneous medical data. In *Human-Computer Interaction and Knowledge Discovery in Complex, Unstructured, Big Data* (2013), vol. 7947, pp. 1–12. 2.2, C.1
- [512] TUSZYNSKA, I., MAGNUS, M., JONAK, K., DAWSON, W., AND BUJNICKI, J. M. NPDock: a web server for protein–nucleic acid docking. *Nucleic Acids Res* 43, W1 (2015), W425–W430, doi: 10.1093/nar/gkv493. □ ■ □ □ . A.5.2, A.13
- [513] ULMER, S., BACKENS, M., AND AHLHELM, F. J. Basic principles and clinical applications of magnetic resonance spectroscopy in neuroradiology. *J Comput Assist Tomogr* 40, 1 (2016), 1–13, doi: 10.1097/RCT.0000000000000322. 2.2, 3.2.1, B.3
- [514] US DEPARTMENT OF AGRICULTURE, A. R. S. USDA national nutrient database for standard reference, release 28. <http://www.ars.usda.gov/nea/bhnrc/mafcl>, may 2016. C.3.2
- [515] VALM, A. M., COHEN, S., LEGANT, W. R., MELUNIS, J., HERSHBERG, U., WAIT, E., COHEN, A. R., DAVIDSON, M. W., BETZIG, E., AND LIPPINCOTT-SCHWARTZ, J. Applying systems-level spectral imaging and analysis to reveal the organelle interactome. *Nature* 546, 7656 (2017), 162–167, doi: 10.1038/nature22369. □ ■ □ □ . A.6.1, A.13
- [516] VAN DEN HEUVEL, M., MANDL, R., AND HULSHOFF POL, H. Normalized cut group clustering of resting-state fMRI data. *PLoS One* 3, 4 (2008), e2001, doi: 10.1371/journal.pone.0002001. ■ ■ □ □ . A.8.4, A.13
- [517] VAN DER GRAAF, M. In vivo magnetic resonance spectroscopy: Basic methodology and clinical applications. *Eur Biophys J* 39, 4 (2010), 527–540, doi: 10.1002/9780470882221. 2.2, 3.2.1, B.1
- [518] VAN LONG, T., AND LINSEN, L. Multiclustertree: interactive visual exploration of hierarchical clusters in multidimensional multivariate data. *Comput Graph Forum* 28, 3 (2009), 823–830, doi: 10.1111/j.1467-8659.2009.01468.x. 2.2, C.1
- [519] VAN PELT, R., BESCOS, J. O., BREEUWER, M., CLOUGH, R. E., GRÖLLER, TER HAAR ROMENY, B., AND VILANOVA, A. Interactive virtual probing of 4D MRI blood-flow. *IEEE Trans Vis Comput Graph* 17, 12 (2011), 2153–2162, doi: 10.1109/TVCG.2011.215. ■ ■ □ □ . A.8.1, A.13
- [520] VAN PELT, R., BESCOS, J. O., BREEUWER, M., CLOUGH, R. E., GRÖLLER, E., TER HAAR ROMENIJ, B., AND VILANOVA, A. Exploration of 4D MRI blood flow using stylistic visualization. *IEEE Trans Vis Comput Graph* 16, 6 (2010), 1339–1347, doi: 10.1109/TVCG.2010.153. ■ ■ □ □ . A.13
- [521] VAN PELT, R., AND VILANOVA, A. Understanding blood-flow dynamics: new challenges for visualization. *Computer* 46, 12 (2013), 60–67, doi: 10.1109/MC.2013.121. ■ . A.8.1, A.13

- [522] VAN PELT, R. F. P., FUSTER, A., CLAASSEN, G. G. H., AND VILANOVA, A. Characterization of blood-flow patterns from phase-contrast mri velocity fields. In *Proc EuroVis* (2014), pp. 61–65. ■ ■ ■ ■ . A.13
- [523] VAQUERO, R. M. M., RZEPECKI, J., FRIESE, K.-I., AND WOLTER, F.-E. Visualization and user interaction methods for multiscale biomedical data. In *3D Multiscale Physiological Human*. 2014, pp. 107–133. ■ . 2.1, A.1, A.13
- [524] VÁZQUEZ, P., HERMOSILLA, P., GUALLAR, V., ESTRADA, J., AND VINACUA, A. Visual analysis of protein-ligand interactions. In *Comput Graph Forum* (2018), vol. 37, pp. 391–402. □ ■ ■ □ . A.5.2, A.13
- [525] VERMA, A., KUMAR, I., VERMA, N., AGGARWAL, P., AND OJHA, R. Magnetic resonance spectroscopy - revisiting the biochemical and molecular milieu of brain tumors. *BBA Clinical* 5 (2016), 170–178, doi: 10.1016/j.bbacli.2016.04.002. 3.2.1, B.1
- [526] VERMATHEN, M., MÜLLER, J., FURRER, J., MÜLLER, N., AND VERMATHEN, P. 1H HR-MAS NMR spectroscopy to study the metabolome of the protozoan parasite giardia lamblia. *Talanta* 188 (2018), 429–441, doi: 10.1016/j.talanta.2018.06.006. B.8
- [527] VICECONTI, M., AND HUNTER, P. The virtual physiological human: ten years after. *Annu Rev Biomed Eng* 18 (2016), 103–123, doi: 10.1146/annurev-bioeng-110915-114742. 1.1, 3.1, A.1
- [528] VICKERY, O. N., AND STANSFELD, P. J. CG2AT2: an enhanced fragment-based approach for serial multi-scale molecular dynamics simulations. *bioRxiv* (2021), doi: 10.1021/acs.jctc.1c00295. □ ■ ■ □ . A.5.2, A.13
- [529] VILANOVA, A., PREIM, B., PELT, R. v., GASTEIGER, R., NEUGEBAUER, M., AND WISCHGOLL, T. *Visual Exploration of Simulated and Measured Blood Flow*. 2014, pp. 305–324. ■ . A.8.1, A.13
- [530] VIOLA, I., CHEN, M., AND ISENBERG, T. Visual abstraction. In *Foundations of Data Visualization*. pp. 15–37. 2.3, 3.3.1, 3.3.1, D.2, D.3, D.6
- [531] VIOLA, I., AND ISENBERG, T. Pondering the concept of abstraction in (illustrative) visualization. *IEEE Trans Vis Comput Graph* 24, 9 (2017), 2573–88. 2.3, D.2
- [532] VOLIOTIS, M., THOMAS, P., GRIMA, R., AND BOWSHER, C. G. Stochastic simulation of biomolecular networks in dynamic environments. *PLoS Comput Biol* 12, 6 (2016), e1004923, doi: 10.1371/journal.pcbi.1004923. A.6
- [533] VOSOUGH, Z., HOGRÄFER, M., ROYER, L. A., GROH, R., AND SCHULZ, H. J. Parallel hierarchies: A visualization for cross-tabulating hierarchical categories. *Comput Graph* 76 (2018), 1–17, doi: 10.1016/j.cag.2018.07.009. 2.2, C.1
- [534] WAIT, E., WINTER, M., BJORNSSON, C., KOKOVAY, E., WANG, Y., GODERIE, S., TEMPLE, S., AND COHEN, A. R. Visualization and correction of automated segmentation, tracking and lineaging from 5-d stem cell image sequences. *BMC Bioinform* 15, 1 (2014), 1–14, doi: 10.1186/1471-2105-15-328. □ ■ ■ □ . A.6.2, A.13
- [535] WALDIN, N., WALDNER, M., LE MUZIC, M., GRÖLLER, E., GOODSSELL, D. S., AUTIN, L., OLSON, A. J., AND VIOLA, I. Cuttlefish: color mapping for dynamic multi-scale visualizations. In *Comput Graph Forum* (2019), vol. 38, pp. 150–164. 2.1, 3.3.2, E.1, E.2, E.5
- [536] WALSH, C., TAFFOREAU, P., WAGNER, W., JAFREE, D., BELLIER, A., WERLEIN, C., KÜHNEL, M., BOLLER, E., WALKER-SAMUEL, S., ROBERTUS, J., ET AL. Imaging intact human organs with local resolution of cellular structures using hierarchical phase-contrast tomography. *Nat Methods* (2021), 1–10, doi: 10.1038/s41592-021-01317-x. 4, A.4, A.11
- [537] WALTEMATE, T., SOMMER, B., AND BOTSCH, M. Membrane mapping: combining mesoscopic and molecular cell visualization. In *Proc EG VCBM* (2014), pp. 89–96. □ ■ ■ □ . A.6.1, A.13

- [538] WALTON, S., BERGER, K., THIYAGALINGAM, J., DUFFY, B., FANG, H., HOLLOWAY, C., TREFETHEN, A. E., AND CHEN, M. Visualizing cardiovascular magnetic resonance (CMR) imagery: challenges and opportunities. *Prog Biophys Mol Biol* 115, 2-3 (2014), 349–358, doi: 10.1016/J.PBIOMOLBIO.2014.07.009. ■ . A.8.2, A.13
- [539] WAN, Y., AND HANSEN, C. Uncertainty footprint: Visualization of nonuniform behavior of iterative algorithms applied to 4d cell tracking. In *Comput Graph Forum* (2017), vol. 36, pp. 479–489. ■ ■ ■ . A.6.2, A.13
- [540] WANG, X. Z., MEDASANI, S., MARHOON, F., AND ALBAZZAZ, H. Multidimensional visualization of principal component scores for process historical data analysis. *Ind Eng Chem Res* 43, 22 (2004), 7036–7048, doi: 10.1021/ie030816j. 2.2, C.1
- [541] WANG, Y., ZHONG, Z., AND HUA, J. Deeporgannet: On-the-fly reconstruction and visualization of 3D/4D lung models from single-view projections by deep deformation network. *IEEE Trans Vis Comput Graph* 26, 1 (2019), 960–970, doi: 10.1109/TVCG.2019.2934369. ■ ■ ■ . A.8.3, A.13
- [542] WARD, M., LEBLANC, J. T., AND TIPNIS, R. N-land: a graphical tool for exploring n-dimensional data. In *Proc CGI* (1994), pp. 131–141. 3.2.1, B.6
- [543] WARD, M. O. Xmdvtool: integrating multiple methods for visualizing multivariate data. In *Proc IEEE VIS* (1994), pp. 326–333. 3.2.1, B.6
- [544] WASHINGTON, P., SAMUEL-GAMA, K. G., GOYAL, S., AND RIEDEL-KRUSE, I. H. Bioty: A cloud-based development toolkit for programming experiments and interactive applications with living cells. *bioRxiv* (2017), doi: 10.1073/pnas.1815367116. ■ ■ ■ . A.6.2, A.13
- [545] WAYLEN, L. N., NIM, H. T., MARTELOTTO, L. G., AND RAMIALISON, M. From whole-mount to single-cell spatial assessment of gene expression in 3D. *Commun Biol* 3, 1 (2020), 1–11, doi: 10.1038/s42003-020-01341-1. A.7
- [546] WEHI.TV. Explore wehi.tv. WEHI.TV, 2022. <https://www.wehi.edu.au/wehi-tv>. E.3.1
- [547] WEI, C.-Y., DIMITROVA, N., AND CHANG, S.-F. Color-mood analysis of films based on syntactic and psychological models. In *Proc IEEE ICME* (2004), vol. 2, pp. 831–834. 2.3, E.2
- [548] WEIDELE, D. K. I. Conditional parallel coordinates. *IEEE Trans Vis Comput Graph* 26 (2019), 221–225, doi: 10.1109/VISUAL.2019.8933632. 2.2, C.1
- [549] WEISKOPF, D. Vis4vis: visualization for (empirical) visualization research. In *Foundations of Data Visualization*. 2020, pp. 209–24. 2.3, D.2
- [550] WEISKOPF, N., VEIT, R., ERB, M., MATHIAK, K., GRODD, W., GOEBEL, R., AND BIRBAUMER, N. Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. *NeuroImage* 19, 3 (2003), 577–586, doi: 10.1016/S1053-8119(03)00145-9. ■ ■ ■ . A.8.4, A.13
- [551] WERNER, E. Strategies for the production of molecular animations. *Front Bioinform*, 48, doi: 10.3389/fbinf.2022.793914. 2.3
- [552] WESOLOWSKA, N., AVILOV, I., MACHADO, P., GEISS, C., KONDO, H., MORI, M., AND LENART, P. Actin assembly ruptures the nuclear envelope by prying the lamina away from nuclear pores and nuclear membranes in starfish oocytes. *Elife* 9 (2020), e49774, doi: 10.7554/eLife.49774. ■ ■ ■ . A.6.1, A.13
- [553] WEXNER, L. B. The degree to which colors (hues) are associated with mood-tones. *J Appl Psychol* 38 (1954), 432–435. 2.3, E.2
- [554] WHITE, I. R., ROYSTON, P., AND WOOD, A. M. Multiple imputation using chained equations: issues and guidance for practice. *Stat Med* 30, 4 (2011), 377–399, doi: 10.1002/sim.4067. 14, 14, C.2.4

- [555] WIECHERT, L., COMERFORD, A., RAUSCH, S., AND WALL, W. A. Advanced multi-scale modelling of the respiratory system. In *Fundamental Medical and Engineering Investigations on Protective Artificial Respiration*. 2011, pp. 1–32. ■ ■ ■ □ □ . A.8.3, A.13
- [556] WILLIAMSON, D. F., PARKER, R. A., AND KENDRICK, J. S. The box plot: a simple visual method to interpret data. *Ann Intern Med* 110, 11 (1989), 916–921, doi: 10.7326/0003-4819-110-11-916. 3.2.1, B.6
- [557] WILSON, M., REYNOLDS, G., KAUPPINEN, R. A., ARVANITIS, T. N., AND PEET, A. A constrained least-squares approach to the automated quantitation of in vivo 1H magnetic resonance spectroscopy data. *Magn Reson Med* 65, 1 (2011), 1–12, doi: 10.1002/mrm.22579. B.1, B.3, B.7
- [558] WITTINK, D. R., AND BAYER, L. R. The measurement imperative. *J Mark Res* 6, 4 (Fall 1994), 14. 3.2.1
- [559] WOLF, I., VETTER, M., WEGNER, I., BÖTTGER, T., NOLDEN, M., SCHÖBINGER, M., HASTEN-TEUFEL, M., KUNERT, T., AND MEINZER, H. P. The medical imaging interaction toolkit. *Med Image Anal* 9, 6 (2005), 594–604, doi: 10.1016/j.media.2005.04.005. 2.2, B.2
- [560] WOLTER, M., ASSENMACHER, I., HENTSCHEL, B., SCHIRSKI, M., AND KUHLEN, T. A time model for time-varying visualization. *Comput Graph Forum* 28, 6 (2009), 1561–1571, doi: 10.1111/j.1467-8659.2008.01314.x. ■ ■ ■ □ □ . A.13
- [561] WONG, B. Points of view: color coding. *Nat Methods* 7, 8 (2010), 573, doi: 10.1038/nmeth0810-573. 2.3, 3.3.2, E.2, E.4
- [562] WONG, K. C., WANG, L., ZHANG, H., AND SHI, P. Physiological fusion of functional and structural data for cardiac deformation recovery. In *Proc MICCAI* (2010). ■ ■ ■ □ □ . A.8.2, A.13
- [563] WORTEL, I. M., AND TEXTOR, J. Artistoo, a library to build, share, and explore simulations of cells and tissues in the web browser. *Elife* 10 (2021), e61288, doi: 10.7554/eLife.61288. ■ ■ ■ □ □ . A.7.1, A.13
- [564] WU, G., WANG, Q., LIAN, J., AND SHEN, D. Estimating the 4d respiratory lung motion by spatiotemporal registration and super-resolution image reconstruction. *Med Phys* 40, 3 (2013), 031710–1–031710–17, doi: 10.1118/1.4790689. ■ ■ ■ □ □ . A.8.3, A.13
- [565] WU, K., CHEN, J., PRUETT, W. A., AND HESTER, R. L. HumMod browser: An exploratory visualization tool for the analysis of whole-body physiology simulation data. In *Proc BioVis* (2013), pp. 97–104. ■ ■ ■ □ □ . A.9, A.13
- [566] XIONG, G., SUN, P., ZHOU, H., HA, S., Ó HARTAIGH, B., TRUONG, Q. A., AND MIN, J. K. Comprehensive modeling and visualization of cardiac anatomy and physiology from CT imaging and computer simulations. *IEEE Trans Vis Comput Graph* 23, 2 (2016), 1014–1028, doi: 10.1109/TVCG.2016.2520946. ■ ■ ■ □ □ . A.8.2, A.13
- [567] XIONG, G., AND TAYLOR, C. A. Physics-based modeling of aortic wall motion from ecg-gated 4d computed tomography. In *Proc MICCAI* (2010), pp. 426–434. ■ ■ ■ □ □ . A.8, A.13
- [568] XU, D., AND VIGNERON, D. Magnetic resonance spectroscopy imaging of the newborn brain—a technical review. *Semin Perinatol* 34, 1 (2010), 20–27, doi: 10.1053/j.semperi.2009.10.003. 3.2.1, B.6, B.6
- [569] XU, F., MA, D., MACPHERSON, K. P., LIU, S., BU, Y., WANG, Y., TANG, Y., BI, C., KWOK, T., CHUBYKIN, A. A., ET AL. Three-dimensional nanoscopy of whole cells and tissues with in situ point spread function retrieval. *Nat Methods* 17, 5 (2020), 531–540, doi: 10.1038/s41592-020-0816-x. ■ ■ ■ □ □ . A.6.1, A.13
- [570] XVIVO. Xvivo services. XVIVO Scientific & Medical Animation, 2022. <https://xvivo.com/scientific-animation-examples/>. E.3.1
- [571] YANG, J., PENG, W., WARD, M., AND RUNDENSTEINER, E. Interactive hierarchical dimension ordering, spacing and filtering for exploration of high dimensional datasets. In *Proc IEEE InfoVis* (2003), pp. 105–112. 2.2, C.1

- [572] YANG, J., WARD, M., AND RUNDENSTEINER, E. InterRing: an interactive tool for visually navigating and manipulating hierarchical structures. In *Proc IEEE InfoVis (2002)*, pp. 77–84. 2.2, C.1
- [573] YANG, Y., GUO, X., VICK, J., TORRES, L. G., AND CAMPBELL, T. F. Physics-based deformable tongue visualization. *IEEE Trans Vis Comput Graph* 19, 5 (2013), 811–823, doi: 10.1109/TVCG.2012.174. ■
■□□. A.8.5, A.13
- [574] YARDIMCI, G. G., AND NOBLE, W. S. Software tools for visualizing hi-c data. *Genome Biol* 18, 1 (2017), 1–9, doi: 10.1101/086017v1. ■. A.5, A.13
- [575] YEOMANS, K. A., AND GOLDBER, P. A. The guttmann-kaiser criterion as a predictor of the number of common factors. *J R Stat Soc Series D* (1982), 221–229, doi: 10.2307/2987988. 14, 14, C.2.2
- [576] YUAN, X., REN, D., WANG, Z., AND GUO, C. Dimension projection matrix/tree: Interactive subspace visual exploration and analysis of high dimensional data. *IEEE Trans Vis Comput Graph* 19, 12 (2013), 2625–2633, doi: 10.1109/TVCG.2013.150. 2.2, C.1
- [577] ZEILEIS, A., HORNIK, K., AND MURRELL, P. Escaping rgbland: selecting colors for statistical graphics. *Comput Stat Data Anal* 53, 9 (2009), 3259–3270, doi: 10.1016/j.csda.2008.11.033. E.5
- [578] ZHANG, Y., LIU, T., WANG, J., ZOU, B., LI, L., YAO, L., CHEN, K., NING, L., WU, B., ZHAO, X., ET AL. Cellinker: a platform of ligand–receptor interactions for intercellular communication analysis. *Bioinformatics* 37, 14 (2021), 2025–2032, doi: 10.1093/bioinformatics/btab036. □ ■□□. A.5.2, A.13
- [579] ZHOU, L., AND HANSEN, C. D. A survey of colormaps in visualization. *IEEE Trans Vis Comput Graph* 22, 8 (2015), 2051–2069, doi: 10.1109/TVCG.2015.2489649. 2.1, E.2
- [580] ZIEMKIEWICZ, C., CHEN, M., LAIDLAW, D. H., PREIM, B., AND WEISKOPF, D. Open challenges in empirical visualization research. In *Foundations of Data Visualization*. 2020, pp. 243–52. 2.3, D.2

Appendix: Paper A

Web of Science Hot and Highly Cited Papers

The two tables on the following pages provide an overview of highly interesting application domain research in the area of physiology, and which provide an indication of the opportunities for visualization research, which we used to drive in part our literature search in PaperA. We used Web of Science's "hot papers" and "highly cited" filters with the keyword "physiology," where a "hot paper" is any paper published in the past two years that has received enough citations to rank in the top 0.1% of papers in its field. A "highly cited" paper ranks in the top 1% of cited papers for its field and publication year.

We took the top 20 papers from each of these filters and excluded works that did not relate to humans or other mammals. Papers are keyed to an area of physiology, e.g., molecular pathways or heart function, following standard medical physiology textbooks [180].

Table A1: Web of Science **Hot Papers**, Top 20 (Non-mammal works excluded)

Rank	Title	Authors	Year	Publication Venue	Topic Area
1	COVID-19 and its implications for thrombosis and anticoagulation	Connors & Levy	2020	Blood	Blood Flow
2	Reactive oxygen species (ROS) as pleiotropic physiological signalling agents	Sies & Jones	2020	Nature Reviews Molecular Cell Biology	Molecular interactions
3	mTOR at the nexus of nutrition, growth, ageing and disease	Liu & Sabatini	2020	Nature Reviews Molecular Cell Biology	Molecular pathways
4	Metabolic regulation of gene expression by histone lactylation	Zhang et al.	2019	Nature	Molecular dynamics, Molecular pathways
5	Angiotensin-converting enzyme 2 (ACE2), SARS-CoV-2 and the pathophysiology of coronavirus disease 2019 (COVID-19)	Bourgonje et al.	2021	Journal of Pathology	Molecular pathways
6	The PI3K-AKT network at the interface of oncogenic signalling and cancer metabolism	Hoxhaj & Manning	2020	Nature Reviews Cancer	Molecular pathways
7	Mitochondrial TCA cycle metabolites control physiology and disease	Martinez-Reyes & Chandel	2020	Nature Communications	Cell dynamics
8	The role of short-chain fatty acids from gut microbiota in gut-brain communication	Silva et al.	2020	Frontiers in Endocrinology	Molecular pathways, *Organism
9	The brain's default network: updated anatomy, physiology and evolving insights	Buckner & DiNicola	2019	Nature Reviews Neuroscience	Brain function
10	Diet-microbiota interactions and personalized nutrition	Kolodziejczyk et al.	2019	Nature Reviews Microbiology	Cell interactions, *Organism
11	Identification of region-specific astrocyte subtypes at single cell resolution	Batiuk et al.	2020	Nature Communications	Cell interactions
12	Ferroptosis: mechanisms, biology and role in disease	Jiang et al.	2021	Nature Reviews Molecular Cell Biology	Cell dynamics
13	Organs-on-chips: into the next decade	Low et al.	2021	Nature Reviews Drug Discovery	Tissue interactions
14	Perivascular spaces in the brain: anatomy, physiology and pathology	Wardlaw et al.	2020	Nature Reviews Neurol.	Brain function
15	Social determinants of health and survival in humans and other animals	Snyder-Mackler et al.	2020	Science	*Organism, Population physio.
16	3-month, 6-month, 9-month, and 12-month respiratory outcomes in patients following COVID-19-related hospitalisation: a prospective study	Wu et al.	2021	Lancet Respiratory Medicine	Lung function
17	Microbial Regulation of Host Physiology by Short-chain Fatty Acids	van der Hee & Wells	2021	Trends in Microbiology	Molecular pathways
18	A human-airway-on-a-chip for the rapid identification of candidate antiviral therapeutics and prophylactics	Si et al.	2021	Nature Biomedical Engineering	Tissue dynamics
19	When the human brain goes diving: using near-infrared spectroscopy to measure cerebral and systemic cardiovascular responses to deep, breath-hold diving in elite freedivers	McKnight et al.	2021	Philosophical Transactions of the Royal Society B: Biological Sciences	Heart function, Brain function

Table A2: Web of Science **Highly Cited Papers**, Top 20 (Non-mammal works excluded)

Rank	Title	Authors	Year	Publication Venue	Topic Area
1	2013 ACCF/AHA Guideline for the Management of Heart Failure	Yancy et al.	2013	Journal of the American College of Cardiology.	Heart function, Blood Flow
2	2014 AHA/ACC/HRS Guideline for the Management of Patients With Atrial Fibrillation	January et al.	2014	Circulation	Heart function
3	2014 AHA/ACC/HRS Guideline for the Management of Patients With Atrial Fibrillation	January et al.	2014	Journal of the American College of Cardiology.	Heart function
4	mTOR Signaling in Growth, Metabolism, and Disease	Saxton & Sabatini	2017	Cell	Molecular pathways
5	Diversity, stability and resilience of the human gut microbiota	Lozupone et al.	2012	Nature	Cell interactions
6	Interactions Between the Microbiota and the Immune System	Hooper et al.	2012	Science	Cell interactions, *Organism
7	Macrophage biology in development, homeostasis and disease	Wynn et al.	2013	Nature	Cell dynamics
8	Physiology of Microglia	Kettenmann et al.	2011	Physiological Reviews	Cell dynamics, interactions
9	Mind-altering microorganisms: the impact of the gut microbiota on brain and behaviour	Cryan & Dinan	2012	Nature Reviews Neuroscience	Cell interactions, *Organism
10	Lipid Peroxidation: Production, Metabolism, and Signaling Mechanisms of Malondialdehyde and 4-Hydroxy-2-Nonenal	Ayala et al.	2014	Oxidative Medicine & Cellular Longevity	Molecular interactions
11	Control of apoptosis by the BCL-2 protein family: implications for physiology and therapy	Czabotar et al.	2014	Nature Reviews Molec. Cell Biology	Cell dynamics
12	From Dietary Fiber to Host Physiology: Short-Chain Fatty Acids as Key Bacterial Metabolites	Koh et al.	2016	Cell	Molecular pathways, *Organism
13	Hypoxia-Inducible Factors in Physiology and Medicine	Semenza	2012	Cell	Molecular pathways
14	Prone Positioning in Severe Acute Respiratory Distress Syndrome	Guerin et al.	2013	New England Journal of Medicine	Lung function
15	Normal gut microbiota modulates brain development and behavior	Heijtza et al.	2011	Proceedings of the National Academy of Sciences of the USA	Cell interactions, Brain function
16	The gut microbiota - masters of host development and physiology	Sommer & Baekhed	2013	Nature Reviews Microbiology	Cell interactions
17	The landscape of long noncoding RNAs in the human transcriptome	Iyer et al.	2015	Nature Genetics	Molecular pathways
18	Microfluidic organs-on-chips	Bhatia & Ingber	2014	Nature Biotechnology	Tissue dynamics
19	Proteomic comparison defines novel markers to characterize heterogeneous populations of extracellular vesicle subtypes	Kowal et al.	2016	Proceedings of the National Academy of Sciences of the USA	Cell interactions
20	Pro-resolving lipid mediators are leads for resolution physiology	Serhan	2014	Nature	Molecular pathways

Appendix: Paper D

Survey Structure

The following describes in detail the structure of our survey that we conducted in PaperD. Screenshots of all pages of the three surveys are available at https://github.com/lauragarrison87/Biomedical_Process_Vis/tree/main/5-survey-screenshots.

1. Welcome screen/GDPR information

Welcome! In this survey we will ask you to compare the different ways that we represent two physiological phenomena: [normal topic] and [pathological topic].

We will present you with two task-based scenarios, (1) expert and (2) general audience, and ask you to identify the (1) most and (2) least effective visual representations for each scenario's communication objective for both blood flow and an aneurysm, respectively. We will furthermore ask you to provide some basic ratings and keywords for your choices. We will use these data to analyze and develop a set of recommendations for more effective physiology representation methods driven by audience and task. Following the conclusion of this study we plan to make these results and assets available to the broader research community.

In compliance with GDPR regulations in the EU, we will not collect or store any personally identifying information about you unless you choose to provide it. We will keep private any personal information that you provide (i.e., profession, age, gender, email). We will use this optional information to identify possible demographic-related patterns in perception; provision of your email is only if you are willing for us to contact you if we have a question for you on any of your survey responses. Your data will be stored on secure servers with multi-factor authentication through Typeform (for more information, visit <https://help.typeform.com/hc/en-us/articles/360029259552-Security-at-Typeform> and will be moved to a secure server at University of Bergen following the completion of this study.

We will only use your responses for the purposes of this research and our stated aims. By proceeding into this survey, you indicate your agreement for our use of your responses for this purpose. If you have any further questions, would like to view, change, or delete your responses, you may contact us (email laura.garrison@uib.no) at any time.

This survey should take around **30 minutes**. Thank you in advance for contributing your time and feedback to this research.

2. User profile

First, please tell us a bit about yourself.

- (a) Please briefly describe your profession or background. For example, "Clinical researcher in immunology, 10 years in the field."

- (b) What gender do you most identify with?
- i. Male
 - ii. Female
 - iii. Prefer not to say
- (c) Please tell us your age using the following categories:
- i. < 24 years old
 - ii. 25-34 years old
 - iii. 35-44 years old
 - iv. 45-54 years old
 - v. 55-64 years old
 - vi. 65-74 years old
 - vii. 75 years or older
 - viii. Prefer not to say
3. How would you rate your expertise regarding [normal topic], relative to the general population?
- (a) **0** - no knowledge
 - (b) **1** - some knowledge, I've heard of this
 - (c) **2** - I enjoy reading/hearing occasionally about this topic
 - (d) **3** - I know this area relatively well
 - (e) **4** - I know this area extremely well
4. **Expert audience scenario – normal topic**
 Consider the following scenario and learning objective describing [normal topic] for an expert audience: [scenario ...]
- (a) In your opinion, which of the following assets **BEST** visually describes and support this scenario and communication objective?
 - (b) What is your **second** choice for this scenario and communication objective?
 - (c) What is your **third** choice for this scenario and communication objective?
 - (d) Now we will ask you to provide a few ratings of only your **top** choice.
NB: Rating is done according to the following scale (these are heavily referencing the AMI salon judging guidelines).
 - 1 - (poor/minimal)
 - 2 - (fair/low)
 - 3 - (average/adequate)
 - 4 - (very good/high)
 - 5 - (excellent/very high)
 - i. Rate your **top** choice for aesthetics (i.e., How visually appealing do you consider this to be?).
 - ii. Rate your **top** choice for visual clarity (i.e., How well does it convey the described scenario?).

- iii. Rate your **top** choice for scientific accuracy.
 - iv. Rate your **top** choice for communication success (i.e., How well does it meet [scenario]’s stated communication/learning objective?).
 - (e) Use the below keywords to describe the strengths of your **top** choice. Choose as many as you feel are relevant:
Accurate, Clear, Detailed, Direct, Easy to read, Informative, Precise, Pretty, Visually pleasing
 - (f) Use the below keywords to describe the weaknesses of your **top** choice. Choose as many as you feel are relevant:
Confusing, Distracting, Excessive, Inaccurate, Intricate, Misleading, Simplistic, Visually unappealing
 - (g) If you have further comments on the options or your **top** choices for this expert scenario, please write them here. Otherwise, simply write *N*.
5. Consider again the same scenario and communication objective in describing [topic] for an **expert** audience. In your opinion, which of the following assets **LEAST** visually describes and supports this scenario and communication objective (i.e., your bottom choice)?

(a) *Repeat questions 4a–g for expert audience bottom choice*

6. **General audience scenario – normal topic**

Now consider the following scenario and communication objective in describing [topic] for a general audience. [scenario . . .]

In your opinion, which of the following assets **BEST** visually describes and supports this scenario and communication objective?

(a) *Repeat questions 4a–g for general audience top choice*

7. Consider again the same scenario and communication objective in describing [topic] for a general audience. In your opinion, which of the following assets **LEAST** visually describes and supports this scenario and communication objective (i.e., your bottom choice)?

(a) *Repeat questions 4a–g for general audience bottom choice*

8. How would you rate your expertise regarding [pathological topic], relative to the general population?

- (a) **0** - no knowledge
- (b) **1** - some knowledge, I’ve heard of this
- (c) **2** - I enjoy reading/hearing occasionally about this topic
- (d) **3** - I know this area relatively well
- (e) **4** - I know this area extremely well

9. **Expert audience scenario – pathological topic**

Consider the following scenario and learning objective describing [pathological topic] for an expert audience: [scenario ...]

(a) *Repeat questions 4a–g for expert audience top choice*

10. Consider again the same scenario and communication objective in describing [pathological topic] for an expert audience. In your opinion, which of the following assets **LEAST** visually describes and supports this scenario and communication objective (i.e., your bottom choice)?

(a) *Repeat questions 4a–g for expert audience bottom choice*

11. **General audience scenario – pathological topic**

Consider the following scenario and learning objective describing [pathological topic] for a general audience:[scenario ...]

(a) *Repeat questions 4a–g for general audience top choice*


12. Consider again the same scenario and communication objective in describing [pathological topic] for a general audience. In your opinion, which of the following assets **LEAST** visually describes and supports this scenario and communication objective (i.e., your bottom choice)?

(a) *Repeat questions 4a–g for general audience bottom choice*

13. Thank you for your participation! If you have questions or are interested in updates on this work please feel free to reach out to laura.garrison@uib.no



Graphic design: Communication Division, UIB / Print: Skjipes Kommunikasjon AS



uib.no

ISBN: 9788230866597 (print)
9788230841389 (PDF)