Scaling Up Medical Visualization

Multi-Modal, Multi-Patient, and Multi-Audience Approaches for Medical Data Exploration, Analysis and Communication

Eric Mörth

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

UNIVERSITY OF BERGEN

Scaling Up Medical Visualization

Multi-Modal, Multi-Patient, and Multi-Audience Approaches for Medical Data Exploration, Analysis and Communication

Eric Mörth

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

Date of defense: 02.09.2022

Year: 2022
Title: Scaling Up Medical Visualization

Name: Eric Mörth
Print: Skipnes Kommunikasjon / University of Bergen

Scientific Environment

The work presented in this thesis was conducted as part of my Ph.D. studies at the Department of Informatics, University of Bergen. In addition, I have been enrolled in the ICT Research School at the Department of Informatics, University of Bergen. Parts of my work have been done in the context of the Mohn Medical Imaging and Visualization Centre (MMIV) at the Department of Radiology at Haukeland University Hospital. Furthermore, I conducted two research stays at the Visualization Group of the Technical University of Vienna. My research was supported via my principal supervisor Noeska N. Smit by the Trond Mohn Foundation (TMS) Grant Number: 811255.

UNIVERSITETET I BERGEN

ICT Research School In
Information and Communication Technology

MMIV

TROND
MOHN
STIFTELSE

TU
WIEN
TECHNISCHE
UNIVERSITÄT
WIEN
Vienna University of Technology

Acknowledgements

First, I would like to thank my parents, Gerlinde and Heinrich Mörth, my girlfriend Sunisa Thongkham, brother Stefan Mörth, grandparents, and the rest of my family for their support. Without them, I would have most definitely not had the courage to apply for the Ph.D. position I am currently finishing. Thank you very much for supporting me throughout my studying time in Vienna and now in Bergen. You made this possible!

I want to thank Noeska Smit for choosing me as the first Ph.D. candidate to supervise and for allowing me to conduct the Ph.D. research of my dreams. Noeska is one of the most engaged professors I have met throughout my studying career. We always had encouraging discussions, pushed through late-night conference deadlines, and polished our papers until they could not shine anymore. I felt like I was in a very fortunate position to not care too much about funding for travel, equipment, or other expenses, which made the Ph.D. experience a great pleasure (except the virus crossing our traveling plans for most of the Ph.D. time). You made this Ph.D. time a pleasure, and I would like to thank you for our great time. I would also like to thank Stefan Bruckner, my co-supervisor, who always had an open ear for me. Stefan contributed significantly to the success of this thesis. I enjoyed our meetings and always felt like being in good hands with my supervisor team. Furthermore, I would like to thank Helwig Hauser for the fruitful discussions and the joyful time at the vis group in Bergen.

My appreciation also extends to my colleagues in the Bergen Visualization Group: Firstly, Thomas Trautner, with whom I spent many long fjellskiing trips, action-packed holidays in the north, and a fair share of my free time in Bergen. Secondly, Yngve Sekse Kriastansen invited us to experience the Norwegian lifestyle, including various fishing trips with pristine cuisine, teaching us how to fillet, salt and dry, and finally smoke fish. Takk skal du ha! Furthermore, I would like to thank Sergej Stoppel, Juraj Pálenik, Sherin Sugathan, Laura Garrison and Oli, Fabian Bolte, Forough Gharbalchi, Jan Byška, and Veronika Šoltészová. Furthermore, I would like to thank Tanja Eichner, Nastasja Steinhauer, Christian Hein, Mathias Bøe, Arthur Andersen, and Evgenia Stupak for trusting me as their project or Master's thesis (co-)supervisor.

I would also like to thank all my colleagues and, foremost also, the co-authors of several of my papers at the Mohn Medical Imaging and Visualization group. Without your open ears, motivation, and engagement, many of our ideas would not have worked out and would not have resulted in the papers we are proud to share here. Thanks to Ingfrid S. Haldorsen, Erlend Hodneland, Kari Wagner-Larsen, Camilla Krakstad, Njål Gjærde Lura, Heidi Espedal and Sathiesh Kaliyugarasan and to all the other collegues at MMIV and at the Kvinneklinikken of the Haukeland University Hospital Bergen, Norway.

Furthermore, I would like to thank the Visualization Group at the TU Vienna for inviting and hosting two short research stays at their group. Thank you very much,

Eduard M. Gröller, for organizing my stays and your valuable input to my scientific work and discussions about Ph.D. life and how to plan an uncertain future. Thank you, Manuela Waldner, Renata Raidou, Ivan Viola, Max Höfferer, Daniel Pahr, Henry Ehlers, Áron Samuel Kovács and Maath Musle. I want to thank Hannes Beiglböck from the Medical University of Vienna and his colleagues for the fruitful collaboration.

Special thanks also go to my friends back in Austria, whom I always look forward to visiting when being there. "Last but not least, I wanna thank me! I wanna thank me for believing in me, I wanna thank me for doing all this hard work, I wanna thank me for having no days off, I wanna thank me for never quitting ..." - Snoop Dogg. It is essential to look back and thank yourself for pulling through and putting in everything needed to fulfill your dreams.

Abstract in English

Medical visualization is one of the most application-oriented areas of visualization research. Close collaboration with medical experts is essential for interpreting medical imaging data and creating meaningful visualization techniques and visualization applications. Cancer is one of the most common causes of death, and with increasing average age in developed countries, gynecological malignancy case numbers are rising. Modern imaging techniques are an essential tool in assessing tumors and produce an increasing number of imaging data radiologists must interpret. Besides the number of imaging modalities, the number of patients is also rising, leading to visualization solutions that must be scaled up to address the rising complexity of multi-modal and multi-patient data. Furthermore, medical visualization is not only targeted toward medical professionals but also has the goal of informing patients, relatives, and the public about the risks of certain diseases and potential treatments. Therefore, we identify the need to scale medical visualization solutions to cope with multi-audience data.

This thesis addresses the scaling of these dimensions in different contributions we made. First, we present our techniques to scale medical visualizations in multiple modalities. We introduced a visualization technique using small multiples to display the data of multiple modalities within one imaging slice. This allows radiologists to explore the data efficiently without having several juxtaposed windows. In the next step, we developed an analysis platform using radiomic tumor profiling on multiple imaging modalities to analyze cohort data and to find new imaging biomarkers. Imaging biomarkers are indicators based on imaging data that predict clinical outcome related variables. Radiomic tumor profiling is a technique that generates potential imaging biomarkers based on first and second-order statistical measurements. The application allows medical experts to analyze the multi-parametric imaging data to find potential correlations between clinical parameters and the radiomic tumor profiling data. This approach scales up in two dimensions, multi-modal and multi-patient. In a later version, we added features to scale the multi-audience dimension by making our application applicable to cervical and prostate cancer data and the endometrial cancer data the application was designed for. In a subsequent contribution, we focus on tumor data on another scale and enable the analysis of tumor sub-parts by using the multi-modal imaging data in a hierarchical clustering approach. Our application finds potentially interesting regions that could inform future treatment decisions. In another contribution, the digital probing interaction, we focus on multi-patient data. The imaging data of multiple patients can be compared to find interesting tumor patterns potentially linked to the aggressiveness of the tumors. Lastly, we scale the multi-audience dimension with our similarity visualization applicable to endometrial cancer research, neurological cancer imaging research, and machine learning research on the automatic segmentation of tumor data. In contrast to the previously highlighted contributions, our last contribution, ScrollyVis, focuses primarily on multi-audience communication. We enable the creation of dynamic scientific scrollytelling experiences for a specific or general audience. Such stories can be used for specific use cases such as patient-doctor communication or communicating scientific results via stories targeting the general audience in a digital museum exhibition.

Our proposed applications and interaction techniques have been demonstrated in application use cases and evaluated with domain experts and focus groups. As a result, we brought some of our contributions to usage in practice at other research institutes. We want to evaluate their impact on other scientific fields and the general public in future work.

Abstract in Norwegian

Medisinsk visualisering er en av de mest applikasjonsrettede områdene av visualiseringsforsking. Tett samarbeid med medisinske eksperter er nødvendig for å tolke medisinsk bildedata og lage betydningsfulle visualiseringsteknikker og visualiseringsapplikasjoner. Kreft er en av de vanligste dødsårsakene, og med økende gjennomsnittsalder i i-land øker også antallet diagnoser av gynekologisk kreft. Moderne avbildningsteknikker er et viktig verktøy for å vurdere svulster og produsere et økende antall bildedata som radiologer må tolke. I tillegg til antallet bildemodaliteter, øker også antallet pasienter, noe som fører til at visualiseringsløsninger må bli skalert opp for å adressere den økende kompleksiteten av multimodal- og multipasientdata. Dessuten er ikke medisinsk visualisering kun tiltenkt medisinsk personale, men har også som mål å informere pasienter, pårørende, og offentligheten om risikoen relatert til visse sykdommer, og mulige behandlinger. Derfor har vi identifisert behovet for å skalere opp medisinske visualiseringsløsninger for å kunne håndtere multipublikumdata.

Denne avhandlingen adresserer skaleringen av disse dimensjonene i forskjellige bidrag vi har kommet med. Først presenterer vi teknikkene våre for å skalere visualiseringer i flere modaliteter. Vi introduserer en visualiseringsteknikk som tar i bruk små multipler for å vise data fra flere modaliteter innenfor et bildesnitt. Dette lar radiologer utforske dataen effektivt uten å måtte bruke flere sidestilte vinduer. I det neste steget utviklet vi en analyseplatform ved å ta i bruk «radiomic tumor profiling» på forskjellige bildemodaliteter for å analysere kohortdata og finne nye biomarkører fra bilder. Biomarkører fra bilder er indikatorer basert på bildedata som kan forutsi variabler relatert til kliniske utfall. «Radiomic tumor profiling» er en teknikk som genererer mulige biomarkører fra bilder basert på første- og andregrads statistiske målinger. Applikasjonen lar medisinske eksperter analysere multiparametrisk bildedata for å finne mulige korrelasjoner mellom kliniske parameter og data fra «radiomic tumor profiling». Denne tilnærmingen skalerer i to dimensjoner, multimodal og multipasient. I en senere versjon la vi til funksjonalitet for å skalere multipublikumdimensjonen ved å gjøre applikasjonen vår anvendelig for livmorhalskreft- og prostatakreftdata, i tillegg til livmorkreftdataen som applikasjonen var designet for. I et senere bidrag fokuserer vi på svulstdata på en annen skala og muliggjør analysen av svulstdeler ved å bruke multimodal bildedata i en tilnærming basert på hierarkisk gruppering. Applikasjonen vår finner mulige interessante regioner som kan informere fremtidige behandlingsavgjørelser. I et annet bidrag, en digital sonderingsinteraksjon, fokuserer vi på multipasientdata. Bildedata fra flere pasienter kan sammenlignes for å finne interessante mønster i svulstene som kan være knyttet til hvor aggressive svulstene er. Til slutt skalerer vi multipublikumdimensjonen med en likhetsvisualisering som er anvendelig for forskning på livmorkreft, på bilder av nevrologisk kreft, og maskinlæringsforskning på automatisk segmentering av svulstdata. Som en kontrast til de allerede fremhevete bidragene, fokuserer vårt siste bidrag, ScrollyVis, hovedsakelig på multipublikumkommunikasjon. Vi muliggjør skapelsen av dynamiske og vitenskapelige “scrollytelling”-opplevelser for spesifikke eller generelle publikum. Slike historien kan bli brukt i spesifikke brukstilfeller som kommunikasjon mellom lege og pasient, eller for å kommunisere vitenskapelige resultater via historier til et generelt publikum i en digital museumsutstilling.

Våre foreslåtte applikasjoner og interaksjonsteknikker har blitt demonstrert i brukstilfeller og evaluert med domeneeksperter og fokusgrupper. Dette har ført til at noen av våre bidrag allerede er i bruk på andre forskingsinstitusjoner. Vi ønsker å evaluere innvirkningen deres på andre vitenskapelige felt og offentligheten i fremtidige arbeid.

List of Papers

This thesis is based on the following publications:

(A) **Eric Mörth**, Ingfrid S. Haldorsen, Stefan Bruckner and Noeska N. Smit. **ParaGlyder: Probe-driven Interactive Visual Analysis for Multiparametric Medical Imaging Data**. In *Magnenat-Thalmann N. et al. (eds) Advances in Computer Graphics.* (CGI, 2020), Lecture Notes in Computer Science, vol 12221. DOI: 10.1007/978-3-030-61864-3_29

The paper has been awarded the **Best Interdisciplinary Presentation Award** at the 37th International Conference on Computer Graphics on 20nd – 23rd of October 2020 in Geneva (online).

(B) **Eric Mörth**, Kari Wagner-Larsen, Erlend Hodneland, Camilla Krakstad, Ingfrid S. Haldorsen, Stefan Bruckner and Noeska N. Smit **RadEx: Integrated Visual Exploration of Multiparametric Studies for Radiomic Tumor Profiling**. In *Computer Graphics Forum* (PG 2020), 39: 611-622, DOI: 10.1111/cgf.14172

(C) **Eric Mörth**, Tanja Eichner, Ingfrid S. Haldorsen, Stefan Bruckner and Noeska N. Smit **Interactive Hierarchical Clustering for Multiparametric Cancer Imaging**. In *VINCI 2022 - Short Papers - In Submission*

(D) **Eric Mörth**, Stefan Bruckner and Noeska N. Smit **ScrollyVis: Guided dynamic narratives for scientific storytelling**. In *IEEE Transactions on Visualizations and Computer Graphics (TVCG) - In Submission*

The following papers are also related to this thesis:

(I) **Eric Mörth**, Renata Raidou, Ivan Viola and Noeska N. Smit **The Vitruvian baby: Interactive reformation of fetal ultrasound data to a T-position**. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (VCBM 2019), DOI: 10.2312/VCBM.20191245

(II) H. Beiglböck, **E. Mörth**, B. Reichardt, T. Stamm, B. Itariu, J. Harreiter, M. Hufgard-Leitner, P. Fellinger, J. Eichelter, G. Prager, A. Kautzky, A. Kautzky-Willer, P. Wolf, and M. Krebs **Sex-Specific Differences in Mortality of Patients with a History of Bariatric Surgery: a Nation-Wide Population-Based Study**. In *Obesity Surgery* (2022), 32: 8-17, DOI: 10.1007/s11695-021-05763-6

All listed publications were written during the Ph.D. studies of the main author in collaboration with Noeska N. Smit as the main supervisor and Stefan Bruckner as the co-supervisor. Both supervisors contributed significantly with fruitful discussions, advice, and guidance to publish the scientific work.

Paper A: Ingfrid S. Haldorsen contributed by providing application context and data and was available as a medical expert throughout the whole process of developing ParaGlyder.

Paper B: Kari Wagner-Larsen, Camilla Krakstad, and Ingfrid S. Haldorson delivered relevant problems and data that formed the basis for the development of the RadEx platform. Kari and Ingfrid introduced us to the analysis problem from a radiologist's perspective, while Camilla gave us insight into challenges associated with the analysis of histopathological data. Erlend Hodneland introduced us to the challenges faced while working with radiomic tumor profiling data and co-registration of several imaging modalities.

Paper C: Tanja Eichner contributed software development work during an internship at the University in Bergen, and Ingfrid S. Haldorsen had similar roles as in the works before.

Contents

Part I

Overview

«The place to improve the world is first in one's own heart and head and hands, and then work outward from there.»
Robert M. Pirsig

Chapter 1

Introduction

Novel techniques, medical research, and advanced clinical practice in the medical domain partially drive medical visualization research. There has been a significant increase in available medical imaging modalities, sequences, and data per patient in recent years. A single tissue analysis conducted under the microscope is now often fully digitized, and one investigation can potentially deliver hundreds of gigabytes of data that must be analyzed. Finding efficient and effective data analysis methods to cope with the sheer amount of data is an ongoing endeavor involving multiple scientific disciplines.

Medical visualization is mainly a domain-driven science. Visualization researchers can envision novel and exciting techniques for visualizing medical imaging data independently. However, often these do not meet the requirements of the domain experts for use in clinical practice or research. Close contact with medical researchers is therefore essential to create medical visualizations that have the potential to impact the patients the data is derived from. Involving domain experts is challenging as they have their own vocabulary, and often time is limited, but it is the only path to results that may finally be used in practice. Radiologists mainly conduct medical image analysis under time pressure. Therefore, the main criterion for bringing medical visualization research into clinical practice is the efficiency and effectiveness of the created algorithm or application. The rest of the criteria can be summarized as follows: it is good if the radiologists can conduct their work faster with your application or tool. If not, it has to have a significant impact on the patient. This rule can be altered when applying medical visualization to medical research. In research, time is not the most crucial aspect. Visualization should instead allow for exploration and hypothesis formation. Besides visualization where the human in the loop is essential, machine learning is another answer to difficult questions in medical image analysis.

In recent times machine learning has evolved into a general problem solver. However, while artificial intelligence can be beneficial for many scenarios, it struggles in cases where, e.g., additional context is needed, problem definitions are unclear, or ethical value judgments are involved. Furthermore, medical experts must still review the results created by machine learning algorithms for quality assurance. People also have the right in EU law under the General Data Protection Regulation (GDPR) to receive an explanation for every decision made by an automated or machine learning created decision [46]. Visualization is therefore still valuable and essential either to accompany machine learning or to show that more straightforward solutions may also be good enough, following Occam's razors.

1.1 Problem Statement

Medical Visualization is a multidisciplinary application-oriented field where visualization researchers often work closely with domain experts. Such interdisciplinary research frequently results in applications addressing specific medical image analysis problems. The increasing complexity in imaging results from modern magnetic resonance imaging (MRI), X-ray computed tomography (CT), or functional MRI (fMRI) introduces challenging data analysis problems. For example, **multi-sequential** image series are the result of using different imaging sequences within an MRI, and **multi-modal** images are derived by combining different imaging techniques such as MRI and CT. Visualizing such data is a challenge and one of the primary endeavors in this thesis.

Medical research often involves the *analysis* of data from multiple patients, so-called cohorts, to *explore* patterns that can later be used for individual cases and treatment decisions. Such cohorts have varying sizes ranging from a handful of patients up to thousands depending on the frequency of the illness being studied and the size of the hospital or consortium that carries out the study. Data compatibility issues are one of the most common problems when working with cohorts. In general, data derived from a CT investigation is more accessible to *analyze* than MRI data that is not based on a shared value range. Both modalities further share comparison challenges when data is derived from machines of different vendors or protocols. The *analysis* of MRI images within a cohort relies on comparing statistical data features. One problem can be linked to differences between scanner manufacturers, especially in studies across hospitals. Furthermore, acquisition protocols can change over time or between institutions, making the data even less comparable. Many of these problems arise in **multi-patient** data *analysis* which is an additional central aspect of this thesis.

Finally, research results which are often funded by public resources, are *communicated* within the research community but also to the public. There are inherent differences in *communicating* such findings based on the target audience. The first pathway of sharing scientific findings is via scientific papers presented at conferences or published in journals. Scientific papers often follow a strict pattern and deliver insights in a structured way which is efficient if the consumers are used to the format. When *communicating* findings with a broader audience such as the general public, other media is more favorable as the general public is often not trained to gather information from structured manuscripts. This form of *communication* is often called outreach and can consist of blog posts, videos, (newspaper) articles, or even television appearances and interviews. When addressing **multiple audiences**, many challenges have to be addressed to reach the goal of a common understanding. Such challenges include different knowledge levels, different vocabularies, and cultural differences. One central part of this thesis is finding efficient ways to *communicate* scientific results to multiple audiences, from peers to the general public.

1.2 Scope and Contributions

The research conducted in the course of this Ph.D. study has the goal of scaling medical visualization in three different dimensions, namely: multi-modal, multi-patient, and multi-audience. We contributed three application papers and one technique-oriented paper. We solve several problems when dealing with multi-modal medical imaging data of multiple patients or even cohorts while fulfilling the needs of experts working in multiple different scientific fields. To show the utility of our approaches, three out of our four contributions have been evaluated with qualitative user studies. The main contributions of this thesis can be summarized as follows:

Multi-Modal

We contribute the following methods and applications to the area of multi-modal medical visualization:

Stixels: Star Glyph Pixels We introduce the so-called Stixels (**St**art glyph pi**xels**), a visualization method to visualize the values of multiple MRI sequences within a single slice view. With this method, we allow, e.g., medical researchers to analyze the value distribution of multiple medical imaging sequences in only one view.

Interactive Clustering Exploration We propose an application called *ICEVis* (**I**nteractive **c**lustering **e**xporation **Vis**ualization) which enables medical researchers to perform subtumor analysis based on the result of a multiparametric hierarchical clustering approach.

Multi-Patient

Within this thesis, we contribute the following techniques and applications to the area of multi-patient medical visualization.

Interactive Visual Probing Within ParaGlyder (the name reflects the probing interaction), we propose a technique to interactively probe multi-modal medical images, facilitating a digital biopsy of the available data. This technique can be used to compare tumor tissue characteristics to the properties of surrounding tissue. Compared to ICEVis, the probing interaction does not need a pre-processing step of clustering and allows for interactive on-demand data exploration. Furthermore, such biopsies can also be used to compare the tumors of different patients to gather generalizable tumor describing features.

Integrated Visual Exploration of Multiparametric Studies We introduce the application RadEx (**Rad**iomic tumor profiling **Ex**ploration), which enables medical researchers to explore multi-parametric studies of radiomic tumor profiling results visually. The application includes calculating radiomic tumor profiling features of a whole cohort of patients, including several imaging sequences. In addition, the application enables hypothesis formulation and hypothesis investigation for such data by combining dimensionality reduction techniques with interactive visual analysis.

Multi-Audience

We propose the following techniques and presentation methods to target multiple different audiences within the field of visualization in general and, more specifically, medical visualization.

Multiparametric Similarity Visualization Within ParaGlyder, we introduce a multi-parametric similarity visualization that takes the digital biopsy described earlier as a reference and presents a highlights voxel-based multi-parametric similarity to this digital biopsy within a medical imaging sequence. The similarity visualization enables radiologists to highlight the extent of the tumor and potential cancer spreading more quickly than by only using single imaging sequences. In addition, it allows for the assessment of tumor extent and inflammatory regions after tumor surgery, as demonstrated in a brain tumor case. Lastly, it allows machine learning experts to assess which imaging sequences are most informative for the development of automatic segmentation methods.

Interactive Visual Authoring of Guided Dynamic Narratives With ScrollyVis (Scrollytelling Visualization), we enable authors to share either scientific or general stories with a broader audience, e.g., the public. Scrollytelling (stories that progress based on the scrolling interaction of the reader) is a storytelling format where content is revealed by a scrolling interaction of the story consumer. Our application allows for the creation of guided dynamic narratives and creates a ready-to-deploy website that enables scrollytelling by focusing on a non-code approach. With this contribution, we enable scientists and other story authors to create stories about scientific contributions or everyday stories like a trip on the weekend. We target multiple audiences to act both as an author of compelling stories and as a story consumer.

1.3 Thesis Structure

This thesis consists of two main parts. The first part provides an overview of the research conducted in the thesis, while the second part consists of the individual publications. The format of the individual contributions was adjusted to fit the thesis layout. Furthermore, the bibliographies of the individual papers were merged into a unified bibliography.

Chapter 1 is an introduction to the covered topics in this thesis. Chapter 2 surveys state of the art related to our contributions and present how our approaches differ from previously conducted research. Chapter 3 outlines the contributions included in this thesis. We present the three scaling dimensions and relate the contributions within the individual papers to them. Chapter 4 demonstrates how our scientific contributions are used or will soon be used in practice. Chapter 5 provides a discussion and an outlook on possible future work. The second part of the thesis includes the four papers providing further details on the contributions of this Ph.D. thesis.

2

«It's said that a wise person learns from his mistakes. A wiser one learns from others' mistakes. But the wisest person of all learns from others' successes.»
John C. Maxwell

2

Chapter 2

Related Work

Application Driven Visualization	Multi-Modal	Multi-Patent	Multi-Audiences	Evaluating Visualization Applications
Becker et al. [6]	Preim et al. [125]	Angelelli et al. [3]	Ma et al. [100]	Ellis et al. [39]
Buja et al. [19]	Klippet et al. [85]	Klemm et al. [83]	Ma et al. [100]	Munzner [108]
Schneiderman [142]	Klippel et al. [84]	Bernard et al. [8]	Kosara et al. [89]	Isenberg et al. [65]
Card et al. [22]	Glaßer et al.[47]	Raidou et al. [130]	Dahlstorm [31]	Sedlmair [139]
Baldonado et al. [168]	Ropinski et al. [133]	Preim et al. [124]	Seyser et al. [141]	Glasßer [48]
Tory et al. [159]	Kehrer and Hauser [78]	Jönssen et al. [73]	Conlen et al. [29]	Smit et al. [144]
Brehmer and Munzner [16]	Raidou et al. [128]		Tong et al. [158]	Preim et al. [122]
Preim et al. [123]	Jäckle et al. [66]		Joubert et al. [70]	Preim et al. [126]
Miksch et al. [103]	Stoppel et al. [150]		Hohmann et al. [61]	
Munzner [109]	Lawonn et al. [95]		Conlen et al. [30]	
	Opach et al. [115]		*Stornaway* [151]	
	Gillman et al. [46]		*Twine* [161]	
			Tableau Stories [42]	

Table 2.1: The related work of this thesis is categorized into five sections: (1) application-driven visualization design, (2) multi-modal, (3) multi-patient, (4) multi-audience, and (5) related work on how to evaluate visualization applications.

In this chapter, we summarize research related to this thesis. In Table 2.1 the related work is listed, categorized by which research aspect this relates to, and ordered by the year of publication. Furthermore, we differentiate between work related to the specific aspects of this thesis, i.e., multi-modal, multi-patient, and multi-audience. In addition to the specific areas we cover within this thesis, we present more general related work regarding the application-driven design of visualization applications and the evaluation of visualization applications.

2.1 Application Driven Visualization

Munzner [109] defined visualizations as tools to enhance the cognitive capabilities of humans, which are particularly suitable in cases where the users would like to explore and discover new knowledge from the data and processes they analyze. There are various application fields where visualization is applicable. One field of interest for this thesis is medical visualization [123]. Medical visualization covers a vast and complex domain where pre-knowledge is essential as it may be applied to areas like anatomy, pathology, and biomedical engineering. Preim and Botha [123] delivered a thorough overview of challenges and opportunities within the field of medical visualization. In a more recent publication by Gillman et al. [46] ten open challenges in the field of medical visualization are discussed, including multi-modal visualization, the focus of this thesis.

Tory et al. [159] defined a common taxonomy of visualization based on the data being used. They introduced the terms *Scientific Visualization* which deals with data with inherent spatial information and *Information Visualization* which works with abstract, nonphysical data [159]. Furthermore, *Visual Analytics* combines the two visualization fields and further incorporates concepts like data mining and analysis of data using statistics. The *human in the loop* is an essential concept within visual analytics [109]. Within this thesis, we will present new approaches in scientific visualization and visual analytics platforms to enable experts from different fields to investigate their data and formulate hypotheses.

General Concepts and Techniques

Many design choices can be effective in various scenarios of visual analytics applications. In this section, we highlight the most important ones for this thesis.

The Information Seeking Mantra was introduced by Schneiderman [142] and is defined as follows: *Overview first - Zoom and filter - Details-on-demand.* The whole process should be enabled via interaction. The mantra is valid for the whole field of visualization research, and an expanded version has been introduced by Keim et al. [79]: *Analyze first, Show the Important, Zoom, filter and analyze further, Details on demand.* Following the mantra [142] and the adaption by Keim et al. [79], visual analytics enables an optimal workflow for users without overloading them with information.

The Data-Users-Tasks Design Triangle introduced by Miksch et al. [103] defines three different aspects as the most important when it comes to the implementation of visual analytics platforms, namely: characteristics of the data, users, and the users' tasks.

A Multi-Level Topology of Abstract Visualization Tasks by Brehmer and Munzner [16] defines the central questions: why, how, and what. *Why* is a task performed, *how* is a task executed, and *what* are the inputs and outputs of the task.

Multiple (Coordinated) Views are often used within visual analytics platforms and were first described by Baldonado et al. [168]. Their work proposed guidelines on how to design multiple view systems effectively.

Brushing and Linking is a concept enabling filtering of data by selecting interesting sub-parts of data in various views, first shown within scatterplots [6] and later in a more general approach discussed by Buja et al. [19].

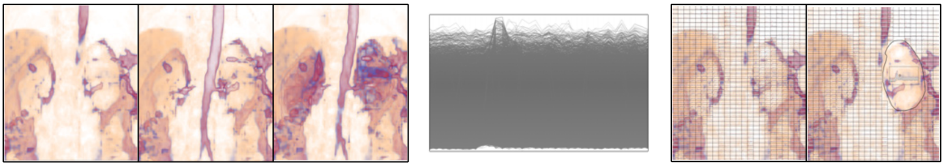

Figure 2.1: Stoppel et al. [150] presented primitives for the visualization of time-dependent spatial data presenting volumetric changes. Traditionally animated volume renderings are used to visualize time-dependent data (left). While showing all time-intensity curves together, a cluttered plot is a result (middle). They used small multiples combined with interaction methods, resulting in a less cluttered, more informative, and usable approach (right).

Focus + Context aims to visualize a focus object perceivable while illustrating the surrounding context without distracting the focus. Techniques to facilitate this include unsaturated colors for shading, different shading methods per object class to varying entire rendering concepts [95].

2.2 Multi-Modal

Kehrer and Hauser [78] stated that visualization and visual analysis is of high importance in exploring, analyzing, and presenting scientific data. Multifaceted data is used more frequently either from different data sources (multi-modal data), from multiple simulation runs (multi-run data), or multi-physics simulations of interacting phenomena (multi-modal data from coupled simulation models). To analyze such data, visualization is more often combined with computational analysis. The authors analyzed existing methods and showed opportunities for new research directions to cope with such data sources.

Lawonn et al. [95] presented a survey on multi-modal medical data visualization. Multi-modal is an ambiguous term in medical visualization and either describes different imaging techniques resulting from different scanners or different imaging sequences within the same machine. An example of the former is MRI and CT imaging, while an example of the latter are different acquisition parameter weighting strategies in MRI imaging, such as T1 or T2. Visualizing multi-modal imaging data is still one of the ten open challenges highlighted by Gillman et al. [46]. Lawonn et al. [95] stated that the goals of multi-modal medical data visualization are the following:

- reduction of complexity and cognitive load
- enabling, improving, or accelerating the decision making processes
- providing tailored visualizations for specific applications

Glyph Based Approaches

Standard visualization methods often reach limitations when combining more than three different imaging modalities. There are different methods of solving this limitation, e.g., by presenting multivariate data, ranging from scatterplot matrices and parallel coordinate plots to glyph-based visualization techniques [66]. One possible solution is to place glyphs or primitives on top of medical images to present multi-dimensional data while keeping the context of the medical data available, following the focus-and-context visualization pattern. One example of such a method is presented by Stoppel et al. [150] named Graxels, and shown in Figure 2.1. The method provides information regarding temporal developments as small multiples placed in their spatial context. Furthermore, the authors allow for multiple interactions in both the spatial and value domain.

Ropinski et al. [133] presented a survey on glyph-based visualization techniques for spatial multivariate medical data. They proposed a classification of glyph techniques into two main groups: those supporting pre-attentive and attentive processing. Furthermore, the authors proposed guidelines to support improved glyph-based visualizations for the medical domain. One of the essential aspects of glyphs is their shape which should be easily perceivable and unambiguous [133]. Therefore, the authors differentiate two different types of glyph shapes:

- *Basic* glyph shapes are geometric shapes that are modified by their geometric properties like size and orientation. Examples of basic glyph shapes are spheres, cuboids, and ellipsoids.
- *Composite* glyph shapes are composed of basic glyph shapes. They are more specialized and defined by mapping functions, e.g., parameters are mapped to geometric properties such as radius and length. Composite glyphs are often used to display multivariate data.

One example of such composite glyphs was presented by Jäckle et al. [66]. The authors used so-called star-glyph insets to achieve overview preservation while visualizing multivariate data. Although star glyphs can be presented in many ways, the authors chose to use a design where data lines radiate from the center, and the ends are then connected by a contour line forming the glyphs. Essential aspects to consider when using star glyphs are discussed by Klippel et al. [84, 85] including how to shape characteristics to influence classification tasks and also how color enhances star plot glyphs. When thinking about glyph designs, one crucial aspect is defining glyph shape. Opach et al. [115] proposed a guide to help decide between polyline-based glyphs as used by Stoppel et al. [150] and star glyphs in a grid plot. They concluded that it is task-dependent: polyline glyphs are better when datapoint values must be read within the glyphs, and star glyphs are better when the task involves a visual search among glyphs.

We also introduce a star glyph-based grid layout that is designed according to the glyph design guidelines by Opach et al. [115]. It enables radiologists to analyze multiple sequences simultaneously without having to display them in multiple juxtaposed views. In comparison to prior work, we use multi-parametric imaging data of one patient at one timepoint and do not visualize data changes over time. Furthermore, our design is focused on identifying pattern differences and not reading individual values. Therefore, we decided to use star-glyphs colored by similarity to a probed area within the data following the guidelines of Klippel et al. [84, 85].

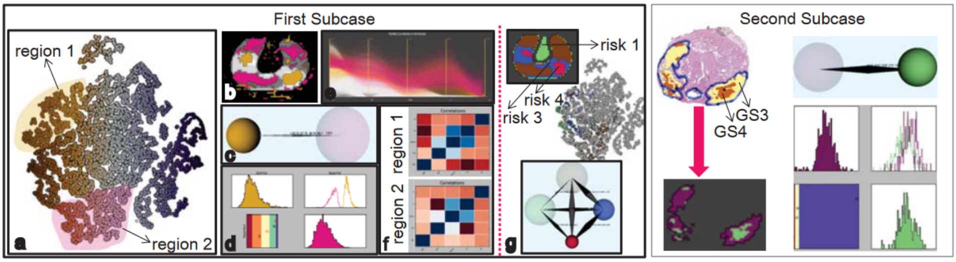

Figure 2.2: Raidou et al. [128] presented a visual tool for the exploration and analysis of the feature space defined by imaging-derived tissue characteristics. They further support knowledge discovery and hypothesis generation, and confirmation. The Figure presents a case study on prostate cancer patient.

Visual Analytics

Besides considering how to visualize multivariate medical imaging data, one also must consider the analysis tasks medical experts must carry out using the data. Raidou et al. [128] presented a visual analytics platform for tumor tissue characterization. An example use case of their approach is presented in Figure 2.2. They proposed an application that enables exploration and visual analysis of image-derived tissue characteristics that enables hypothesis discovery, generation, and confirmation. Clinical researchers can analyze tumor tissue characteristics in high-dimensional feature space by including multiple linked interactive views.

Similar to their approach, Glaßer et al. [47] presented a visual analytics application for breast tumor analysis in dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) that facilitates a voxel-wise glyph-based overview and region-based analysis. Their approach is depicted in Figure 2.3. In contrast to the approach of Glaßer et al. [47] we visualize multiple sequences on top of one imaging slice. Preim et al. [125] published a survey on visual exploration and analysis of perfusion data. Their work outlined how 2D parameter maps, 3D visualizations of parameter volumes, and exploration techniques can be combined. In addition, they discussed related work based on three major application areas: ischemic stroke diagnosis, breast tumor diagnosis, and the diagnosis of coronary heart diseases.

We further enable the exploration of medical imaging data by interactive probing and allow for the analysis of radiomic tumor profiling data within a cohort of patients while focusing on the imaging data. The experts can still inspect the imaging data of all patients while analyzing statistical features of the whole tumor data throughout the cohort.

2.3 Multi-Patient

In contrast to the previous section, where we focused on visualization techniques and visual analytics applications mainly targeted at multi-modal medical imaging data, this section focuses on the analysis of the data of multiple patients. In medical research, groups of patients are also called cohorts. Preim et al. [124] showed an extensive

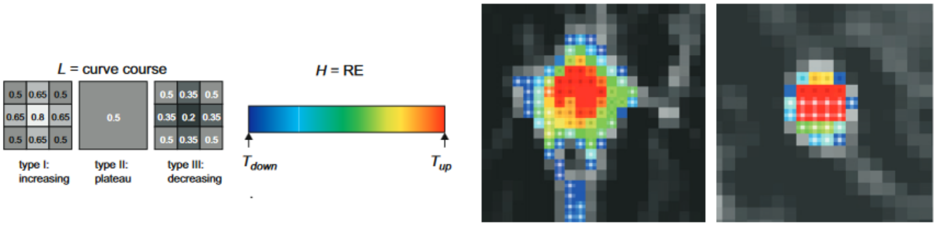

Figure 2.3: Glaßer et al. [47] presented a glyph-based visualization of two lesions. The relative enhancement and the relative enhancement curves, based on contrast-enhanced MRI, are mapped to the voxels of the images using a rainbow colormap.

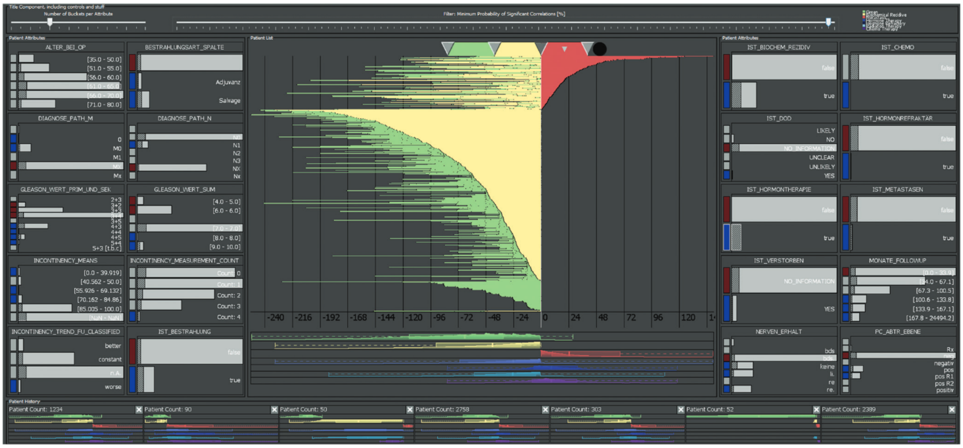

Figure 2.4: Bernard et al. [8] introduced a system enabling physicians to define and analyze cohorts of prostate cancer patients. All the visualizations used in the application are synchronized, and a centered list-based visualization delivers an overview of large sets of patient histories.

overview of visual analytics solutions aiming to support public health professionals in analyzing public health-related data. They described requirements, tasks, and visual analytics techniques often used in public health-related applications.

Angelelli et al. [3] presented a visual analytics approach that enables the visual exploration and analysis of large amounts of heterogeneous data to help generate and validate hypotheses. They implemented data-cube-based models to handle overlapping data subsets and seamlessly integrate data during visualization, linking spatial and non-spatial data views.

In contrast to Angelelli et al.'s work [3], which targeted data on cognitive aging, Bernard et al. [8] developed a visual-interactive system for prostate cancer cohort analysis. Their application is visible in Figure 2.4. It was developed in close collaboration with medical researchers and helped the clinical experts to efficiently and effectively analyze single and multiple patient histories at a glance. They used several linked views and information visualization techniques combined with guidance concepts.

Klemm et al. [83] focused on epidemiological data, which enables experts in the

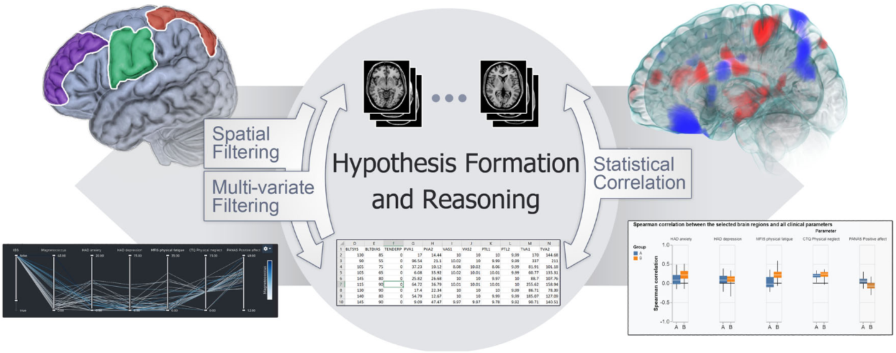

Figure 2.5: Jönsson et al. [73] introduced an application for the interactive visual analysis of brain imaging data and clinical measurements. Targeted toward neuroscientists, the authors enable the analysis of correlations between active brain regions and physiological and psychological factors.

field to investigate their data pool for hypothesis validation and generation. The authors combined image-based and non-image data in a visual analytics platform that extensively features interaction methods to analyze lower back pain-related studies. In contrast to Klemm et al. [83], Raidou et al. [130] focused on the exploration of radiotherapy-induced bladder toxicity in a cohort study. The authors introduced an interactive application that provides multiple linked views where inter-patient and temporal exploration, analysis, and comparison are supported. They evaluated their approach with clinical experts who positively assessed the functionality and the design of the visualizations.

Jönssen et al. [73] focused on the analysis of brain imaging and clinical measurements as shown in Figure 2.5. They introduced VisualNeuro, a hypothesis formation and reasoning application for multivariate brain cohort study data. The authors used a parallel coordinate plot to enable effective subject group selection. Furthermore, they employed the Welch's t-test for brain region filtering and multiple visualizations based on the Pearson correlation between brain regions and clinical parameters to enable correlation analysis. Finally, they performed a qualitative user study with three neuroscientists from different domains.

Our work further enables cohort data analysis by using multiple linked views presenting both imaging data and clinical parameters. In contrast to the presented approaches, we deal with multi-sequential MRI data for each patient within the cohort and enable the radiologists to drill down to the imaging level for each patient while keeping the context of the whole cohort available. Furthermore, we allow for the analysis of radiomic tumor profiling data, a recent research field that aims to find new imaging biomarkers correlated to outcome-related variables.

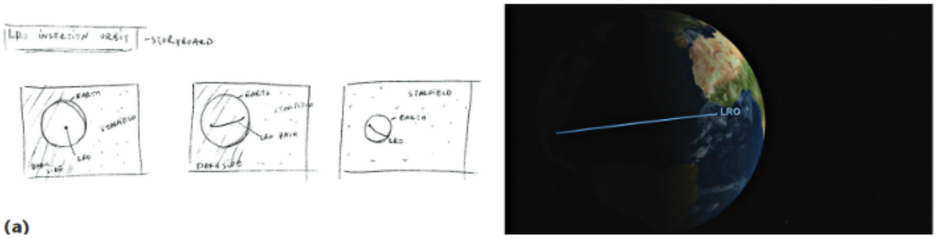

Figure 2.6: From storyboard to visualization: the story of NASA's Lunar Reconnaissance Orbiter. Scientific Storytelling Using Visualization referred by Ma et al. [100].

2.4 Multi-Audience

This section focuses on multi-audience-related approaches related to scientific storytelling, which is writing about scientific observations to create a captivating story and scrollytelling where a digital story progresses by a scrolling interaction of the story viewer. Storytelling is the soul of scientific communication, according to Joubert et al. [70]. The authors stated that we must go beyond presenting facts and evidence. Instead, we must create emotional connections between scientists and the public. According to Joubert et al. [70], storytelling can be a powerful way to increase engagement in science and stories to help people understand, process, and recall science-related information.

Storytelling is rapidly gaining momentum within the field of visualization, and with this trend, techniques are introduced which enhance understanding [158]. As a result, more and more storytellers integrate complex visualizations into their narratives. The authors presented a classification of literature on storytelling in visualizations in two different dimensions. The first dimension is defined by the questions: *Who?*, *How?* and *Why?*. The following shows the first category for each of the first dimensions:

- **Authoring-Tools:** This class addresses the question *who* creates the story and the narrative.
- **Narrative:** Narratives define *how* an author tells a story.
- **Memorability:** Memorability of the presented information defines *why* we use storytelling.

The second dimension of *How?* mainly focuses on the possibilities for a story viewer to traverse the story, which can be dynamic, static, or random access based. The authors also provide an extensive overview of visualization methods used within storytelling and note that almost all use information visualization and very little use of scientific visualization methods.

Ma et al. [100] stated that there is a clear need to consider how storytelling and visualization can make scientific findings more comprehensible and accessible to the public. The authors point out that scientific visualization can learn from information visualization as it is already broadly used in storytelling, whereas scientific visualization is less often used. Ma et al. [100] outline that visualization and storytelling is one key concept at the Scientific Visualization Studio at NASA's Goddard Space Flight Center, as also shown in Figure 2.6. Furthermore, when producing visualizations for science museums, storytelling is a crucial aspect [100].

Kosara et al. [89] wrote already in 2013 that storytelling could be the next big step for visualization. It allows a natural flow from exploration to analysis and finally to presentation. The presentation should be seen with at least equal importance to exploration and analysis. The authors mention that based on the maturing of the visualization research field, researchers must focus on presentation as visualization is increasingly used for decision making.

Dahlstorm [31] argued that narratives and storytelling are essential instruments to communicate to a non-expert audience as one of their main channel of science information is mass media content based on narrative formats. The author also raises awareness that narratives are intrinsically persuasive, which can be used to persuade resistant audiences but could also raise ethical concerns. Dahlstrom emphasizes the importance of scientific communication in a new media environment, especially on the web.

Narratives and Storytelling on the Web

Seyser et al. [141] analyzed visual storytelling in online journalism and concluded that storytelling and, more explicitly, scrollytelling mainly use text in combination with multimedia content. Most authors use infographics to visualize complex data in long-form journalism published online. The authors found two main categories of such articles: *text-based*, where multimedia content enhances the text with further contextual information, and *image-centric*, where text complements photos, graphics, and animations. The authors also mention that the Gestalt laws should be considered to structure visual elements to be easily understood.

Hohmann et al. [61] examined the design of interactive articles in their article. The authors state that interactive articles are applicable in many domains, e.g., research dissemination, journalism, education, and policy and decision making. Each of these domains comes with its opportunities and challenges. In Figure 2.7, the authors presented an extensive overview of interactive articles published on the web. Hohmann et al. [61] further discuss the theory and practice of interactive articles based on several online articles and found the following aspects fundamental:

Connecting People and Data: Such a connection can be facilitated in many ways, including animation, which can also be used to improve engagement. Animations effectively communicate state transitions, uncertainty, causality, and constructive narratives. Unit visualizations evoke empathy in the reader, and games efficiently convey information and build empathy.

Making Systems Playful: Interactive visualizations can allow readers to build an intuition about the behavior of a described system which ultimately leads to a fundamentally different understanding compared to looking at a static set of equations. In addition, sandbox simulations featuring multiple changeable parameters allow readers to experience the behavior changes on the fly and therefore get a deeper understanding of the underlying processes.

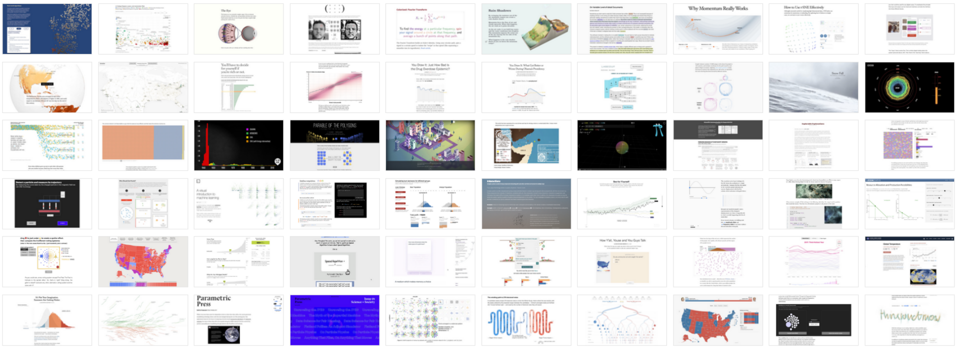

Figure 2.7: Hohman et al. [61] presented exemplary interactive articles from around the web. The authors described the essence of interactive articles and reflected critically on their own experience of publishing interactive open source at scale. They further discuss challenges and open research directions in the authoring, designing, and publishing interactive articles.

Prompting Self-Reflection: Asking story or article consumers to reflect on reading or studied material back on themselves is known to impact learning outcomes positively. This mechanism can also be used in interactive articles by letting readers make predictions or reflect on the material presented.

Personalizing Reading: One option is to personalize the story's content based on the reader's input, e.g., date of birth. Text and multimedia can change automatically based on the readers' input creating an individualized story for every story consumer. Such techniques have been shown to increase engagement and learning outcomes for the readers. Segmenting information into smaller pieces and letting readers consume it at their own pace is another way to facilitate personalized reading.

Reducing Cognitive Load: Authors must bear in mind the readers' expertise and not overload them. Details on demand are vital in reducing the readers' cognitive load.

The following section will discuss research in story authoring tools and the development of an extensive story that is part of a museum exhibition.

Interactive Visualization Story in a Museum Exhibition

Ma et al. [99] presented Living Liquid: Design and Evaluation of an Exploratory Visualization Tool for Museum Visitors. The authors faced several challenges during the development of their interactive visualization. First, it must engage visitors personally, and visitors often lack the background to interpret visualizations of scientific data. Often time to spend at an individual exhibits within a museum is limited. The authors conclude the following lessons learned which they derived from visitor evaluations: do not distract visitors from the data they should explore, include background information about the visualizations, making visualizations understandable is more important than

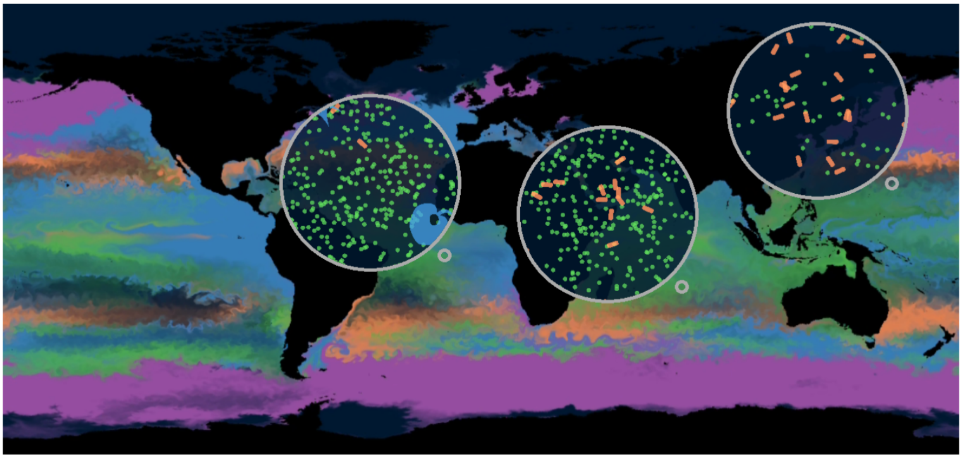

Figure 2.8: Ma et al. [99] present Living Liquid: Design and Evaluation of an Exploratory Visualization Tool for Museum Visitors. The authors designed an interactive visualization about simulated marine microbes using an iterative development process and provided guidance on how to best design visualizations for public audiences.

scientific accuracy, and layer data accessibility to structure inquiry. A picture of their visualization is presented in Figure 2.8.

Story Authoring Applications

Conlen et al. [29] introduced Idyll, a novel "compile-to-the-web" language targeted at web-based interactive narratives. They allow authors to control document style, layout, reader-driven events, and a structured interface to JavaScript components. By evaluating their approach with undergraduate computer science students, they were able to show that Idyll reduces the effort of creating interactive articles. As an extension, Conlen et al. [30] later introduced Idyll Studio, a structured editor for authoring interactive and data-driven articles. An image of the Idyll Studio framework is visible in Figure 2.9. They now support reflective documents which enable inspection and modification of their programs at runtime. The authors conclude that Idyll Studio enables non-technical users to complete tasks more rapidly than expert users using tools they are familiar with while having a reduced cognitive demand compared to existing tools.

Two additional well known authoring tools are *Stornaway* [151] and *Twine* [161]. *Stornaway* enables authors to create dynamic narratives which can be exported as videos and can be embedded in websites. Dynamic narratives allow the story viewer to decide on the fly which direction the story should pursue. *Twine* allows for dynamic narratives and creates stories that are web-based and mainly focus on text and image integration. Tableau [42] stories are another compelling way to tell data stories that can be defined directly in the Tableau software, which is often used to analyze and explore data. The stories created with the application are presented in a slideshow fashion and allow for interactivity.

We contribute an editor for guided dynamic narratives, and in contrast to the afore-

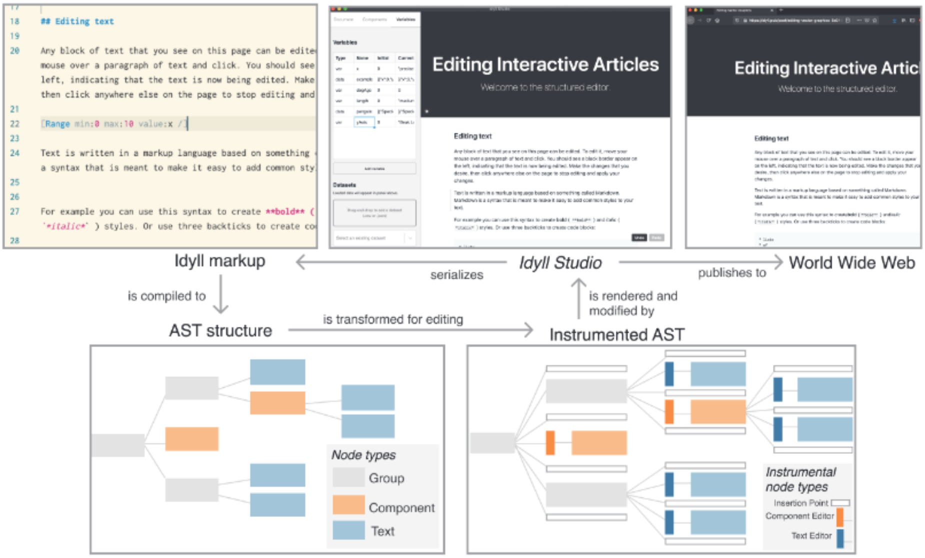

```
## Editing text

Any block of text that you see on this page can be edite
mouse over a paragraph of text and click. You should see
left, indicating that the text is now being edited. Make
then click anywhere else on the page to stop editing and

[Range min:0 max:10 value:x /]

Text is written in a markup language based on something 
a syntax that is meant to make it easy to add common sty


For example you can use this syntax to create **bold** (
'*italic*' ) styles. Or use three backticks to create co
```

Figure 2.9: Conlen et al. [30] presented Idyll Studio: A Structured Editor for Authoring Interactive & Data-Driven Articles. Idyll Studio enables non-technical users to author data-driven articles in a shorter time than expert users using standard tools.

mentioned approaches, we use scrollytelling as a storytelling format. This allows users to choose the story's path and enables authors to include various media formats, including scientific visualizations such as 3D volume rendering. In addition, our editor is based on a non-code paradigm and exports the website ready-to-deploy with the website code still available to the author to adapt if needed.

2.5 Evaluation of Visualizations

When developing visualization applications, one crucial aspect is the evaluation of design and implementation choices. In this section, we discuss several publications on how to perform evaluations in the visualization domain and, more specifically, in medical visualization. Munzner [108] introduced a model for visualization design and validation which consists of four layers:

- characterize tasks and data in the words of the problem domain
- abstract tasks and data into operations and data types
- design the visual encoding and interaction techniques
- create algorithms to execute techniques efficiently

She further provides guidance to determine which evaluation approaches are appropriate based on the different levels of the visualization design. Ellis et al. [39] discussed why evaluation of visualization applications is complex and proposed explorative evaluation to discover new things about visualization techniques. The authors further provide practical guidance on how to evaluate visualization applications. Isenberg et al. [65] outlined current and historical evaluation practices and developments and concluded that there was an increased emphasis on user experience and user performance in contrast to algorithmic performance and qualitative result inspection through images. Furthermore, they conclude that there is an increasing interest in how new visualizations help in data analysis and reasoning. More recently, Sedlmair [139] proposed seven guiding scenarios which aim to characterize better different contributions resulting from design study projects. Furthermore, these scenarios should help provide better guidance on designing the evaluation of design study contributions.

In the field of medical visualization, one problematic aspect is that often there is no ground truth to compare to [48]. This results in the limitation that medical visualizations can often only be evaluated in a qualitative comparison that neglects objective measures like accuracy and task completion time and focuses more on user preference. The authors use guidance to enrich qualitative evaluations by presenting quantitative analysis in their work. Preim et al. [122, 126] discussed how evaluation practice in medical visualization should be conducted carefully and should be performed over a more extended period. The authors argued that individual situations must be assessed, and the adaption to new techniques should be analyzed.

3

«I cannot guarantee that my answer causally correlates with your question.»
Eduard M. Gröller

Chapter 3

Contributions

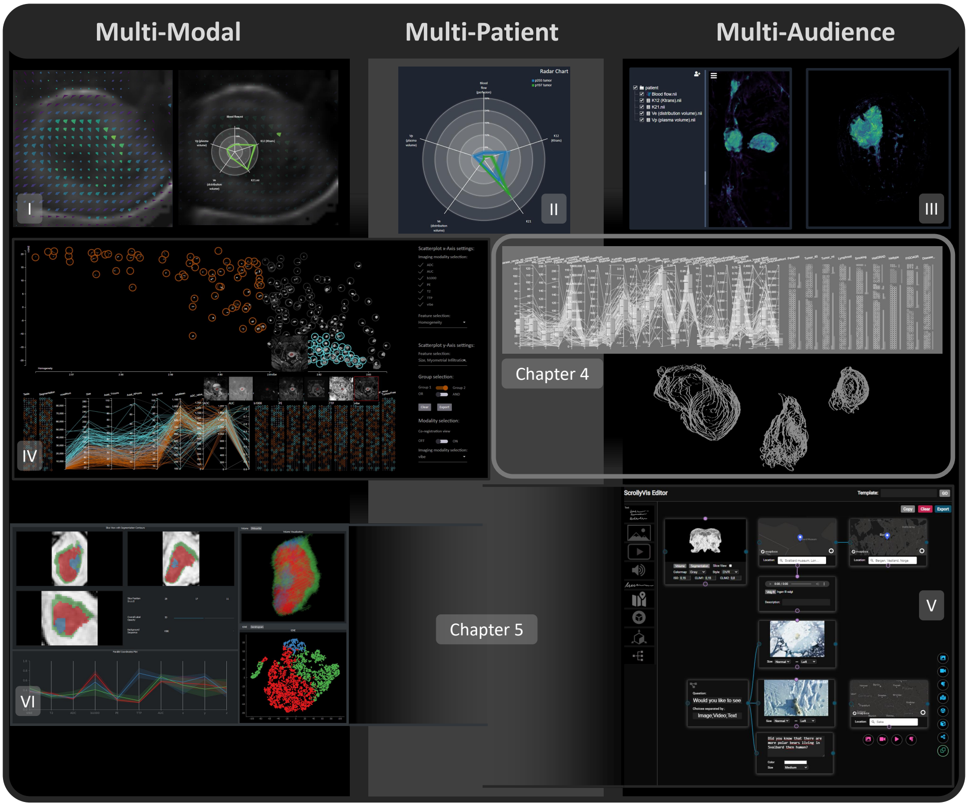

Figure 3.1: Overview of the contributions contained in this thesis. The contributions are positioned according to the dimensions they scale in. Part I, II, and III are contributions introduced in paper A. Image IV depicts paper B, and V represents paper C in this thesis. Lastly, image VI represents paper D. The gray images to the right of image IV present the extension of paper B for making our contribution ready to be used for cervical and prostate cancer data described in Chapter 4. The region between images V and VI presents the area which we describe as future work of the related papers and this thesis discussed in Chapter 5.

The contributions described in this thesis can be grouped according to the scale dimensions they focus on. Overall, this thesis aims to scale up medical visualization in the dimensions of number of modalities, number of patients, and number of target audiences. An overview of the contributions and their categorization is shown in Figure 3.1. Paper A proposes an application that allows for the exploration and analysis of multi-parametric studies targeting multiple audiences, including gynecological and neurological cancer imaging experts and machine learning researchers. The contributions of the paper are presented in components I, II, and III in Figure 3.1. Paper A scales up across all three dimensions. Paper B builds upon paper A and describes an application that allows clinical experts to formulate and prove hypotheses based on multiple imaging modalities and clinical parameters. Dimensions one and two are therefore targeted in paper B and visualized by component IV in Figure 3.1. The focus of paper C is to make scientific communication more accessible both for the storytellers and the story consumers. Our proposed technique enables the creation and authoring of guided interactive scrollytelling websites targeted at scientific communication. Paper C focuses on the multi-audience dimensions intending to reach various communities, including scientists as well as the general public, as visible in component V in Figure 3.1. Paper D describes an application to enable interactive sub-tumor analysis in gynecological cancer data. In the application, multiple modalities are used in a hierarchical clustering approach, as seen in component VI in Figure 3.1. The gray images to the right of part IV represent an unpublished extension of RadEx described in Chapter 4 and the area between components V and VI represents future work we discuss in Chapter 5. In the following, we will present details about the individual contributions.

3.1 Multi-Modal

When analyzing medical imaging data in endometrial cancer research, medical researchers review multiple sequences they must analyze simultaneously. For example, the MRI imaging results of a single investigation of a cancer patient deliver different pictures of the same structures depicting different metabolic properties based on the imaging protocol. These properties hold information necessary for risk stratification and treatment decisions. However, analyzing up to seven image stacks simultaneously is not easy and could be improved by advanced visualization methods. Interesting data exploration and analysis avenues include tumor homogeneity, sub-tumor analysis, and tumor texture exploration with radiomic tumor profiling.

3.1.1 Tumor Homogeneity

Reflecting on the related work in multi-modal medical visualization, we developed a visualization method that visualizes multiple sequences overlaid on a single slice. We call this method Stixels which stands for Star Glyph Pixels, and we show an example in Figure 3.2. Stixels are glyphs presented on top of a single image slice in a regular grid, representing the data given by all sequences available for analysis within a single grid cell per glyph. With this technique, we present multi-sequential imaging data as non-spatial data on top of spatial data allowing for the exploration of imaging patterns.

The basis of the Stixels is a radial chart represented in B in Figure 3.2. First, we place a regular grid on top of the imaging slice where the size of the grid cells is variable and defined by the radiologist on the fly. Next, we consider every pixel within the grid cell and calculate the mean value of every sequence. Afterward, we mark the values on

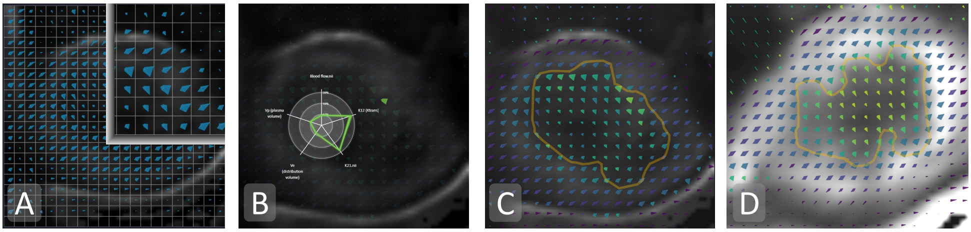

Figure 3.2: In ParaGlyder, we introduce a visualization method using star glyphs to present multiple multi-modal imaging sequences within one slice. As presented in A, the grid in which the glyphs are placed is adaptable. In B, the radial boxplot used to create the glyphs is shown, which appears when hovering over one glyph. C and D show Stixels colored by similarity to a selected region in two different patients highlighting the capability of showing multi-modal heterogeneity within the tumor.

every axis within the radial chart presented in B in Figure 3.2. Finally, the marks are connected and thereby form the Stixel figure.

To enable a more detailed analysis of a grid cell of interest, hovering the presented Stixels reveals the underlying radial plot as shown in image B of Figure 3.2. In addition, we also enable the radiologists to probe areas of interest within the medical imaging data for further analysis. This feature is discussed in more detail in Section 3.2 as it also enables radiologists to compare the tumor data of different patients with each other. Furthermore, the Stixels can also be colored, reflecting similarity to a probed area when probing the data as depicted in image C and image D in Figure 3.2. Radiologists can use this to analyze the homogeneity of the tumor and how it differs throughout the whole imaging volume.

Depending on the type of cancer, tumor tissue can be either homogeneous or heterogeneous. When comparing image C and image D in Figure 3.2 one can note that the two tumors differ in their imaging value homogeneity. All Stixels within the tumor, within the outline marked in D, have a similar shape. In contrast, the Stixels within the tumor in C are different. Tumor homogeneity is potentially linked to aggressiveness within the tumor [171]. The Stixels are an easy and potentially effective visualization to depict homogeneity within the tumor and could therefore influence treatment decisions. Furthermore, they might identify relevant regions within the tumor where a biopsy should be performed for further tissue analysis.

3.1.2 Sub-Tumor Analysis

In the case of endometrial cancer or cervical cancer, tumor tissue is often quite heterogeneous. Tumor tissue homogeneity might be linked to aggressiveness, making it a vital characteristic to study in cancer imaging research. Besides statistical measures based on the imaging characteristics of the whole tumor, another way to analyze it is by deriving sub-tumor parts. Our medical collaborators informed us that tumors often consist of multiple parts, such as necrotic cores, i.e., parts of the tumor consisting of dead tissue. Furthermore, the tumor could also be surrounded by inflamed tissue. These tumor parts could also be linked to the tumor's aggressiveness and can be used to predict outcomes and influence treatment decisions.

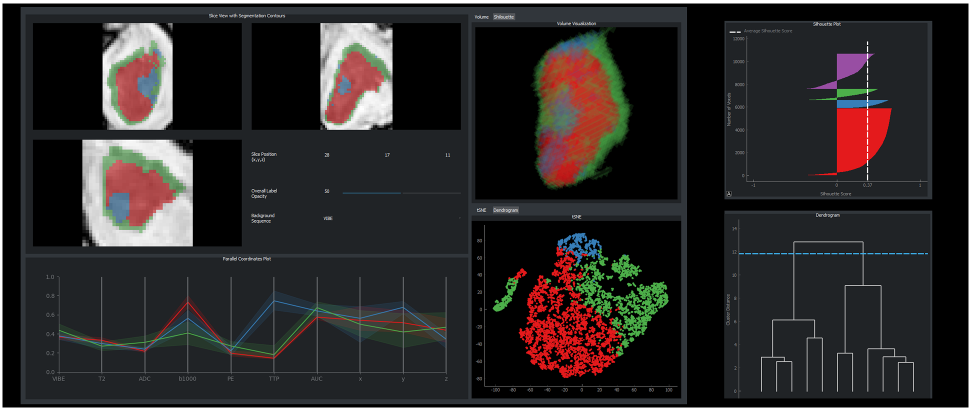
Slice Position (x,y,z) 28 17 11
Overall Label Opacity 50
Background Sequence VIBE

Figure 3.3: ICEVis consists of a 3D tumor view and calculated cluster information. A parallel coordinate plot, a dendrogram, a silhouette plot, and a t-SNE representation of the tumor segmentation are embedded to support the clustering exploration process.

We introduce an interactive clustering exploration for tumor sub-region analysis in multi-parametric cancer imaging to analyze such sub-tumor parts. ICEVis (paper D) enables medical researchers working in cancer imaging research to analyze sub-parts of tumors. A screenshot of the application is visible in Figure 3.3. Our tool facilitates a hierarchical-clustering approach to derive up to ten tumor sub-parts displayed in various linked visualizations. Furthermore, the results can be used in cohort analysis in the future to investigate if specific patterns can be linked to outcome-related indicators.

3.1.3 Radiomic Tumor Profiling

In addition to the already mentioned methods for analyzing tumor imaging data, radiomic tumor profiling is one of the most promising methods in cancer outcome prediction-related research based on imaging data. Radiomic tumor profiling is a method that is based on statistical measurements derived from tumor imaging results linked to clinical variables to find meaningful predictors. It follows the following steps, as depicted in Figure 3.4. First, a segmentation of the tumor is needed, as only data belonging to the tumor should be considered, and typically, the method is applied to multiple sequences. Then, statistical measurements of different levels are derived from the segmented tumor in multiple modalities, resulting in various parameters ready for further analysis. These parameters hold information about the homogeneity of the tumor tissue and more superficial measurements such as the variance or mean value for each sequence. The next and final step in the analysis pipeline is the correlation of the derived biomarkers to clinical variables such as survival time or whether the patient died from active disease.

To support radiomic tumor profiling analysis and, in particular, the analysis of homogeneity measurements as they might be linked to tumor aggressiveness, we developed RadEx presented in paper B. RadEx allows medical experts to analyze cohort tumor imaging data within one application using visual analytics. First, the application incorporates pre-processing steps for data cleaning, creating the radiomic tumor

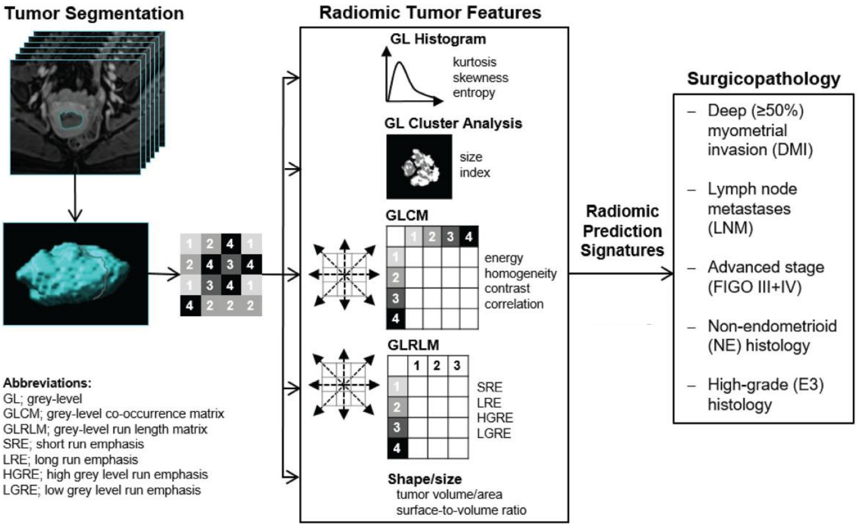

Figure 3.4: The radiomic tumor profiling pipeline, adapted from Fasmer et al. [41]. The process begins with MRI data capturing structures of interest. Secondly, a segmentation of the tumor data is required to calculate statistical features, only taking the target structure into account. The third step in the pipeline calculates the statistical features that can range from first-level statistics such as the mean value and the variance of the data up to higher-level statistics such as the gray level co-occurrence matrix and the gray level runtime matrix. As the last step, the features are correlated with the clinical outcome variables to find correlations of interest.

profiling measurements and visual representations for each tumor. Second, it enables data analysis by providing a visual analytics dashboard. As this method is designed for analyzing cohort data, we describe further details of the approach in Section 3.2.

3.2 Multi-Patient

While treatment decisions and diagnoses are primarily based on the data of single patients, research often relies on cohorts of patients and therefore involves the analysis of multi-patient data. Within paper A which presents ParaGlyder, we introduced a novel technique to probe the imaging data and visualize the probing result in a radial boxplot as visualized in Figure 3.5. We enable the radiologists to place one or two spheres within the medical imaging result and analyze the voxels located within the placed spheres. Radiologists can now compare the data within the tumor and its surroundings or data describing the tumor of different patients. In paper B, which presents RadEx, we presented a multi-parametric tumor analysis platform for analyzing tumor texture features of medical cohort data.

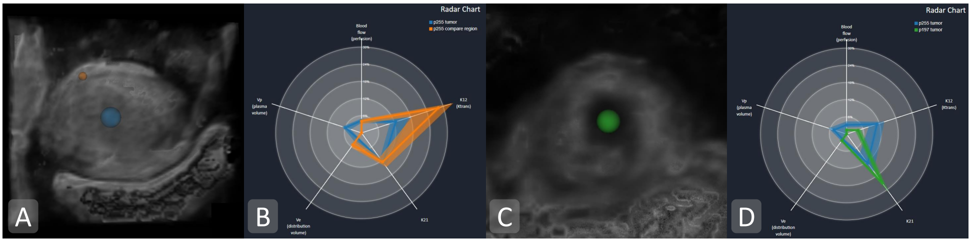

Figure 3.5: The probing interaction implemented in ParaGlyder (paper A) performs a digital biopsy within the patient's imaging data. In A, two spheres are placed within the imaging data of one patient probing the tumor and its surroundings. We present the resulting radial boxplot in B. In C, an additional sphere is placed in the tumor data of another patient. D presents the result of comparing the tumor within the first patient (A) and the second patient (C).

3.2.1 Interactive Probing for Tumor Pattern Analysis

Analyzing tumor patterns within a patient and across multiple patients is an essential part of gynecological cancer imaging research. Such patterns might help characterize tumors based on imaging, leading to better outcome predictions. Comparing the imaging statistics of the tumor in all sequences within a defined probing area forms a digital biopsy of the data presenting an imaging footprint of the tumor and its surroundings. With this feature, colleagues who work in machine learning mentioned they could use the feature to find out which imaging sequences are the most important for developing automatic tumor segmentation methods. For this reason, this contribution also addresses multiple audiences, which we will discuss further in Section 3.3.

Besides analyzing the data within one patient, the probing feature within ParaGlyder can also be used to compare the imaging characteristics of tumors between two or more patients. However, this approach faces one challenge: MRI data is difficult to compare as data ranges vary more than CT data. Therefore, we decided to use pharmacokinetic parameter maps provided by our medical collaborators. These maps contain physiological parameters such as the blood flow within the tissue and make the data comparable. Our approach also works with standard modalities but makes the comparison between patients less accurate and more dependent on overall shape differences within the radial boxplot and less on absolute values.

3.2.2 Multi-Parametric Tumor Analysis Platform

As mentioned in Section 3.1.3, we developed RadEx to enable the analysis of multiparametric tumor data across a cohort. In this section, we focus on the visual analytics platform we developed. RadEx is presented in Figure 3.6 and consists of three major parts. A scatterplot presents every patient as a circle with a glyph representing the patient's tumor. A slice view is shown as a tooltip when hovering one of the patient's circles, including additional slice views for every imaging sequence available. While hovering over the patient's circle, the slice views are scrollable enabling viewing the whole image stack. At the bottom center, we present all available clinical parameters in either a unit chart or a parallel coordinate plot, depending on whether it is categorical or numerical data. We have a settings section on the right that allows the user to make

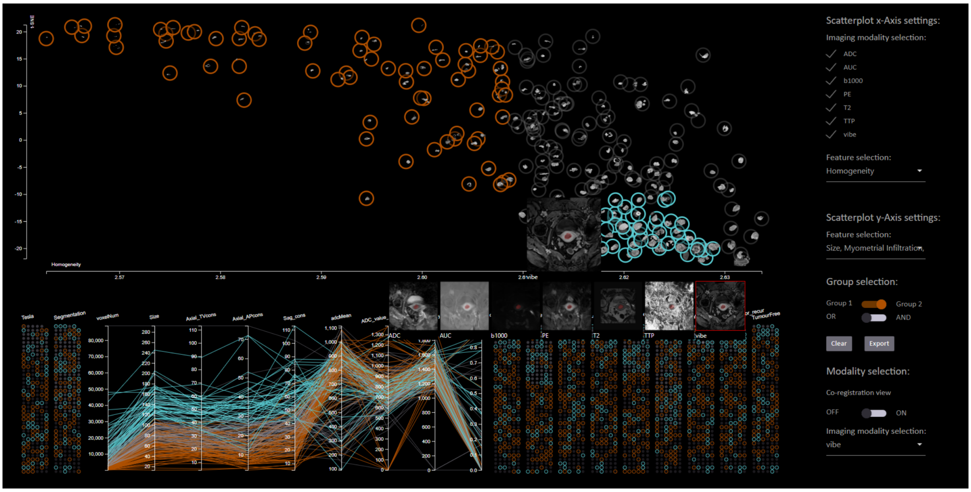

Scatterplot x-Axis settings:

Imaging modality selection:

- ✓ ADC
- ✓ AUC
- ✓ b1000
- ✓ PE
- ✓ T2
- ✓ TTP
- ✓ vibe

Feature selection:
Homogeneity

Scatterplot y-Axis settings:

Feature selection:
Size, Myometrial Infiltration,

Group selection:

Group 1 ● Group 2
OR ○ AND

Clear Export

Modality selection:

Co-registration view

OFF ○ ON

Imaging modality selection:
vibe

Figure 3.6: RadEx consists of three different components: the scatterplot (center top), the parameter visualizations (bottom row), and the settings section (right panel). The scatterplot represents all patients in the cohort visualized by a circle. Within these, a glyph represents the tumor of the patient. All important clinical variables are visualized in either unit charts or a parallel coordinate plot. On the right relevant settings and options are available.

group selections or change the scatterplot settings. RadEx allows group selections in every visualization to enable hypothesis formulation and checks. One novel feature which has not been used before in similar analysis platforms is the implementation of a 1D dimensionality reduction used for one axis of the scatterplot. This allows the radiologists to use multiple radiomic tumor features of interest within one scatterplot axis. In contrast, the other axis can be used to show, e.g., clinical outcome-related parameters of interest. Compared to a 2D dimensionality reduction, the benefit of our approach is that one axis still has clinical meaning and a meaningful value range.

One hypothesis that the medical experts involved in the development of RadEx wanted to investigate within our application was the possible connection between the homogeneity of the tumor and the survival of the patient. The two parameters are believed to correlate with each other. One exciting aspect of this analysis is that many parameters measure the homogeneity of tissue data within radiomic tumor profiling. Therefore, we decided to implement an animated transitioning method when selecting the parameters to present on the scatterplot axes. One example is presented in Figure 3.7, the change of the data is animated between states A and B. The radiologists can now define groups of interest and see how the data changes when analyzing different homogeneity measurements. Overall, RadEx received positive feedback from the medical experts involved in evaluating the application. As we further discuss in Section 4.2, it is currently being used in research practice at another institute.

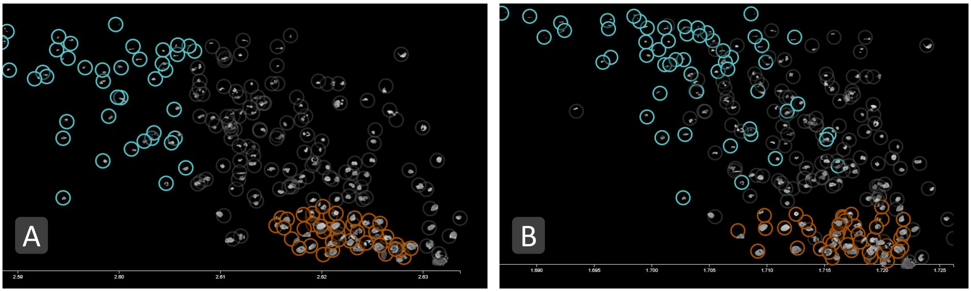

Figure 3.7: Changes of the scatterplot axes within RadEx are animated, and in this example, A represents the state before and B after the change. We animate the transition to simplify the perception of how the data changed from A to B.

3.3 Multi-Audience

Medical applications are often developed for a very targeted specialized audience, in our case, medical researchers. Nevertheless, this does not always have to be the case, as we show with our application ParaGlyder, which further targets machine learning experts. By involving experts of each specific field in the development of the application, we were able to provide a helpful tool, as demonstrated in the evaluation Section A.6 of ParaGlyder. Furthermore, when communicating medical or general research results, the target audience ranges from science peers to the general public, making the task even more challenging. Paper C presents ScrollyVis to fulfill these communication goals.

3.3.1 Similarity Visualization for Parameter Exploration

We introduced the similarity visualization in ParaGlyder to show areas of interest for different target audiences. The similarity visualization is based on the probing mechanism we discussed in Section A.4. The user first defines an area of interest, for example, within tumor tissue. We then define a new volume with the same extent as the imaging volumes given, and for each voxel within this volume, we calculate how similar it is to the probed area. Next, we take the multi-parametric Euclidean distance between the original voxel and the mean value within the probed area. This results in a new volume where each voxel value reflects how similar it is in the multi-sequential data space to the multi-sequential probing area of the user. Finally, we present the similarity volume by coloring it with the Viridis colormap, as we did before with the Stixels 3.1.1. Several results of this approach can be seen in Figure 3.8.

The radiologists use this feature to analyze the tumor's shape and extent. By placing a sphere within the tumor, the similarity visualization highlights all similar regions in the images, such as the whole tumor, presented in part A of Figure 3.8. Furthermore, cancer spread, also known as metastases, can share the same imaging characteristics as the tumor leading to highlighting in the similarity visualization as shown in Figure 3.8 A.

Machine learning experts who research how to segment tumors automatically are interested in finding the sequences with the most prominent information about the tu-

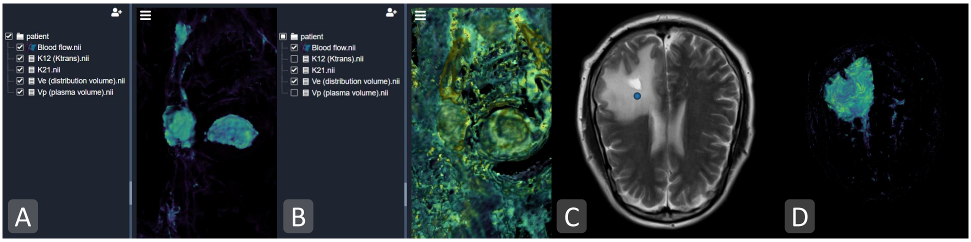

A

☑ patient
- ☑ Blood flow.nii
- ☑ K12 (Ktrans).nii
- ☑ K21.nii
- ☑ Ve (distribution volume).nii
- ☑ Vp (plasma volume).nii

B

☐ patient
- ☑ Blood flow.nii
- ☐ K12 (Ktrans).nii
- ☑ K21.nii
- ☑ Ve (distribution volume).nii
- ☐ Vp (plasma volume).nii

Figure 3.8: The similarity visualization in ParaGlyder enables different target audiences the analysis of their tissues of interest. For example, radiologists can analyze the tumor and cancer spread (A). Machine learning experts can use the same functionality to determine which imaging sequences best describe the tumor (B). Users can select which imaging sequences are part of the similarity calculation. An image presented in B will be displayed when selecting those who do not provide sufficient information to describe the tumor. Finally, neurological cancer imaging experts may find inflammatory regions within the brain after tumor removal surgery (C) and (D). Inflammation occurs when the body sends out first responder cells to cure an injury or to fight agents like viruses.

mor. Therefore, we enable them to select and de-select which sequences should be considered for the similarity calculation. Using only the essential sequences in the calculation, the tumor is still nicely visible. However, when too few sequences or sequences that do not describe the tumor sufficiently are selected, a result as visible in Figure 3.8B is displayed. In this example, the similarity visualization is not exclusively highlighting the tumor anymore as there is not enough information to distinguish it from the surrounding image information.

Neurological tumor imaging experts are, besides studying brain tumors, interested in visualizing inflammatory regions within the brain, which can occur after surgery to remove the tumor. We present such a case in Figure 3.8 C and D. In this case, the similarity visualization has the benefit of being able to differentiate between data representing the skull and data representing the inflammatory region when placing the probing sphere in the region of interest. By using the data of all sequences, differentiation is possible.

3.3.2 Scientific Scrollytelling for Diverse Audiences

With ScrollyVis (paper D), we introduce a novel approach for authoring, editing, and presenting data-driven scientific narratives. Our system connects story authors and story viewers. As shown in Figure 3.9, we enable authors to edit their stories in our interactive ScrollyVis story editor. The editor, in turn, creates an XML representation of the created story and hands it over to the ScrollyVis compilation pipeline. This pipeline provides a ready-to-deploy web page that can finally be presented in any standard web browser. Furthermore, the generated code enables dynamic stories where the viewers can decide which path they would like to traverse. Therefore, our system allows fast, incremental, and reversible scrollytelling experiences. It also provides the opportunity of building dynamic narratives where the viewer is involved in story traversing decisions.

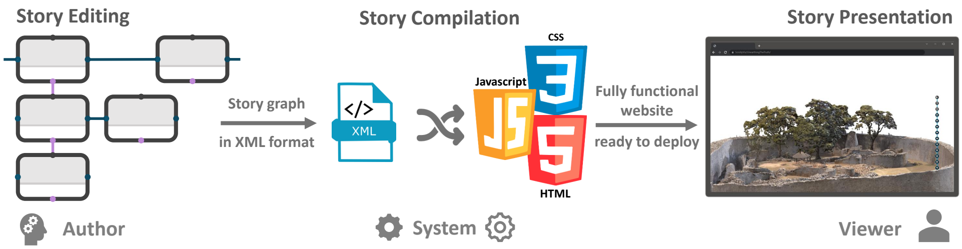

Figure 3.9: ScrollyVis consists of three parts namely, story editing, story compilation and story presentation. Authors edit their story, afterwards the story gets compiled within the ScrollyVis system to a ready-to-deploy website, and finally the website is presented in any standard web browser to the story viewer.

Story Editing

With ScrollyVis, we enable data types that have not been used before in scrollytelling editors and allow for rapid and intuitive creation of stories. The editor interface is depicted in Figure 3.10. Novel media types include volume visualization, slice visualization, and surface representations. In comparison to traditional node-link diagrams as also used by well-known editors such as *Twine* [161], our editor presents the story author with previews wherever possible. This is reminiscent of a storyboard design where authors would sketch scenes before they are finally filmed, animated, or programmed. The previews are also interactive, e.g., a surface visualization can be turned and zoomed by the author to define the viewport they would like to present interactively. All parameter modifications update the previews on the fly, and therefore the author always knows what the result on the deployed website will look like.

By defining different pathways within the editor, namely a main path and a sub-path, we allow for intuitive layering of content. Content replaces preceding content while traversing the main path (the blue connection between nodes see Figure 3.10). Connecting nodes on the sub-path presents content on top of the already displayed content. Story content can now be presented dynamically without complex linking interactions within the editor. One additional key aspect unique to ScrollyVis is the integration of decision nodes where the story path can split into several different directions. This concept has not been shown in combination with scrollytelling experiences before. Within the editor, a decision node allows the author to define a story splitting point. One path leads into such a decision node, but several paths emerge, depending on the number of answer possibilities. A split path can stay split and present different stories in parallel; the path can join a common path at some point or lead to dead ends.

Story Compilation

After the editing, we introduce the story compilation, split into two different steps. The first step deals with the story contents, and the second step with the story transitions. At the end of the story compilation step, the website is ready-to-deploy in HTML, CSS, and Javascript code. This allows for further editing of web developers and customization if desired. Our compilation phase is unique because the websites we create are

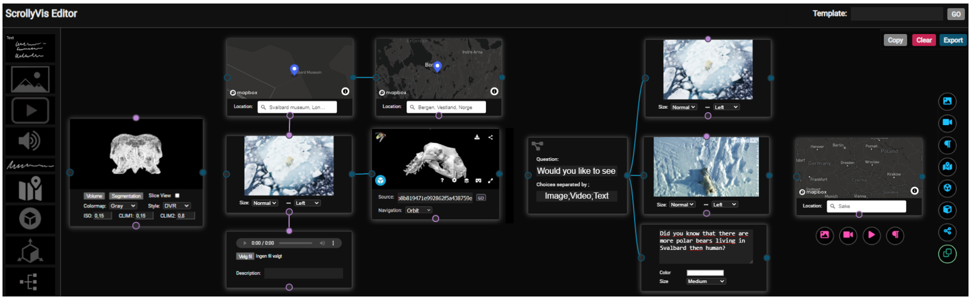

Template:
Slice View ☐
Colormap: Gray
Style: DVR
ISO: 0,15
CLIM1: 0,15
CLIM2: 0,8
Location: Svalbard museum, Lon...
Location: Bergen, Vestland, Norge
Size: Normal
Left
0:00 / 0:00
Ingen fil valgt
Description:
Source: b8b819471e992862f5a438759e
Navigation: Orbit
Question:
Would you like to see
Choices separated by ;
Image;Video;Text
Did you know that there are more polar bears living in Svalbard then human?
Color
Size Medium
Location: Søke

Figure 3.10: Screenshot of the ScrollyVis editor. Various media formats are supported and can be layered to combine them. Decision nodes allow the story flow to split into different paths and a context menu allows for quick and intuitive story editing.

not static. The story flow defined within the created website is dependent on the story viewer interaction. If a story contains decision nodes, the traversal of the representing code must be flexible as well. We solved this by defining the story traversal on the fly within the code in real-time while the story viewer consumes the narrative. Based on the story graph defined by the author, we create code that handles the transitioning between the contents. When traversing on the main path, previous nodes will be blended out, and new ones will be blended in. While traversing on the sub-path previous nodes will not be blended out, and new content will be presented on top of the old one.

Story Presentation

Finally, we can present the created stories in any standard web browser. ScrollyVis enables the traversing of dynamic narratives and presents the story viewer with an abstract visualization of the story to keep track of the traversed path. With ScrollyVis, we introduce story node transitions based on the content presented. We feature adaptive transition methods for complex content like map views, direct volume rendering, slice views, and 3D surface visualizations, including camera movement, parameter adaption, and slice index changes. This allows for the easy and intuitive creation of animations by just defining critical points within the editor. In addition, we use opacity change as a transition method for standard content like images and videos. Finally, we enable dynamic narratives by defining successor story nodes on the fly while the story viewer consumes the story.

Kosara [88] discussed that it is vital for scrollytelling viewers to keep in mind the overall story length and where within the story they are currently. Therefore, we introduce an abstract depiction of the overall story within the final website where the viewer can see the current position within the story and the decisions they have made. The story tree view also replaces the standard scrollbar as it does not have any meaning within dynamic scrollytelling experiences. We present an example in Figure 3.11. Left, the story within the editor is presented, and right, the tree on the final website is shown.

With ScrollyVis, we introduce an editor for dynamic scientific scrollytelling, which allows experts from various fields to tell their research questions or outcomes in an appealing way to target other scientists and reach the public. Furthermore, ScrollyVis

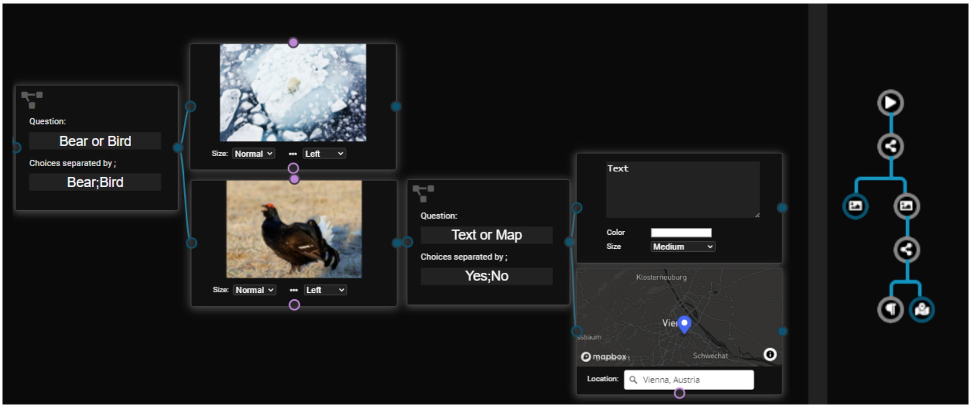

Question: Bear or Bird

Choices separated by ; Bear;Bird

Size: Normal

Left

Size: Normal

Left

Question: Text or Map

Choices separated by ; Yes;No

Text

Color

Size Medium

Location: Vienna, Austria

Figure 3.11: Example showing the story tree view of a created story. On the left the story in the ScrollyVis editor is presented and, on the right, we show the resulting story tree on the final website.

can also be an interesting application for scientists and everybody who would like to tell a story on the web.

4

«Erfolg ist kein Glück sondern nur das Ergebnis von Blut, Schweiß und Tränen ...»
«Success is not luck it is simply the results of blood sweat and tears ...»
Kontra K

4

Chapter 4

Demonstration and Use in Practice

With our work, we contributed to the field of medical visualization and focused mainly on application-driven research. In this chapter, we reflect on using our tools in practice and the steps we carried out to get them there. Furthermore, we discuss how specific techniques we introduced are planned to be used and evaluated in the future.

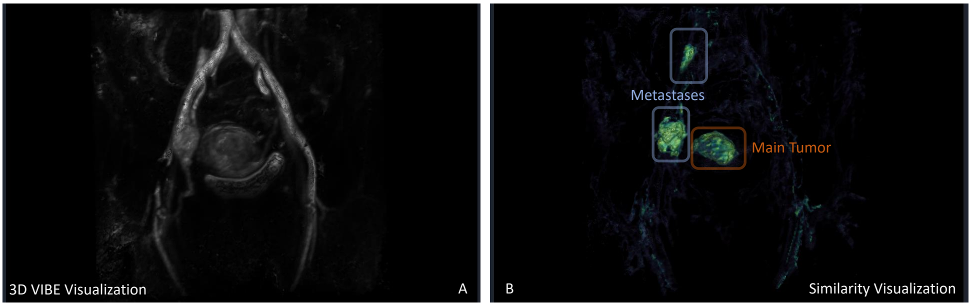

Figure 4.1: Our introduced similarity visualization has the potential to highlight the primary tumor as well as metastases within the same imaging volume. On the right, the original 3D visualization is shown. After probing the data with our probing interaction, the right visualization presents the similarity of each voxel within the volume to the probed area using the Viridis colormap and opacity adjustments.

4.1 Similarity Visualization

Paper A (ParaGlyder) introduced two different visualization methods to work with multi-modal medical imaging data. Our clinical collaborators were especially interested in the similarity visualization. We identified cases where the similarity visualization revealed the tumor and the associated metastases in the same MRI image. Our collaborators told us that they had never seen such a clear representation of metastases and the tumor before, without prior segmentation of the structures. An example is shown in Figure 4.1. Our collaborators are interested in exploring the possible impact further in a cohort of patients. We received similar feedback from tumor imaging researchers from the Netherlands Cancer Institute (NKI) when presenting the application at their facility.

To evaluate the impact of our approach, we are planning to conduct a study on similarity visualization and its potential for clinical impact. The study is planned at the

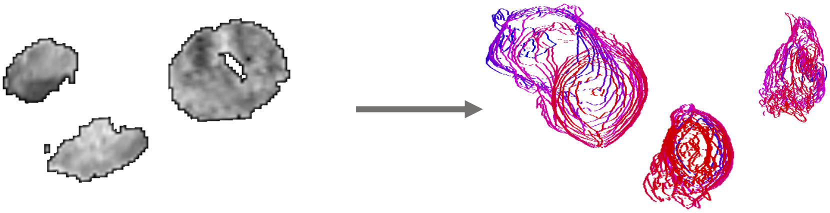

Figure 4.2: We improved the original tumor glyph design used in RadEx with a new design visible on the right side of the image. We create the new glyph by taking the outline of the tumor data across all slices, coloring it with the plasma colormap by the depth within the volume, and stacking the lines. This allows for complete tumor visualization and visualizes the size and shape of the tumor.

Mohn Medical Imaging and Visualization (MMIV) center in Bergen, Norway. It will include several patients with metastases in the same imaging volume as the primary tumor and several patients without. The goal is to define the extent to which similarity visualization can consistently highlight metastases in the same imaging volume as the primary tumor across patients. Therefore, several radiologists will be invited to use the tool and afterward define if there are metastases in the same imaging volume, where they are located, and how many are present. The results will then be compared to the initial reading of the same imaging volume and with the outcomes of surgery which is the standard procedure for endometrial cancer patients.

The potential impact of the similarity visualization is manifold. We see the most significant potential in highlighting metastases close to the tumor. If such metastases are present, they must be considered during surgery, and a biopsy is typically carried out. Currently, small metastases are hard to see on the MRI images as the resolution is also limited. Typically, PET-CT imaging is used to define if the tumor has spread and confirm metastasis throughout the whole body. This imaging technique uses radiation and may be harmful to the body if carried out too often. Metastases in lymph nodes near the tumor indicate a spread throughout the whole body. In case the similarity visualization can be used to define if there is a spread in the close circumference of the tumor, this could also be used as an indicator if a PET-CT scan is necessary or not. Our collaborators generally see great potential in the similarity visualization shown in paper A, ParaGlyder, and plan to evaluate the potential further.

4.2 Integrated Visual Exploration of Multiparametric Studies

We developed the application RadEx (paper B) in close collaboration with experts in cancer imaging research at the MMIV in Bergen, Norway. The application was tailored to their needs and expectations and fulfilled the requirements we defined together. This is also reflected in the results of the user evaluation we conducted together with the experts involved in developing the application. We involved four additional experts in evaluating the usefulness of our approach and received positive feedback. Applications developed for a targeted use case, endometrial cancer imaging research, are often difficult to apply within other institutes or other datasets as they are within a Ph.D.

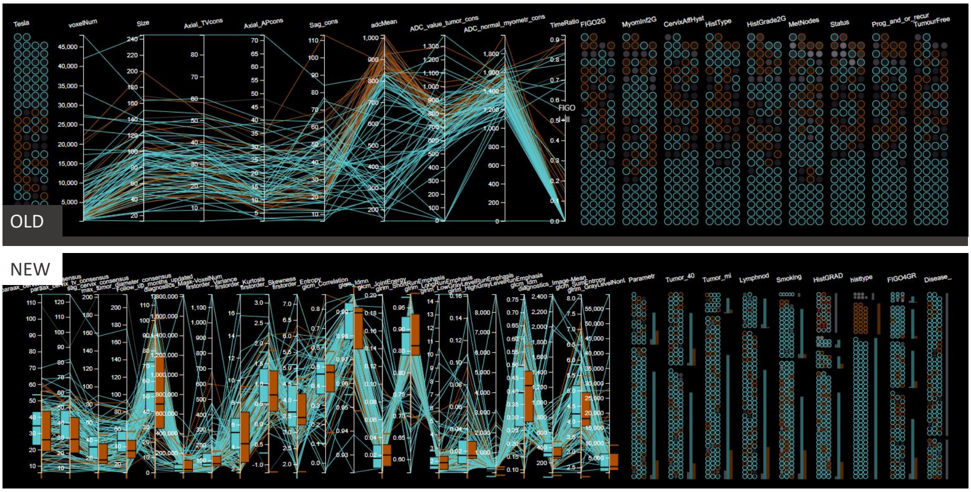

Figure 4.3: The parallel coordinate plot and unit chart visualizations have been extended by split box plot visualizations and bar charts to summarize group information. This allows for effective group comparison while still retaining the same interaction possibilities as before.

project often developed as a prototype application serving specific needs. In the case of RadEx, radiologists at the MMIV researching cervical cancer and researchers of the Netherlands Cancer Institute (NKI) working on prostate cancer data also showed interest in working with RadEx. We assessed the customization needed to make it ready to work with the new data and at another facility and implemented them. This section discusses which steps were necessary to make the tool usable for new use cases.

In general, the analysis workflow in RadEx is split into two phases: the data quality check and the search and query phase. Phase one starts with pre-processing the data. In this phase, radiomic tumor profiling data is calculated, and a tumor thumbnail image and glyph are created. We adapted the glyphs to be more representative and show more information about the tumor. In Figure 4.2 the original version is visible on the left and the new design on the right. The first glyph design represented a slice of the tumor in one of the MRI sequences at the slice position where the tumor has the largest extent. This represents the shape and the size of the tumor. As we have the 3D data available and the shape of the cancer is unique for each patient, we introduce a new glyph design which is calculated as follows: 1) The outline of the tumor is calculated in each slice, 2) the outlines are colored using the magma colormap which represents the depth within the volume and finally 3) the outlines are stacked on top of each other. This glyph design allows them to see the tumor shape and size in 3D and delivers more contextual information than the initial glyph design. The pre-processing pipeline worked without changes for both the cervical and prostate cancer data to accommodate different data.

After adapting the pre-processing phase, we also made the visual analytics application more flexible to suit the use cases of the NKI and other institutes. As a first step, we increased the flexibility of axis data selection. We now enable the selection of dimensionality reduction for each axis independently or combined as a 2D dimensionality reduction. Furthermore, both axes can also show single parameters. This supports a broader range of analysis tasks and better suits the needs of different radiologists.

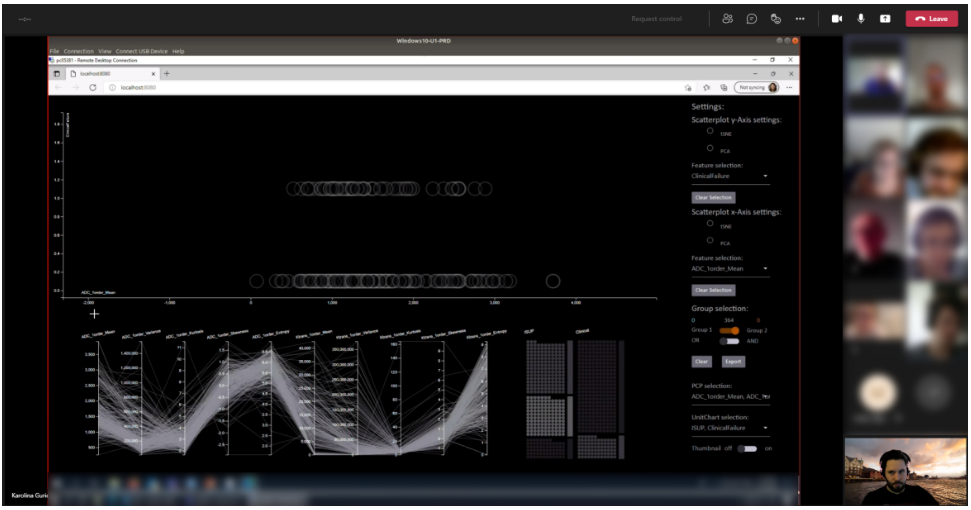

4

Figure 4.4: During a meeting with the Netherlands Cancer Institute (NKI), the improved version of RadEx was used in a live setup. A published hypothesis has been validated on the data of the NKI. The researchers of the NKI found the application effective and easy to use and will use the tool in the future to work with their prostate cancer data.

In addition to the scatterplot, we also adapted the bottom section of the application, namely the parallel coordinate plot and the unit chart section. First, we enable the users to select which data they want to see in these sections. For our initial use case, data availability was limited. At the same time, we were able to show all available data. However, too much data was available for the NKI use case to visualize effectively without scrolling. To address this, we implemented data selection settings that differentiate automatically between categorical and numerical data. In addition, we integrated boxplots in the parallel coordinate plot to enhance group comparison and added bar charts to the unit charts. The results are visualized in Figure 4.3. These features are mentioned as future work in the original publication as our medical collaborators at MMIV already expressed interest in improved group comparison capabilities.

The researchers at NKI currently use RadEx to analyze prostate cancer data and aim to find new imaging biomarkers to include in the prostate cancer diagnosis staging process. One indicator of the usability and usefulness of our application was the live usage of the tool during a cancer imaging research meeting at NKI to test a hypothesis. Figure 4.4 shows a screenshot of this meeting.

4.3 Co-Registration Validation

Radiomic tumor profiling, as discussed in RadEx, is shown to be an effective indicator for survival in both cervical and endometrial cancer imaging [41, 62]. To perform such an analysis, the tumor must be segmented. This process can be performed either manually or with machine learning algorithms. In most cases, the segmentation is only available on one modality since a manual segmentation is rather time-consuming,

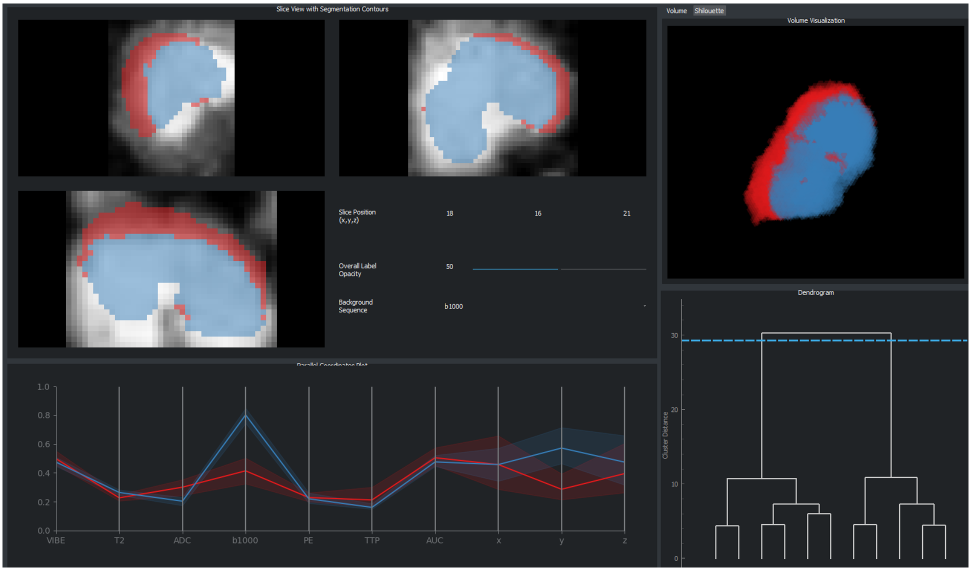

Slice Position (x,y,z) 18 16 21

Overall Label Opacity 50

Background Sequence b1000

Figure 4.5: ICEVis is primarily designed for the analysis of tumor sub-parts that might be linked to interesting regions within the tumor. However, this case shows that it can also find incorrect co-registrations, as seen in the parallel coordinate plot and the imaging data. The two clusters differ significantly in one of the sequences and one of the coordinate axes. Furthermore, in the slice view, the viewer can see that the segmentation mask is misplaced.

and the machine learning algorithms are, in their current state, not well suited for a transfer learning approach. As the radiomic tumor profiling data of more than one sequence might be necessary to find predictive patterns, the medical images must be co-registered [62]. Co-registration aligns the images, so they perfectly fit on top of each other. This enables the usage of only one tumor mask for several co-registered imaging sequences or modalities.

Co-registration is not an easy task to carry out automatically as the medical sequences differ in contrast and image content. There are many automatic methods to solve the problem, but they are often not working correctly or deliver unsatisfactory results. A landmark-based method, including manual steps, can be used in such cases. However, before putting the human in the loop, one must determine if the co-registration has gone wrong. This can again be performed manually by visually inspecting the results or automatically. There are different strategies for finding cases where the co-registration has gone wrong, and some work better than others. We can also identify co-registration issues with our ICEVis application (paper D). Besides defining interesting pathological patterns within the tumor, ICEVis shows a clear pattern in the parallel coordinate plot of the given multi-sequential data in cases where the co-registration did not work. We present such a result in Figure 4.5. This pattern inspired us to tackle the co-registration problem in a follow-up project we are currently working on. We will use the segmentation mask aligned with one of the available sequences as a landmark source and align it with the other imaging sequences by rotating and translating it. This allows us to perform a landmark-based co-registration without the medical expert needing to place landmarks on medical images, which is an untypical work for them.

After discussing with our medical collaborators, we know that movements during the imaging process cause the images not to be perfectly aligned. This includes moving the body within the scanner, inducing translation and rotations, bowel movement, and bladder filling during the investigation, further shifting structures of interest. However, as the tumor is a relatively hard tissue, the tumor itself will not be deformed significantly. Keeping this in mind, we know that the tumor segmentation mask does not have to be changed in its shape to be applied to other modalities but must be shifted and rotated. We follow up on this hypothesis in future work and hope to find a satisfactory solution across imaging sequences.

4.4 Dynamic Narratives for Scientific Scrollytelling

Our research on dynamic narratives for scientific scrollytelling is practically-oriented work. We aimed to enable researchers from various areas of expertise to tell stories about their science. We designed ScrollyVis as an abstract extensible system and realized it with our ScrollyVis editor. ScrollyVis can facilitate the creation of a wide variety of stories in various domains. In Figure 4.6 we present snippets of different stories we created together with experts in their respective fields. For example, as shown in Figure 4.6 A, we created a story with an expert in anatomical education research. The group she is working with has surface scans of various bones uploaded to Sketchfab, a platform providing surface rendering capabilities and annotations. Our approach also delivers the possibility of creating a dynamic story about the data. The professor can add auditive narrations for specific bones they would like to describe in detail. Furthermore, a quiz about specific parts of the bone can also be included. The expert creating the story with us will use ScrollyVis to analyze further the differences in learning outcomes when using a free exploration form as Sketchfab provides or a guided scrollytelling approach as we introduce ScrollyVis.

Figure 4.6 B presents a story that we created together with an osteological expert who works at the natural history museum in Bergen. The story's topic is the polar bear. Specific about this story is the usage of direct volume rendering techniques, making it unique and different from other scrollytelling approaches available on the web. With ScrollyVis, storytellers can quickly create direct volume renderings, switching between a 3D visualization of the data or a slice view-based one. The expert particularly liked the possibility of using stories created with ScrollyVis to accompany museum exhibitions where visitors can share their experiences with others and gather more insight into the skulls while visiting the museum. Another significant benefit is that creating a story with ScrollyVis does not involve multiple people like developers and designers. Therefore, our approach is cost-efficient as the scientist can already create appealing stories alone.

In part C of Figure 4.6, we present two snippets of a story we created together with a meteorological visualization expert. Hurricane data is presented on the left, and there is hail data on the right. This shows how versatile our approach is as our expert provided the data, and we adapted the map view of ScrollyVis specifically to present their data. The collaborating expert plans to use ScrollyVis to research communication methods and their effect on the public.

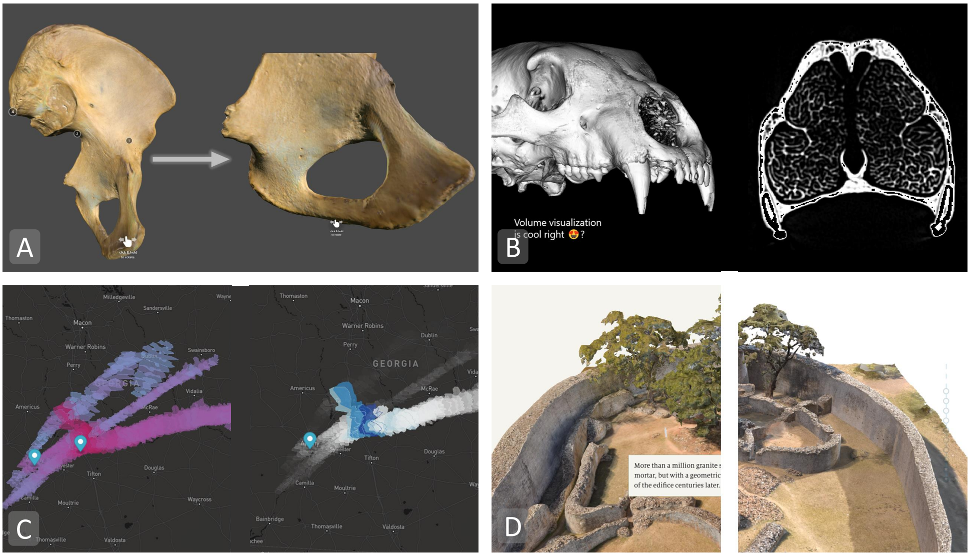

Figure 4.6: Four different story examples we created together with experts of the various fields. A presents a story about anatomical education facilitating the animated camera transition between different views. B is a story about the polar bear by an expert in osteological research and presents a 3D iso-surface visualization of the bear's skull and a slice view of the same skull. C presents two snippets of a story we created with a meteorological visualization expert presenting hurricane and hail data. Finally, D presents a story by the Economist about archaeological science where we show that our editor can reproduce the story without the need for a team of developers.

Medical Narratives

Medical narrative visualization has gained interest in recent times, and target audiences of such narratives include patients, their relatives, and people interested in medicine [102]. After presenting ScrollyVis on different occasions, two different collaborations to use and further develop the features of the ScrollyVis editor and the story compilation were established. Two example stories that will be created using ScrollyVis are presented in Figure 4.7.

The first story is called: "Is there a tornado in Alex's blood?" developed by a master's student in close collaboration with the Otto von Guericke University in Magdeburg, Germany, and the University Leipzig center for heart-related diseases. The story is about developing an aneurism in the bloodstream of an artificial patient called Alex. The story's target audience consists of patients and people who would like to learn more about the disease, the factors that lead to the formation of an aneurism, and prevention. In creating the story, the ScrollyVis editor will be extended with additional functionalities, including creating slide-like presentations within a scrollytelling-based presentation format. The story will be presented at the "Lange Nacht der Forschung" in Leipzig, Germany, and the educational effect of the story will be evaluated during the exhibition. This demonstrates that ScrollyVis might impact patient and public communication and serve research and science dissemination purposes.

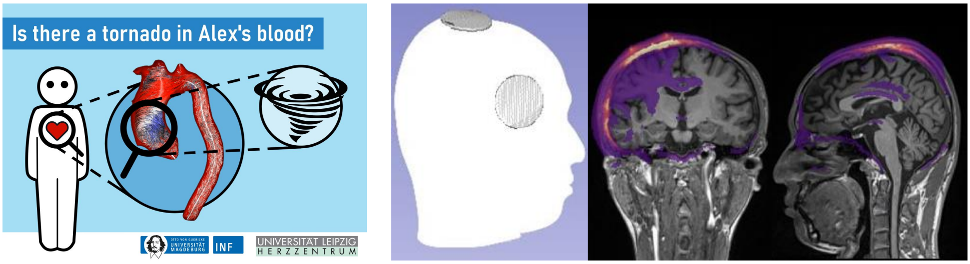

Figure 4.7: Two stories about common diseases which aim to teach a general audience. Both stories include medical visualization and facilitate guided dynamic narratives with the goal of increasing learning outcome over basic non interactive story design.

The second story presented in Figure 4.7 focuses on electroconvulsive therapy (ECT) used in cases of severe depression and is developed by a master's student. ECT therapy is an active research area at the MMIV. It is an effective but controversial therapy where electrical currents are applied to some regions of the brain to relieve patients of severe depression [14, 116, 166]. To enable patients to form an informed decision about this treatment, Sofia is creating a story about the therapy in close collaboration with researchers from the MMIV. She will evaluate how different story components affect the experience of the story viewers. The research goal is to find the most effective ways to convey information about the therapy, so readers are adequately informed but not misled.

5

«You look at where you're going and where you are and it never makes sense, but then you look back at where you've been and a pattern seems to emerge.»
Robert M. Pirsig

Chapter 5

Conclusion and Future Work

We scaled up medical visualization in multi-modal, multi-patient, and multi-audience dimensions with our work. We contributed three design studies focusing on different target areas of medical visualization and one interaction and presentation technique for scientific outreach. Visualizing multi-modal imaging data with more than three dimensions without dimensionality reduction is still a challenging problem and an open research field. With our Stixels, we contributed a visualization technique based on small multiples and radar charts. Our Stixels enable researchers to view the imaging data of several modalities, in our case, seven on one slice, without juxtaposed views. This allows pattern analysis to find heterogeneous and homogeneous regions within a tumor and inform treatment decisions. We further contribute an interactive probing interaction that facilitates a digital biopsy of the imaging data and can be used in a similarity visualization highlighting tumor tissue and metastases in the same imaging volume. The same feature allows for examining inflammatory brain regions and machine learning experts to evaluate the contribution of different imaging modalities to the segmentation of tissue belonging to the tumor.

Next, we contributed an interactive analysis platform for radiomic tumor profiling data on a cohort level consisting of several linked views. Using a 1D t-SNE allows analyzing the high dimensional radiomic tumor profiling data while still having one axis of the scatterplot free to present single clinical variables. This allows for interactive and practical hypothesis generation and hypothesis validation within a single application. We further retain the imaging data context by allowing radiology researchers to display every individual tumor in all available imaging modalities on demand. Our application RadEx was further adapted to not only be used for endometrial cancer but also for cervical cancer and in another research institute in the Netherlands for prostate cancer. This shows us the practical relevance of our approach. In addition, we focused on sub-tumor analysis using all available imaging modalities of tumor regions within the endometrium in a hierarchical clustering approach. The first results show good alignment between discovered clusters and interesting histopathology regions, such as a necrotic core within the tumor and a potential inflammatory area at the tumor border. Our final contribution proposes a data-driven narrative editor. It allows for the efficient creation of dynamic narratives to allow for immersive scientific scrollytelling on the web by employing a no-code paradigm. ScrollyVis allows scientists and the public to create stories about any topic of interest and creates a ready-to-deploy scrollytelling experience without writing code. The presentation of ScrollyVis on different occasions showed us that there is a great interest in using such tools. Purposes include anatomy education, education about meteorological phenomena, osteological research dissemination, medical visualization narratives, and the creation of stories about research results to a general and specific audience.

This thesis delivers a new answer to visualizing multi-modal spatial data for joint exploration and analysis. We introduced a novel imaging probing interaction helping to find global patterns, especially interesting when applied to cancer imaging data. Furthermore, the novel application of a 1D t-SNE dimensionality reduction allows presenting several relevant imaging markers on one scatterplot axis while the other axis still retains a meaningful representation of clinical parameters. With our ScrollyVis contributions, we combined scrollytelling and scientific storytelling. We provided a powerful editor for scientists to distribute their work in an intuitive and accessible way to their peers and the general public. The work in this thesis revealed that there are more questions to answer than before we started with our research. Therefore, in this section of the thesis, we elaborate on potential future avenues that could potentially impact the work of the visualization community.

First, the similarity measurement and visualization we used in ParaGlyder could potentially help define a general measurement of similarity between different patients. Besides radiologists who examine the imaging data of a tumor, general practitioners could benefit from comparing patient histories to form informed decisions about the treatment of the patients and the expected quality of life after treatment. Comparing patients on a cohort scale helps find intriguing patterns and brings the insights back to the level of an individual patient. Cohort analysis is often done in a research context, and after publication, there are opportunities for clinical impact by revising treatment guidelines. More detailed predictions about possible treatment outcomes could be formed by directly comparing patients with similar tumors within the same hospital, geographical region, and perhaps even of similar age.

In recent research, global tumor characteristics are an effective measure for risk stratification [60]. ICEVis (paper D) shows a potential use case of applying hierarchical clustering on multi-parametric medical imaging data and finding potentially interesting regions within the tumor, such as a necrotic core of inflammatory regions. Furthermore, by correlating the clusters with the result of a microscopic investigation of the removed tumor, we could validate if our findings correlate in the future. This could aid more personalized treatment strategies and influence either to which extent the tumor will be removed or, in the case of cervical cancer or prostate cancer, influence treatment decisions if surgery is performed or not. Furthermore, patterns found with hierarchical clustering could be analyzed on a cohort level and might be correlated to patient survival.

Radiomic tumor profiling is a promising predictor of survival in both endometrial cancers and cervical cancer patients [41, 62]. Predictive imaging biomarkers have been discovered by analyzing whole tumor characteristics based on first- and second-order statistics. With RadEx, we contributed to interactive visual analysis of radiomic tumor profiling information widely employed in cancer research. In addition to calculating these parameters, these textural features should also be visualized concerning the patient's imaging and clinical data. The radiologist should be able to see the textural features of the tumor beyond reading one value derived by the radiomic tumor profiling analysis.

As the last avenue of future work, scientific storytelling is still a wide-open field with increased attention to visualization for communication in recent times. With ScrollyVis, we present an essential step towards making dynamic storytelling in a scrollytelling form more widely available. In the future, we would like to enhance our editor further to make it more flexible and generally applicable. Examples would be the free placement of elements within the webpage and combined previews of nodes. Every scientist should be able to tell a story about their research for different audiences like their peers up to a vast audience such as the general public. The recent pandemic showed us how vital science communication is to present different concepts behind scientific findings. Making research outcomes available in an engaging and meaningful way is just the beginning of this field of research.

«The way to get started is to quit talking and begin doing.»
Walt Disney

Part II

Included Papers

Paper A

ParaGlyder: Probe-driven Interactive Visual Analysis for Multiparametric Medical Imaging Data

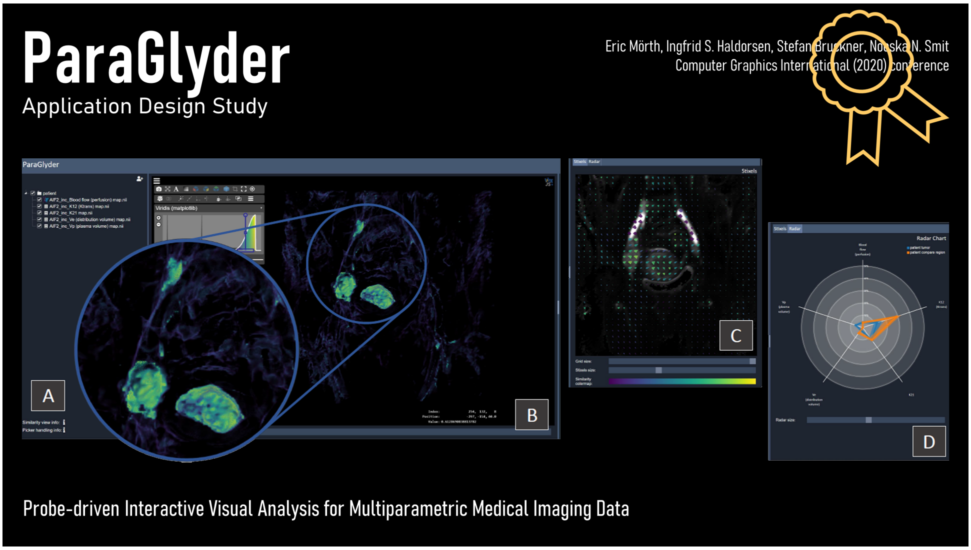

Eric Mörth[1,2], Ingfrid S. Haldorsen[2,3], Stefan Bruckner[1,2] and Noeska N. Smit[1,2]

[1]Department of Informatics, University of Bergen, Norway
[2]Mohn Medical Imaging and Visualization Centre, Department of Radiology, Haukeland University Hospital, Norway
[3]Department of Clinical Medicine, University of Bergen, Norway

---

This article was published in *Advances in Computer Graphics: 37th Computer Graphics International Conference, CGI 2020*, pages 351–363, 2020

Abstract

Multiparametric imaging in cancer has been shown to be useful for tumor detection and may also depict functional tumor characteristics relevant for clinical phenotypes. However, when confronted with datasets consisting of multiple values per voxel, traditional reading of the imaging series fails to capture complicated patterns. These patterns of potentially important imaging properties of the parameter space may be critical for the analysis, but standard approaches do not deliver sufficient details. Therefore, in this paper, we present an approach that aims to enable the exploration and analysis of such multiparametric studies using an interactive visual analysis application to remedy the trade-offs between details in the value domain and in spatial resolution. This may aid in the discrimination between healthy and cancerous tissue and potentially highlight metastases that evolved from the primary tumor. We conducted an evaluation with eleven domain experts from different fields of research to confirm the utility of our approach.

A.1 Introduction

Multiparametric medical imaging scans are commonly used in screening procedures and in targeted diagnostics. Basing decisions on the analysis of these datasets is not an easy task and often involves visual inspection of different juxtaposed representations [40]. Multiparametric datasets are generated in medical imaging by, e.g., Magnetic Resonance Imaging (MRI) scanners, by varying acquisition parameters resulting in imaging data with varying contrasts. In the analysis of medical imaging data, the main task is usually to identify discernible patterns to distinguish pathologic from healthy tissue, and identify, e.g., malignant tumors. The identification of metastases, likely to share characteristic imaging properties with the primary tumor, may be difficult to spot only using one modality, although identifying them at primary diagnostic work-up is essential to develop more tailored and targeted treatment strategies in various cancers. In order to improve the workflow of tumor diagnosis and metastases identification, we have developed a tool for analyzing multiparametric medical imaging data together with gynecological cancer, machine learning and neurological cancer research experts. By employing different views displaying multiparametric data at different levels of detail, we can present imaging data without having to visually compare several modalities in side-by-side views. We enable highlighting of target structures, based on multiparametric similarity, which was not possible before. Medical experts are used to working with 2D slice views. Overlaying multiparametric data on top of these views produces insights which are easy for them to put into a spatial context. Showing multiparametric images in one view reduces the cognitive load and allows the medical experts to see the relevant information at a glance. Our main contributions are the following: (1) We present visualizations that remedy the trade-offs between revealing details in the multiparametric value domain and spatial resolution by introducing a multiparametric star glyph map-based visualization. (2) We present an interactive analysis application primarily targeting cancer imaging, as well as additional workflows in different application areas. (3) We evaluate our system with eleven experts using

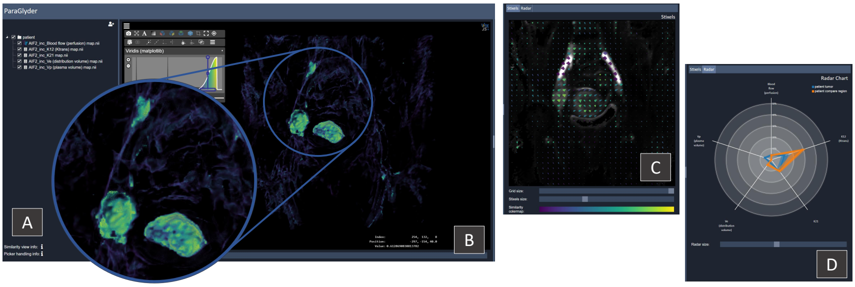

```
Index:         254, 132,   8
Position:     -297, -154, 40.0
Value: 0.6128690838813782
```

Figure A.1: The ParaGlyder prototype application, featuring a subject overview (A), central view (B), Stixels view (C), and radial boxplot view (D).

the System Usability Scale (SUS) [17] and a qualitative evaluation to demonstrate the utility of our approach.

Modern imaging techniques are routinely used at many centers in the preoperative diagnostic work-up in endometrial cancer. Imaging markers derived from these advanced MRI techniques have been shown to be linked to endometrial cancer subtype and stage [7, 40, 55–57]. According to previous findings, low tumor blood flow and a low rate constant for contrast agent intravasation, meaning the backflow of injected contrast into the close vessels, based on dynamic contrast-enhanced (DCE)-MRI, are associated with high-risk histologic subtypes and poor prognosis. Gathering information from parametric maps based on DCE-MRI is usually done using juxtaposed images of the same slice in the different modalities. These maps are derived from a single dynamic acquisition and are therefore co-registered by nature. Examining the images involves comparing the images mentally or by using a manually placed region of interest (ROI). If advanced imaging methods can be utilized to validly predict the aggressiveness of a tumor, this could lead to better risk-stratified treatment algorithms that may be beneficial for the patients. Less invasive treatment regimens may then be given in low-risk patients, and the more invasive treatments can be reserved for high-risk patients in whom the expected survival benefit justifies the increased side effects.

A.2 Related Work

Lawonn et al. [95] provide an extensive overview of different visualization techniques for multimodal medical imaging datasets. Gleicher et al. [49] introduced a taxonomy of visual comparison approaches and surveyed existing methods according it. Friendly et al. [43] proposed radial boxplots, as a means to visualize data variations. Ropinski et al. [133] provide a thorough overview of different glyph-based visualization techniques in the field of multivariate medical data visualization. Wickham et al. [169] introduced a visualization technique called glyph maps. Opach et al. [115] described that the effectiveness of polyline versus star glyphs is task-dependent. The effective combination of star glyphs presenting non-spatial data and geospatial data has been demonstrated by Friendly et al. [43] and more recently by Jäckle et al. [66]. In contrast to this, we use star glyphs to present an abstract version of multiparametric spatial data on top of spatial data. Smit et al. [146] presented a method to spatially query data by placing a sphere in a 3D view, and interaction techniques to effectively place spheres in volume renderings [145]. Bruckner et al. [18] introduced a probing tool for enabling visual queries. Mlejnek et al. [105] presented interactive glyphs for probing tissue characteristics in medical data. In contrast to these approaches, we provide a probing interaction that acts like a digital biopsy of our multiparametric medical imaging datasets. More closely related to our approach, Stoppel et al. [150] used small multiples to visualize spatio-temporal data in a spatial context. Malik et al. [101] introduced a comparative visualization technique that visualizes up to five modalities together in one view. Jönsson et al.[74] presented a visual environment for hypothesis generation using spatial and abstract data. In contrast to these related publications, our approach enables the exploration and analysis of multiparametric medical imaging datasets of more than five modalities. We provide targeted functionality for the analysis of pathology, which allow for inspection of the multiparametric imaging data in linked spatial and non-spatial data visualizations.

A.3 Requirement Analysis

Following the nested model for visualization design by Munzner [108], we characterized the problem domain. To meet the requirements and the demands of the target audience, we consulted experts in gynecological cancer imaging, neurological imaging, and machine learning. We identified application related challenges they face in their research practice. Cancer imaging is performed to assess tumors and metastases, in gynecological cancer imaging in the pelvic area and for neurological imaging in the brain. Cancerous tissue is discernible because it differs from its surrounding healthy tissue. Besides analysis of the extent and size of the tumor, analyzing different sub-regions within a tumor may be of interest. Finding abdominal lymph node metastases is a challenging task, as the metastases have variable size, ranging from a few millimeters to sizes exceeding the primary tumor. Metastases often share some of the characteristic imaging features of the primary tumor. Based on our analysis we present the following requirements for our interactive analysis application:

- R1: Visual analysis of multiparametric imaging data in a single view
- R2: Multiparametric inhomogeneity analysis
- R3: Comparing regions within the multiparametric imaging data
- R4: Comparing multiparametric imaging data of multiple subjects
- R5: Multiparametric similarity analysis based on a digital biopsy
- R6: Interactive parameter selection for automatic multiparametric segmentation tasks

When satisfying these requirements, we support gynecological imaging researchers, neurological imaging experts and machine learning experts in their research or clinical routine with the ultimate goal of improving patient care by providing better diagnostic tools that can guide tailored and individual treatment strategies.

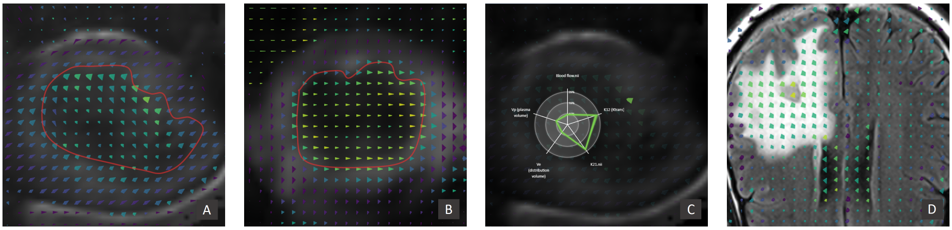

Figure A.3: The Stixels view reveals an inhomogeneous tumor in one subject (a) and a more homogeneous tumor in another subject (b). The outline in red shows the tumor extent for illustration purposes. A tooltip provides details on demand in a radial boxplot (c), The Stixels view reveals oedema in the brain after surgery (d).

A.4 ParaGlyder

In this section, we present our visualization and interaction design decisions based on the requirement analysis. In Figure A.2, we present the different components of our method and their interplay. Our design combines spatial and non-spatial visualizations, linked by a view combining a non-spatial visualization in spatial context. Our approach consists of several visualization and interaction methods for the interactive analysis of multiparametric data described in the following.

A.4.1 Data Processing

Our method relies on multidimensional co-registered volumetric data. Our gynecological cancer experts already deliver co-registered volumes due to the nature of the data source. Co-registration is therefore not part of our application but may be performed by using state of the art applications such as Elastix [82]. Standard MRI imaging data cannot be converted to physical units and therefore is highly dependent on the scanner and sequences employed. In order to allow for comparison normalization is required. In our application, we perform two types of normalization. When we use a slice view, we normalize the data of the slice using a min-max normalization of the selected slice. In the 3D volume visualization, we normalize the whole volume by using the min-max value of the volume. This results in the most appropriate normalization based on the tasks the visualizations support.

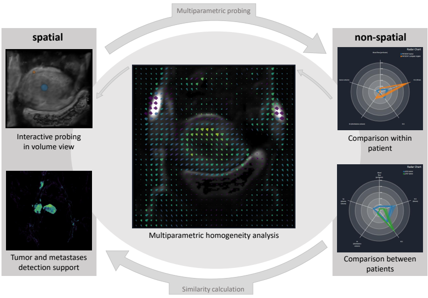

Figure A.2: The ParaGlyder application combines spatial (volumetric view) and non-spatial (radial boxplot) visualization to enable multiparametric analysis and exploration. In between, the Stixels view depicts a combination of both.

A.4.2 The Stixels View

Based on requirement R1, the goal is to raise the level of detail in the value domain but still keep the details in the spatial resolution. To facilitate this, we employ a glyph map approach, presented in the middle of Figure A.2, which is called the Stixels (**St**ar glyph p**ixels**) view. The glyph map is based on a regular grid which is overlaid on a 2D view of a slice. For every grid cell, we calculate statistics of the multiparametric medical imaging data. The star glyphs are then created by summarizing the statistics within each of the cells. The grid size and the star glyph size can be adapted, depending on the granularity of the structure of interest. By cropping the slice view to a region of interest, the glyph maps also adapt to the selection and allows for an even more detailed view of the selected structures. We use star glyphs instead of polyline-based glyphs since according to Opach et al. [115] star glyphs are a better choice for finding differences. For the star glyph design, we display the average value of each parameter within the grid cell on the axes. The area described by connecting these points forms a glyph which describes the relation of average parameters within the cell. When designing a star glyph, a homogeneous shape is favorable [85, 118]. Therefore, the order in which the parameters are presented is adjustable. While even more information could be encoded on the axes of the star glyph, we opted for a design that is easier to interpret and presents all necessary information at a glance to prevent a steep learning curve. The star glyph map provides an overview which allows the user to identify the tumor since the tissue differs from healthy tissue in the multiparametric dimensions. In addition, the inhomogeneity of the tumor can be analyzed. When spotting interesting parts of the tumor, a closer investigation of the area using the interactive probing can be employed.

A.4.3 3D Probing Visualization

Requirements R3 and R4 support analyzing different parts of the tumor independently, enabling identification of tumor patterns. Probing spheres deliver detailed information from data within selected regions. Regions of interest can be specified by using multiple probing spheres. This enables a comparison of different regions within the imaging data for a single patient, e.g., healthy tissue and cancerous tissue. All voxels from all parameters within the spheres can be used in the statistical analysis, like the approach used for the star glyph map. For the visual encoding of the probed regions a radial boxplot is used. It shows the user the summary statistics for selected regions of interest at a glance. Comparison is enabled by the superposition of multiple radial boxplots. Radial boxplots are favorable because they align with the use of star glyphs in the Stixels view and represent a more detailed view of selected areas. Differences and similarities over all modalities can be analyzed by placing multiple spheres either within the data of a single patient or multiple patients. To establish visual correspondence between the probing spheres and radial boxplots, both the spheres and boxplot share the same color hue. Interactive probing can be used to define a multiparametric pattern which describes different tumor characteristics based only on imaging data and may also be found in other patients suffering from a similar tumor type.

A.4.4 Interaction

To support requirements R3 and R4 various interaction methods are provided. The placement of the probing sphere can be performed either in 2D images or in 3D volumes. The size of the sphere can be adapted to fit the scale of the region of interest

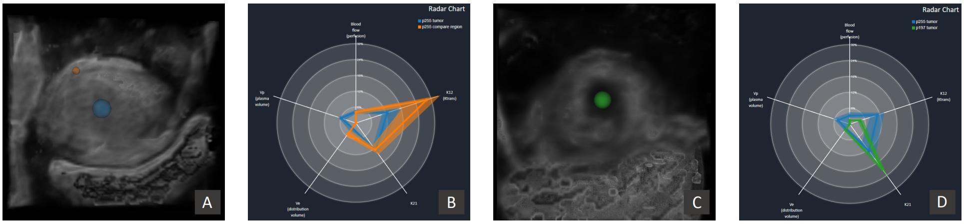

Figure A.4: Volume probing using two different probing spheres (a) results in live updates to the radial boxplot view (b). Probing interaction within another subject (c) results in a radial boxplot comparing data across subjects (d).

and the sphere can be placed freely. The quickest option is the free placement where the sphere is placed according to the intersection of a ray going from the screen position, where the mouse is located, towards the volume based on the closest visible point in the volume. In addition to this quick initial placement of the sphere, we introduce a mode where the sphere can only be translated within the current X-Y plane the sphere is located at. Another option only adapts the depth of the sphere along the Z-axis. When using a 2D view, it may occur that the probing sphere is behind the current slice and thus occluded. To remedy this, we provide an option to snap the sphere back to the slice. To support working with brain data, placing a sphere that is automatically mirrored to the other hemisphere is also possible.

A.4.5 Similarity Visualization

Requirements R5 and R6 state that a similarity analysis and an interactive parameter selection is beneficial in tumor analysis. Analyzing the tumor extent and possible metastases in surrounding tissue is a typical task for radiologists. In addition, segmentation of tumors is an active field of machine learning research, where some algorithms require feature selection. To support these tasks, we employ the multiparametric contents of a probed area in a similarity function. We decided to use the Euclidean distance over all dimensions because they are equally important. When applying this function to each multi-parametric voxel in the volume, we derive a new volume consisting of similarity values between 0 and 1 which can be displayed with an appropriate transfer function. A transfer function that highlights regions of high similarity through color and opacity enables users to identify structures such as tumors and possible metastases and enables a visual clustering with soft boundaries. Metastases which share the same imaging properties as the primary tumor are highlighted using direct volume rendering. Editing the transfer function enables the user to explore the inhomogeneity (R2) and the extent of different parts of the tumor. In addition, this similarity function-based visual encoding is also applied to the star glyph map. The fact that the similarity is based on the user-selected parameters enables the user to perform interactive feature selection (R6).

A.5 Results

The ParaGlyder application is depicted in Figure A.1 and consists of a center view (Figure A.1B), which provides common spatial visualization features, such as a 3D view, 2D slice-based views, cropping, and transfer function editing, and a probing functionality. Next to the main view, the Stixels view is located (Figure A.1C), which consists of a 2D slice view and an overlaid glyph map consisting of star glyphs. The last view is the probing view, component D in Figure A.1. It consists of a radial boxplot based on probing sphere input. We analyzed different datasets of endometrial cancer patients, provided by one of our co-authors, as well as a brain tumor dataset publicly available and provided by Schmainda and Prah [136] via the Cancer Imaging Archive (TCIA) [26]. The endometrial cancer dataset comprises standard multiparametric MR sequences and derived parameter maps visualizing physical parameters, e.g., blood flow and plasma volume. The data is co-registered due to its origin. For the brain tumor and inflammation data, we have access to the standard parameters acquired in multiparametric MR, such as T1-, T2- and diffusion-weighted images.

A.5.1 Tumor Detection and Multiparametric Homogeneity Assessment

To detect tumors and assess their multiparametric homogeneity, the Stixels view is used. The user selects the slice and the parameter to show. A detailed view of individual Stixels is presented when the user hovers the mouse over the specific Stixel. A detailed tooltip is shown, visualized in Figure A.3c. In order to support region of interest (ROI) selection, we employ volumetric cropping to select an appropriate Stixel window. The grid of the Stixels adapts accordingly and then probes smaller regions determined by the ROI. When placing a probing sphere, the Stixels are colored by the multiparametric similarity, measured based on Euclidean distance, using the Viridis colormap. The similarity Stixels view, visible in Figures A.3a and A.3b, additionally reveals the inhomogeneity of the tumor. The red line marks the outline of the tumor and the color and shape variations of the star glyphs represents the inhomogeneity within the primary tumor. In Figure A.3a, a tumor with a high degree of inhomogeneity is visible, while Figure A.3b reveals a more homogeneous tumor. The inhomogeneity analysis enables the user to spot distinct parts within the tumor, e.g., a necrotic core and allows for further analysis of these specific parts in detail.

A.5.2 Region Comparison for Tumor Characteristic Assessment

Probing spheres are used to analyze different parts within one patient or across multiple patients. This probing interaction is conceptually similar to a digital biopsy. The result of the probing interaction is a radial boxplot, visible in component D in Figure A.1. Figure A.4a showcases placement of two spheres for a single subject, while Figure A.4b shows a sphere placed to compare regions across subjects. The radial boxplot is shown in Figure A.4b and A.4d. On each axis, the median value is presented as a dot, and these dots are connected by lines. In addition to the median value, the 25% and the 75% quantile ranges are visible as an overlaid band. This representation allows the user to see the inhomogeneity of the data within the sphere. The maximum values of the axes can be adapted to fit the selected data range. The spheres are used to characterize tumor

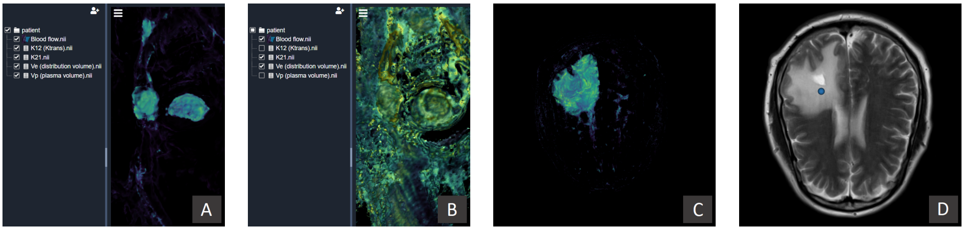

Figure A.5: The similarity view highlighting the uterine primary tumor in the center and two metastatic lymph nodes (a). When an insufficient number of dimensions is selected, the similarity view fails to capture the tumor and metastases (b). The similarity view captures brain inflammation (c), while simple thresholding on one modality would capture the skull as well.

tissue and to come up with specific signature shapes in the radial boxplot that can be used to classify the imaging data of new patients. The interaction responsiveness is ensured by providing a real-time update of the radar chart with the probed values of the volumetric multiparametric imaging data while the sphere is moved interactively through the volume.

A.5.3 Similarity Visualization for Metastases Detection and Feature Selection

The similarity view, visible in Figure A.1B and Figure A.5, visualizes the extent of a tumor and potential nearby metastases. FigureA.5(a) shows the similarity volume when using all multiparametric images and FigureA.5(b) shows the similarity volume with only three out of five of the multiparametric images. The Figure shows that the three selected images do not contain enough information to segment the tumor and the metastases. The colored Stixels are presented in Figure A.3c. For both approaches the Viridis colormap is chosen as a transfer function, where opacity is mapped to similarity, i.e., the visibility of regions that differ from the current selection is reduced. In component B of Figure A.1, the similarity view of parameter maps of a patient with endometrial cancer is visible. This similarity analysis enables a clear and distinct visualization of the tumor (the lower right structure in the inset), by placing a probing sphere inside the tumor tissue. Due to their multiparametric similarity, metastases in the lymphatic system (structures to the left and above the primary tumor) are also highlighted. When analyzing only one of the multiparametric images at a time the detection of metastases is much more difficult because they are not clearly visible. When probing inflammatory data within the brain, the similarity view provides a quick segmentation of inflamed tissue. The segmentation does not include the bone as a standard thresholding operation based on T2 Flair data only would, visualized in Figure A.5(c) and Figure A.5(d). This demonstrates that the multiparametric similarity function facilitates a rapid multiparametric segmentation, which could be used in diagnosis or treatment planning, as well as feature analysis as input to automatic segmentation methods in a machine learning context.

Table A.1: The response of the experts on a 5-point Liker scale. The values range from 1: Strongly disagree to 5: Strongly agree. Statements marked with a star were rephrased to present the positive form in this table, also the scores have been inverted. On the right end of the table the average value over all experts is presented and in the last row the result of the SUS questionnaire is presented.

	Statement	N1	N2	N3	Gy1	Gy2	Gy3	Gy4	Gy5	M1	M2	M3	Avg.
G1	The linked interactions between the center view and the radar chart are well established and intuitive	3	5	4	5	5	5	5	3	5	5	4	4,45
G2	The linked interactions between the analyze view and the Stixel view are well established and intuitive	5	4	4	5	5	5	5	4	5	5	5	4,73
G3	I see myself using the MRI Explorer in the future*	3	5	3	5	5	3	5	5	5	4	4	4,27
G4	I would like to contribute in the future development of the application	5	5	5	5	5	4	5	3	3	5	4	4,45
G5	I can see the application as a part of my daily work routine*	3	5	2	5	4	1	3	5	4	1	1	3,09
G6	The application is more applicable for research than for daily clinical practice	3	4	5	5	4	5	4	3	5	4	2	4,00
G7	The application should be part of the software used in clinical practice*	1	5	1	3	4	3	5	5	4	4	5	3,64
P1	The navigation in 3D is easy to understand and I can place the sphere where I want*	5	4	2	2	5	5	5	4	5	5	5	4,27
P2	The resizing operation of the sphere is easy to understand and to carry out*	4	4	2	5	4	5	5	5	5	4	5	4,36
P3	I can place the sphere anywhere on the plane using the provided keyboard interactions	5	5	3	4	5	4	5	5	5	5	5	4,64
P4	Setting the probing sphere to a specific depth in the volume is intuitive	2	5	2	3	5	4	4	4	5	5	4	3,91
P5	Snapping the probing sphere to the current slice selection is useful	5	4	3	5	5	5	5	4	5	5	3	4,45
P6	The probing interaction is responsive*	3	5	4	5	5	5	5	5	5	5	5	4,73
P7	The automatic update of the Radar chart is beneficial*	5	5	4	5	5	5	4	5	5	5	5	4,82
P8	The radar chart helps me to interpret multimodal data	5	5	5	5	5	5	5	5	5	4	5	4,91
P9	I am confident in interpreting the values that the radar chart presents	3	5	3	4	5	3	4	3	5	5	3	3,91
P10	With the probing functionality, I am able to compare different regions within one subject*	5	5	3	4	5	5	5	5	4	5	4	4,55
P11	The probing function enables me to compare regions between different subjects	3	5	4	5	5	5	4	5	5	3	4	4,36
St1	I understand what the Stixels view shows me and can interpret the star glyphs used.	3	5	4	4	5	3	4	3	5	5	3	4,00
St2	The Stixels view helps me to gather insight of the inhomogeneity of the data*	5	5	3	5	5	5	5	5	5	3	4	4,55
St3	The cropping functionality helps me to focus the Stixel view on the most important region of the subjects data*	1	5	2	5	5	5	5	5	5	5	4	4,27
St4	The different grid sizes help me to first get an overview and add details on demand	5	5	4	5	5	4	5	5	4	5	4	4,64
St5	The tooltip helps me to see more details in the Stixels view when I need them	5	5	4	5	5	5	4	5	5	4	4	4,64
S1	I understand the color coding of the Stixels in terms of similarity*	3	5	4	5	5	5	5	5	5	5	5	4,73
S2	The similarity coloring of the Stixels helps me to adapt my probing selection	3	5	3	5	5	4	4	5	4	5	4	4,27
S3	The similarity volume visualization shows me interesting parts of the volumetric data	5	4	5	5	5	5	5	5	5	5	4	4,82
S4	The similarity view is useful to me and I would like to use it in my work routine/research*	5	5	5	5	5	5	5	5	5	1	4	4,55
Gys1	The application can improve the analysis of the inhomogeneity of gynecological cancer				5	5	5	5	5				5,00
Gys2	The application can support hypothesis generation for linking parameters with aggressiveness of gynecological cancer				5	5	5	5	5				5,00
Gys3	I would find this application useful when analyzing patients gynecological cancer MR data*				5	5	5	5	5				5,00
Gys4	I would like to use this application to explain pathology and treatment to patients				4	5	1	3	3				3,20
Gys5	I would like to use this application to plan a biopsy for analyzing biomarkers of the tumor				X	5	1	3	5				3,50
Gys6	The application is useful for finding metastases*				5	4	2	5	5				4,20
Gys7	The similarity view shows me the structure of the tumor				5	4	5	5	4				4,60
Gys8	The similarity view shows me the size and structure of possible metastases				5	5	3	5	5				4,60
Ns1	The application helps me to visualize lesions in the brain	5	4	1									3,33
Ns2	Comparing different regions within the brain using the comparison picker is particularly useful for me*	5	5	5									5,00
Ns3	The similarity view helps me to get a better volumetric view of the lesion*	5	5	4									4,67
Ns4	I would like to use this tool to further analyse multiparametric brain imaging data	3	5	4									4,00
Ns5	The interaction with the comparison tool is suitable for brain images	5	5	3									4,33
Ns6	The application helps me see the intensity relations of the different tissue types between modalities*	5	5	5									5,00
Ms1	The application helps me to carry out feature selection prior to applying my machine learning algorithms									4	4	3	3,67
Ms2	I find the similarity view useful to identify which modalities are important for me*									4	4	5	4,33
Ms3	I can imagine using this tool before applying machine learning algorithms*									5	4	4	4,33
Ms4	This application is particularly useful for segmentation based on machine learning									5	4	5	4,67
SUS	System usability scale result	97,5	85	40	75	85	80	87,5	80	95	92,5	85	81,75

A.6 Evaluation

We conducted a qualitative evaluation with eleven experts (6 male, 5 female) from the scientific fields of neurological imaging (N1-3), gynecological cancer imaging (Gy1-5) and machine learning research (M1-3). One expert is co-author and provided us with clinical data of gynecological cancer patients and one expert of each domain (N1, Gy2, M2) was included in the interviews during the development of our application. We were especially interested in validating the effectiveness of the various visualization components and identifying opportunities to make our application more suitable for daily research or even clinical practice. The individual evaluation started with a short demonstration of the tool, afterwards experts were encouraged to explore and analyze the multiparametric data themselves. They were invited to comment using a think-aloud protocol. The gynecological cancer and machine learning experts worked with endometrial cancer data and the neurological imaging experts with data provided by Schmainda and Prah [136] via the Cancer Imaging Archive (TCIA) [26]. After this phase, which lasted around 30 minutes, we conducted a semi-structured interview with the experts. Finally, a questionnaire consisting of 27 generally applicable statements and 4-8 targeted statements for the different expert groups was conducted. The experts were asked to indicate their level of agreement using a five-point Likert scale. In addition to our targeted evaluation form, we asked the experts to fill out the system usability scale (SUS) provided by Brook et al. [17]. The evaluation results of the eleven participants are shown in Table A.1.

We conclude from the results presented in Table A.1 that the application is overall valuable for the experts. The probing interaction was rated favorably, two participants would appreciate a guided 3D placement of the probe. All study participants think that the Stixels view helps them to see inhomogeneous regions within the Slice view. The similarity view received the most positive feedback and is potentially useful for all involved experts. The targeted statements demonstrate that the application is applicable different scenarios, albeit for different reasons. The gynecological experts envision that the application could improve the assessment of tumor heterogeneity both in primary tumors and metastases. The SUS scores range from 40 to 97,5, where the second lowest score is 75. On average, the SUS score is 81,75. According to Bangor et al. [5], the score can be interpreted to be between good and excellent.

A.7 Conclusion and Future Work

We present ParaGlyder, a multiparametric image visualization tool. The tool provides different views for tumor detection, inhomogeneity analysis, feature selection, and diagnosis in multiparametric medical images, by a tight coupling of spatial and non-spatial data visualization techniques. Our tool is based on a combination of star glyph maps and radar charts. A built-in similarity visualization of the volumetric data enables the visualization of, e.g., primary tumor and the corresponding metastases. The qualitative evaluation confirmed the utility of our application for diverse application areas. In the future, we plan to extend our approach to analysis of larger patient cohorts in order to assess whether this visualization tool could aid in the detection of metastases. Furthermore, the application has the potential to unravel patient-specific imaging features that may be linked to specific clinical phenotypes and outcomes, thus representing a promising tool to facilitate more personalized treatment strategies.

A

Paper B

RadEx: Integrated Visual Exploration of Multiparametric Studies for Radiomic Tumor Profiling

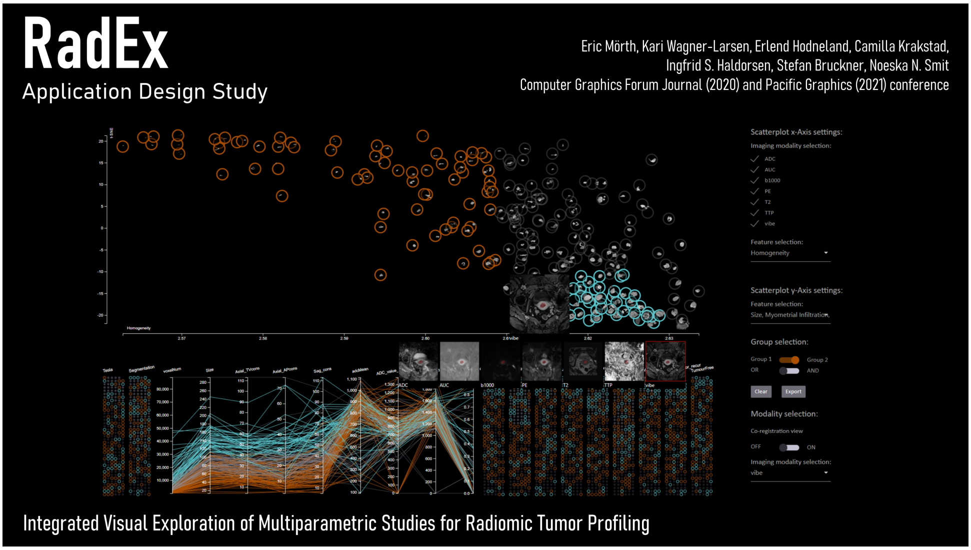

E. Mörth[1,2], K. Wagner-Larsen[2,3], E. Hodneland[2,4], C. Krakstad[2,5], I. S. Haldorsen[2,3], S. Bruckner[1,2] and N. N. Smit[1,2]

[1]Department of Informatics, University of Bergen, Norway
[2]Mohn Medical Imaging and Visualization Centre, Department of Radiology, Haukeland University Hospital, Norway
[3]Department of Clinical Medicine, University of Bergen, Norway
[4]Norwegian Research Centre (NORCE), Bergen, Norway
[5]Centre for Cancer Biomarkers, Department of Clinical Science, University of Bergen

---

This article was published in *Computer Graphics Forum 39 (Pacific Graphics 2020*, pages 611-622, 2020

Abstract

Better understanding of the complex processes driving tumor growth and metastases is critical for developing targeted treatment strategies in cancer. Radiomics extracts large amounts of features from medical images which enables radiomic tumor profiling in combination with clinical markers. However, analyzing complex imaging data in combination with clinical data is not trivial and supporting tools aiding in these exploratory analyses are presently missing. In this paper, we present an approach that aims to enable the analysis of multiparametric medical imaging data in combination with numerical, ordinal, and categorical clinical parameters to validate established and unravel novel biomarkers. We propose a hybrid approach where dimensionality reduction to a single axis is combined with multiple linked views allowing clinical experts to formulate hypotheses based on all available imaging data and clinical parameters. This may help to reveal novel tumor characteristics in relation to molecular targets for treatment, thus providing better tools for enabling more personalized targeted treatment strategies. To confirm the utility of our approach, we closely collaborate with experts from the field of gynecological cancer imaging and conducted an evaluation with six experts in this field.

B.1 Introduction

The World Health Organization announced in 2018 that cancer is globally the second leading cause of death after cardiovascular disease [15]. There are numerous forms of cancer and they arise in all kinds of cells in the human body. When exploring imaging features, radiomic tumor profiling may be performed [4]. Gillies et al. state that the goal of radiomics is to harvest high dimensional data from clinical images which serve as a basis for further analysis, e.g., in terms of predictive value for predicting clinical outcome and response to targeted therapy [45]. This analysis is by no means an easy task and often involves working with multiple modalities or multiple parametric images. For endometrial cancer, which is the most common gynecological tumor in industrialized countries [2], preoperative staging by multiparametric magnetic resonance images (MRI) and results from endometrial biopsy routinely guide choice of surgical procedure and adjuvant therapy.

Tumor regions of interest (ROIs) can be manually placed on the different MRI sequences to quantify tissue features (e.g., diffusion properties on apparent diffusion coefficient (ADC) maps, perfusion markers on dynamic contrast enhanced (DCE)-MRI/parametric maps), and these tumor characteristics have been shown to aid in predicting metastases or tumor progression [7, 40, 56, 57]. However, manual tumor segmentations and analyses of ROI data are time consuming, and thus presently unfeasible in daily routine. A main challenge when analyzing cohorts of patients is to unravel the most relevant imaging patterns that are linked to clinical parameters and patient outcome. Furthermore, analysis of the high dimensional value domain resulting from the combination of the clinical parameters consisting of histological markers, radiological findings, and outcome related parameters and multiparametric imaging data is challenging.

In current clinical research, imaging modalities are typically analyzed individually, and the tumors are analyzed based on manually placed ROIs. As a part of ongoing research of our collaborators, convolutional neural networks are used to perform a machine learning based segmentation of tumors. In addition, an automatic co-registration of all available modalities is performed. This allows integrative analysis of all voxels of the tumor, with the potential to empower clinical researchers with more targeted and capable analysis platforms. Tumor textural features based on manually placed ROIs and only for single modalities have already been proven to be related to high-risk histological subtypes in endometrial cancer [40, 57, 171]. Analysis tools for different ROI measurements taking multiple modalities and all tumor voxels into account are not available. Nonetheless such an analysis would likely improve comparability between different patients and across different hospitals.

To further support cohort analysis in radiomic tumor profiling, we provide a tool which supports both data preprocessing steps for cohort analysis as well as an integrated dashboard for cohort analysis and hypothesis generation. The application targets cancer imaging research where radiomic tumor profiling is carried out and where multiparametric imaging data is acquired. We provide an interactive analysis platform that enables gynecological cancer researchers to analyze cohort data with the goal of hypothesis formation. The overall aim of the application is to provide a tool for visualization, exploration, and identification of complex relations between radiomic tumor profiles and clinical and histological markers.

Our main contributions are the following: (1) We present an interactive cohort analysis application that targets hypothesis formation in radiomic tumor profiling workflows that includes imaging data and clinical parameters. (2) We provide a workflow that enables validation of manual or automatic machine learning-based tumor segmentations and validation of automatic co-registration of multi-parametric imaging data. (3) To show the utility of our approach, we evaluated our system with six experts in gynecological cancer imaging research, using the System Usability Scale (SUS) [17] and a qualitative evaluation.

B.2 Medical Background

Endometrial cancer is the most common pelvic gynecological malignancy in high-income countries. The endometrium comprises the innermost lining of the uterine cavity, and patients typically experience abnormal vaginal bleeding. The diagnosis is confirmed by an endometrial biopsy establishing the histological-/molecular tumor subtype/grade, and subsequently a preoperative pelvic MRI is routinely performed for local staging. However, analyzing complex imaging data in combination with clinical-/tissue data is not trivial. Imaging data is often multiparametric, allowing the visualization of different aspects of tumor physiology related to, e.g., tumor microcirculation and microstructure, which reportedly are closely linked to the observed clinical phenotype in cancer [40, 57]. Extracting whole volume multiparametric imaging data simultaneously may also be done through radiomic tumor profiling.

Radiomic tumor profiling plays an emerging important role in the new era of precision medicine [45]. This field of research aims to produce high-dimensional feature vectors from clinical images [131] to find tumor markers with a higher predictive value. The typical workflow consists of image acquisition, image reconstruction, tumor segmentation, feature extraction and qualification, and finally analysis and model building. One challenge in this regard is feature selection which may be performed a priori or from dimensionality reduction. Radiomics features are extracted from a region of interest representing the tumor. Hence, a high quality tumor segmentation is a crucial step [131]. Manual segmentation is often used as the ground truth for tumor segmentation although some inter-reader variability is inevitably present [131].

Tumor texture parameters derived from MRI scans have been shown to be associated with high-risk disease and reduced survival in endometrial cancer [171]. Parameters like kurtosis, entropy and mean of positive pixels (MPP) from the ADC map and the T1 contrast-enhanced images have been shown to predict high-risk histological subtypes and advanced stage, e.g., deep myometrial invasion in endometrial cancer [40, 171]. Furthermore, tumor tissue markers have been compared with tumor markers from parametric maps (based on dynamic contrast-enhanced MRI) finding that reduced tumor blood flow in MRI reflects increased microvascular proliferation in the tumor samples and predicts poor survival in endometrial cancer [40, 57].

B.3 Related Work

Traditionally, gynecological cancer biomarker research employs statistical analysis tools like SPSS [52] or RStudio [134]. Visualization approaches for targeted cancer diagnosis often focuses on techniques for representing biomarkers or imaging data but there is only a small body of work combining clinical cohort data with multiple imaging sequences per patient. Raidou et al. [128] introduced a visual analytics approach for the exploration of tumor tissue characterization featuring a 2D t-Distributed Stochastic Neighbor Embedding (t-SNE) [163] dimensionality reduction also taking tumor characteristics into account. Although their approach is similar, we differ in multiple aspects in that we do not focus on pharmacokinetic parameter maps but rather visualize high-dimensional radiomic tumor texture features. In contrast to their approach, we use a 1D t-SNE dimensionality reduction and our visual exploration approach aims for the identification of complex relations between radiomic tumor profile and clinical and histological markers. To the best of our knowledge, there is no related literature where a 1D t-SNE dimensionality reduction has been used in combination with a clinically meaningful parameter on the other axis. Compared to the other approaches presented here, we preserve one axis of the 2D visualization to present a meaningful value selection, which is not the case when applying a 2D t-SNE dimensionality reduction.

Image-Centric Cohort Visualization Closely related to our approach is the work by Steenwijk et al. [149], where the authors used multiple linked views including scatterplots and parallel coordinate plots combined with imaging data for each patient. One major difference is that they are limited to displaying two images per patient without a link to the imaging data in their visualization application. Klemm et al. [83] introduced an epidemiological approach which is also image-centric and involves segmentations and hypothesis formulation. However, they do not visualize the original imaging data in their application and mainly focus on model-based visualization in their application.

Jönsson et al. [74] introduced a cohort analysis platform which also allows for group comparison and incorporates imaging data for each patient as well as clinical parameters. In contrast to our approach, they do not use multiparametric imaging data, and they do not work with radiomic tumor features.

Visual Cohort Analysis Preim et al. [124] provided an extensive overview of visual analytics approaches for public health data in their survey. Further related work in the field of visual analytics of patient cohorts includes the work of Angelelli et al. [3]. The authors presented a prototype aiming for cohort-based hypothesis formulation for heterogenous data. Eckelt et al. [36] presented a visual analysis tool enabling statistical analysis of tabular data, cancer drug target discovery, and closing the gap between visualization and statistical analysis. Raidou et al. [130] introduced a visual analytics application which allows for analysis on both cohort level and patient level for radiotherapy induced bladder toxicity. Bernard et al. [8] presented a data centered approach for analyzing large amounts of patients using multiple linked views and selective analysis. When dealing with large amounts of patient data, user guidance could be implemented, as discussed by Ceneda et al. [23, 24]. Further related work in the field of cohort construction includes the work of Krause et al. [91]. In the field of cancer characterization, the work of Turkay et al. [160] is relevant to our approach. In contrast to these approaches, we combine multiparametric imaging data, radiomic tumor profiling data and clinical parameters in one application.

Co-Registration and Segmentation Validation As we work with co-registered data which partially features machine learning-based segmentation masks, the data needs to be validated by experts before further analysis. Hastreiter et al. [59] proposed to use fused visualization methods. Jenkins et al. [67] suggested that overlaying of prominent edges provides a usable co-registration check. More complex and automatic co-registration validation methods include the work of Schnabel et al. [137], which validates nonrigid image registration using finite-element methods. Our approach adds a parameter-based pre-selection of cases where the co-registration outcome is suspicious to first present these cases to medical researchers for further analysis.

Automatic segmentation methods are available, but often they do not meet the acceptance criteria needed for usage in clinical cohort studies. Therefore, a validation step has to be employed before using them as a source of analysis [165]. Von Landsberger et al. [165] introduced a user guided automatic segmentation method where algorithm parameters are set intuitively by using visual analytics tools. Karimov et al. [75] introduced an approach for interactive segmentation correction based on histogram dissimilarities. Haehn et al. [54] proposed a segmentation proofreading technique and demonstrated that expert proofreading has increased performance and speed over manual expert segmentation. In contrast to these approaches, we benefit from various features measured by radiologists which allow us to spot potentially incorrect automatic segmentation. Adding these cases to the training set of automatic algorithms allows for an incremental improvement of the segmentation mask result.

B.4 Data and Tasks

Endometrial cancer classification is highly dependent on the data used. For our clinical collaborators, multiparametric imaging data and histological data as well as further clinical parameters of the patients are available. Due to our close collaboration with both radiologists and experts in the field of molecular biomarkers we were able to gather insight in this highly specialized field of research. In a clinical setting, potential patients who typically face symptoms such as vaginal bleeding, have an endometrium biopsy which serves as the basis for a histological investigation. If the biopsy confirms an endometrial cancer diagnosis, preoperative pelvic multiparametric contrast-enhanced MRI is routinely performed. Imaging findings guide the choice of treatment, normally consisting of surgery in all cases, followed by adjuvant chemo- and/or radiation therapy in high-risk patients. After treatment, the patients have regular clinical follow-ups to detect recurrent disease/tumor progression, and progression-free survival is recorded.

B.4.1 Clinical Parameters

The following clinical parameters are available for analysis:

- Tesla: Field strength value of the MRI scanner used for the screening, either 1.5 or 3 Tesla.
- Segmentation: Indicator of manual or machine learning segmentation of the tumor.
- FIGO2G: International Federation of Gynecology and Obstetrics (FIGO) [9] classification of the tumor. FIGO I and II are one group and FIGO III and IV represent the other group. The grouping is performed based on tumor aggressiveness.
- MyomInf2G: Myometrial infiltration of the tumor with infiltration of <=50% and >50%.
- CervixAffHyst: Tumor extending to the uterine cervix.
- HistType: Endometrial or non-endometrial subtype.
- HistGrade2G: High-grade and low-grade tumors.
- MetNodes: Histologically confirmed lymph node metastasis, no metastases or not investigated.
- Status: Last known status of the patient, category one combines the following possible states of the patient: alive and well, dead from other causes or dead with but not due to active disease, second category is alive with active disease, and third category indicates dead from disease.
- Prog_and_or_recur: Progression or recurrence of disease after surgery.
- TumorFree: Tumor free at the most recent follow-up meeting.

B.4.2 MRI Specifics

MRI imaging includes different sequences depicting tumor extent (using T1- and T2-weighted MRI) and microstructural tumor characteristics (e.g., in diffusion weighted imaging (DWI)). First introduced by Rofsky et al. [132], the Volumetric Interpolated Breath-hold Examination (VIBE) facilitates a 3D gradient-echo sequences that produces T1 weighted images. The advantage of this approach is improved resolution in the Z-axis which enables high-quality multiplanar reconstruction. In DWI, highly cellular tissue features a lower diffusion coefficient [81]. A quantitative assessment of the diffusion may be performed with the generation of apparent diffusion coefficient (ADC) values obtained at different b-values [87]. When using an intravenous contrast agent, the dynamic contrast-enhanced (DCE) MR perfusion is recorded. Typical measurements during this examination are the peak enhancement (PE) measuring the relative enhancement in contrast after the update of the contrast agent and the time to peak (TTP) [51]. The area under the peak enhancement curve (AUC) is also a typical measurement [51]. In total, we have seven MR imaging sequences available.

B.4.3 Specifics of the Application Domain

Gynecological cancer imaging research consists of multiple steps. Data of several MRI sequences must be analyzed, currently done separately and partly only for specific regions within the tumor. The overall goal of our medical collaborators is to examine and explore tumor biomarkers which potentially have a larger predictive value for clinical outcome than well-established ones. These biomarkers may help to further improve treatment of patients and increase personalization. Recent research of our collaborators already includes the analysis of tumor texture, which relies on prior segmentation of the tumor [40, 57, 171]. The gold standard in this segmentation approach is a manual segmentation performed by a trained radiologist. One aspect to consider in this regard is that this step comprises intra-operator variance. Due to the very time-consuming process of segmenting each volumetric slice this process is only performed using one of the seven MRI sequences. Therefore, our collaborators aim for a complete automation using machine learning algorithms. Another step already performed by our collaborators is the co-registration of all seven sequences present. For both steps our collaborators are looking for a validation possibility to assess data quality before analysis. The analysis of the imaging data and clinical parameters involves several tasks: group selection, tumor texture feature analysis, data quality validation, and hypothesis formation. To the best of our knowledge there is no application available which combines these tasks in an easy and intuitive way without having to export and import data multiple times.

Machine Learning Segmentation

One of our co-authors applied a 3D convolutional neural network (UNet3D [25]), using Keras [53] and Tensorflow [1] as backend engine, to facilitate automatic segmentation of the tumor data in endometrial cancer patients. The network was trained using 139 expert segmentations based on preoperative pelvic imaging. The network can retrieve tumor volumes which are comparable to human expert level and a set of segmentation masks with human agreement not differing from inter-rater agreement. Although this algorithm is very promising for further analysis of the segmented tumor volumes and masks, proofreading by a radiologist is still necessary. Common tools used in clinical practice and research allow for such a validation but quickly finding cases where the segmentation might be wrong is desirable but not supported yet. After this processing step, the tumor segmentation is still only available on one sequence. To analyze the tumor on every given sequence, co-registration is necessary. This process was performed automatically but needs to be validated since a segmentation of non-tumor-regions would potentially introduce a critical error in data analysis. Examples of incorrect segmentations can be seen in the accompanying video and figures.

Automatic Co-registration

Our collaborators performed the co-registration automatically using FMRIB's Linear Image Registration Tool (FLIRT) [67, 68] without optimization and only performing geometric alignment in scanner coordinates. However, this automatic registration method may not always find a perfect transformation for each modality and therefore must be validated. The employed co-registration algorithm features a relatively low failure rate but for a meaningful analysis spotting cases where it might have failed is crucial. Our medical collaborators request for an intuitive way to find and validate these cases. Having all sequences co-registered and the segmentation prepared, radiomics feature extraction is the next step. Examples of incorrect co-registrations are presented in the additional materials.

Radiomics Feature Extraction

Radiomics feature extraction takes volumetric imaging data and the volumetric tumor mask as input and generates a high-dimensional feature vector describing the tumor in each parametric imaging sequence. We merge the generated data afterwards with the clinical parameters. In recent work by our collaborators [171] tumors textural features were analyzed using TexRad [97] software. The number of features in this approach is limited and the feature generation algorithms are not open source. Therefore, our medical collaborators expressed interest in a transparent and more controllable data handling method. Based on prior research, tumor texture features are interesting measurements, believed to be correlated with aggressiveness of tumors [40, 171]. To further support this hypothesis, we calculate potential features for homogeneity analysis, including normalized inverse difference moment, contrast, short run emphasis and long run emphasis. The information content contribution of each sequence is not known a priori and therefore an explorative analysis of single sequences and their combinations is of interest for our collaborators.

B.4.4 Task Abstraction

We performed a task abstraction using the task framework proposed by Brehmer and Munzner [16]. We assessed the current status of clinical research in gynecological cancer during multiple interviews with our collaborators. We also encouraged them to envision new workflows including results of convolutional neural network-performed segmentations and the possible parallel analysis of all sequences after co-registration.

During the interviews, we identified two phases. Phase one deals with ensuring data quality and phase two with cohort analysis. Two tasks (T1 and T2) handle the need for a segmentation and co-registration validation. Task T3 reflects a common practice in medical research, namely group selection. Finally, tasks T4 and T5 provide analysis functionality. During discussions, our collaborators mentioned that they commonly use R or SPSS for statistical analysis and that they would like to continue doing so. Therefore, we exclude statistical analysis capabilities from our application design.

T1–Discover Invalid Co-Registrations Our medical collaborators classify co-registrations as invalid based on the misalignment of the given volume. Bladder filling and other physiological processes in the body cause a shift of the organs and therefore automatic methods may not result in a completely perfect alignment of these images. However, finding cases where automatic co-registration fails to find a sufficient transformation is crucial to support productive and time-efficient analysis. The analysis platform should allow the user to *discover* misaligned volumes and to analyze data in detail to *identify* potential causes of the misalignment.

T2–Discover Incorrect Machine Learning Tumor Segmentations Manual volumetric segmentation of endometrial tumors is a tedious and time-consuming task. This task could be automated using machine learning techniques, for example based on convolutional neural networks. Although the algorithm employed by our collaborators features a low failure rate and a comparable precision as the medical experts, the results still must be validated before further analysis. Spotting cases where the segmentation is potentially wrong is challenging and browsing through all patients is not time efficient. The user wants to *discover* potential faulty cases and *identify* the cause of the incorrect segmentation mask.

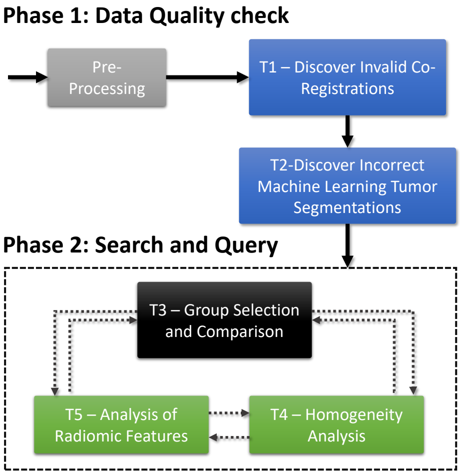

Figure B.1: The RadEx workflow based on our task abstraction. The flow is denoted by arrows consists of two phases. Phase One: **Discover** invalid co-registrations and segmentation masks. Phase Two (**Search and Query**): When the data quality is ensured ,users are able to search and query the whole cohort. The arrows depict that there is no pre-defined ordering of the tasks: each task can be executed in any order in the search and query section.

T3–Group Selection and Comparison Group identification and selection is a common and important task in clinical research. When analyzing a cohort of patients, it is of great interest to spot patients which share similar features, e.g., in imaging or histological

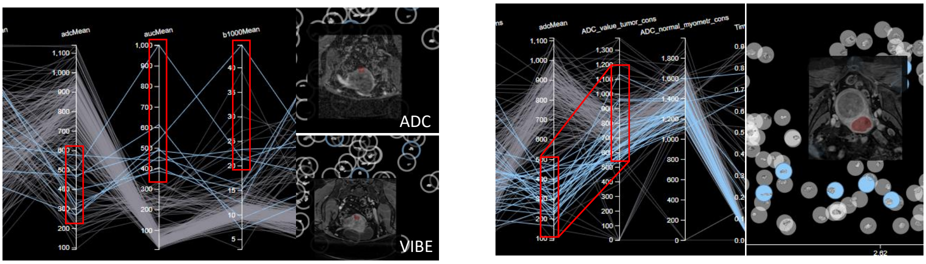

(a) Outliers in the co-registration view of the parallel coordinate plot indicate misaligned volumes after co-registrations. If acquisitions are not registered, segmentation output based on a single sequence will not be correct for other sequences. This leads to unexpected derived values.

(b) The ADC map value is comparable between patients and therefore routinely measured in gynecological cancer cases. A correlation analysis to the mean ADC value within the tumor segmentation could raise suspicion with regards to segmentation accuracy.

Figure B.2: Typical cases where co-registration (a) or tumor segmentation (b) failed. Finding these cases is not an easy task and browsing through all patients would be very time-consuming. Therefore, the parallel coordinate plot in the co-registration view enables a quick search for potential erroneous cases.

analysis results. During the analysis of such cohort data, different groups can be *selected* and the medical researchers would like to *compare* them with each other.

T4–Homogeneity Analysis Heterogeneity is putatively linked to aggressive cancer phenotype supported by previous studies linking specific textural features to high-risk histological subtypes in endometrial cancer [40, 171]. An exploratory platform allowing an assessment of textural features reflecting tumor heterogeneity derived from the different sequences/parametric maps interactively would potentially be clinically useful. *Browsing* through further homogeneity measurements over all available imaging modalities and *comparison* of predictive value is interesting for our collaborators.

T5–Analysis of Radiomics Features Radiomics feature generation is a promising way to generate characteristics which might facilitate a predictive value of specific clinical parameters. Experts would like to have a look at individual parameters, but also at the analysis of multiple combined parameters. They want to *locate* specific combinations or single parameters to *identify* a predictive value to previously defined groups, based on clinical parameters.

B.5 RadEx Workflow and Interface

Figure B.1 illustrates the workflow of RadEx when analyzing unprocessed cohort data. The first two interactions with the application ensure data quality for further analysis steps. Before analyzing the data within RadEx, a pre-processing step is required pre-calculating all slices per patient and tumor extent. The analysis workflow within the RadEx application starts with co-registration validation, followed by machine learning

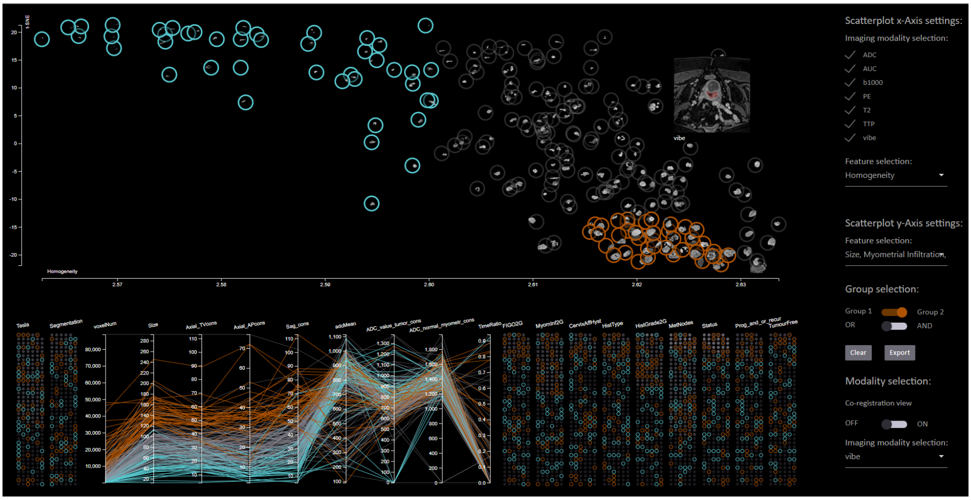

Scatterplot x-Axis settings:
Imaging modality selection:
✓ ADC
✓ AUC
✓ b1000
✓ PE
✓ T2
✓ TTP
✓ vibe
Feature selection: Homogeneity

Scatterplot y-Axis settings:
Feature selection: Size, Myometrial Infiltration,

Group selection:
Group 1 / Group 2
OR / AND

Modality selection:
Co-registration view
OFF / ON
Imaging modality selection: vibe

Figure B.3: Overview of the RadEx interface, an integrated analysis and exploration platform for gynecological cancer data.

segmentation validation. These two steps ensure the data quality and deliver valuable feedback to our collaborators to further improve the segmentation and registration output. After this step, clinical researchers can analyze the cohort data. They can focus on different aspects, for example, browsing, exploring, or locating specific feature. Group selection is performed as a first step, users can explore different homogeneity measurements or radiomics feature combinations. If the users are interested in specific characteristics, they can browse for possible groupings. The number of patients currently included in our application is already a high number for this type of studies, therefore we chose our visualization techniques to cope with the given number of patients.

B.5.1 Central Scatterplot View

The central scatterplot view presented in Figure B.3 reveals the whole cohort at a glance. This view plots a homogeneity measurement against a one-dimensional t-SNE dimensionality reduction of selected radiomics features allowing for an overview of the data. Each patient is marked with a gray circle. According to Cleveland et al. [27], three factors determine effective scatterplot design: (1) the marks are designed with preattentive features in mind, (2) the detection of individual objects is in focus, and (3) the distance between the objects presents a notion of similarity. We use these features to guide our scatterplot design. As shown in Figure B.3, each mark representing a single patient contains a small glyph representing the shape and the size of the tumor. We generate this small image by finding the slice with the largest amount of tumor voxels and extract VIBE pixels within the segmentation mask. If the tumor consists of multiple parts, these are still visible within the circle. An example of these glyphs is shown in Figure B.7. Hovering over a glyph reveals a tooltip showing an image slice with the tumor segmentation as a color overlay. The imaging modality as well as the slice can be selected by the user. This gives the clinical experts a direct relation to underlying imaging data and allows for a detailed co-registration and segmentation validation. The tooltip view presents details on demand by holding the Shift key while hovering over a mark. In this detailed version of the tooltip, an overview over all modalities is presented allowing the user to compare all seven modalities (T1). Both tooltip versions are shown in Figure B.4. The tooltip is only shown on mouse hover over a glyph of interest representing a single patient. At this point the user is interested in exploring data of this single patient, occlusion of other patient glyphs is therefore less problematic. This method for tooltip display does not introduce a visual focus change for the user and is therefore the appropriate placing for it. By scrolling through the modalities or by using the detailed tooltip version, a co-registration check and a segmentation validation can be performed (T1-2).

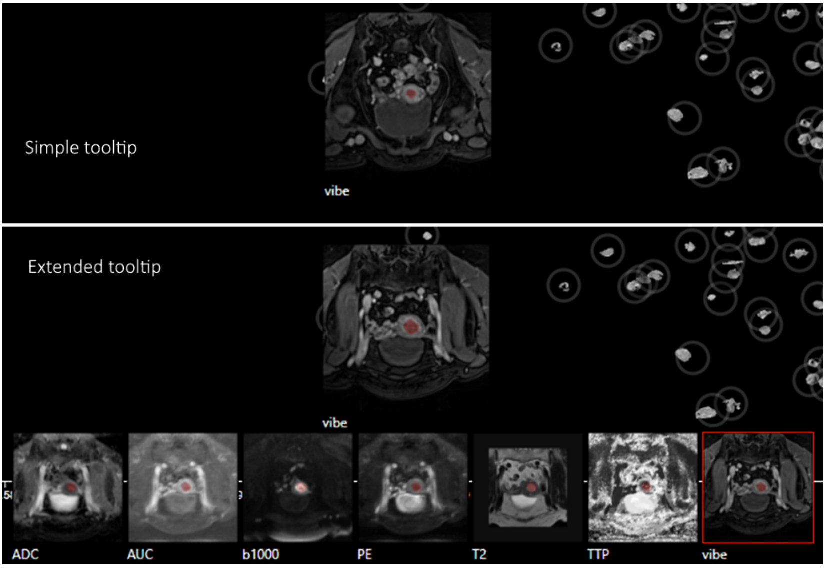

Figure B.4: Top: The simple tooltip reveals an image slice with segmentation information as a color overlay. Bottom: The extended tooltip visualizing all available modalities with tumor segmentation overlay and a red border around the selected modality.

A large amount of data and large marks in the scatterplot lead to overplotting, which can be avoided by various methods. Marks or position could be changed locally to reduce overplotting on demand. We use a simple zoom and pan interaction because it is efficient, and our collaborators are already familiar with these interactions. The methods are also easy to understand and execute while keeping the position of the dots in the scatterplot space static. We also added an option to reset the zoom and pan to its original state on demand.

Selection of the axes is crucial in scatterplots. As inhomogeneity is an essential but rather new measurement that is believed to correlate with tumor aggressiveness, we display this on the x-axis (T4). Following the description of Cleveland et al. [27] we use the x-axis to bring patients with a similar homogeneity closer to each other to enable clustering. For the y-axis we aim to allow for clustering patients according to similarity in the higher dimensional feature space consisting of clinical parameters and radiomic

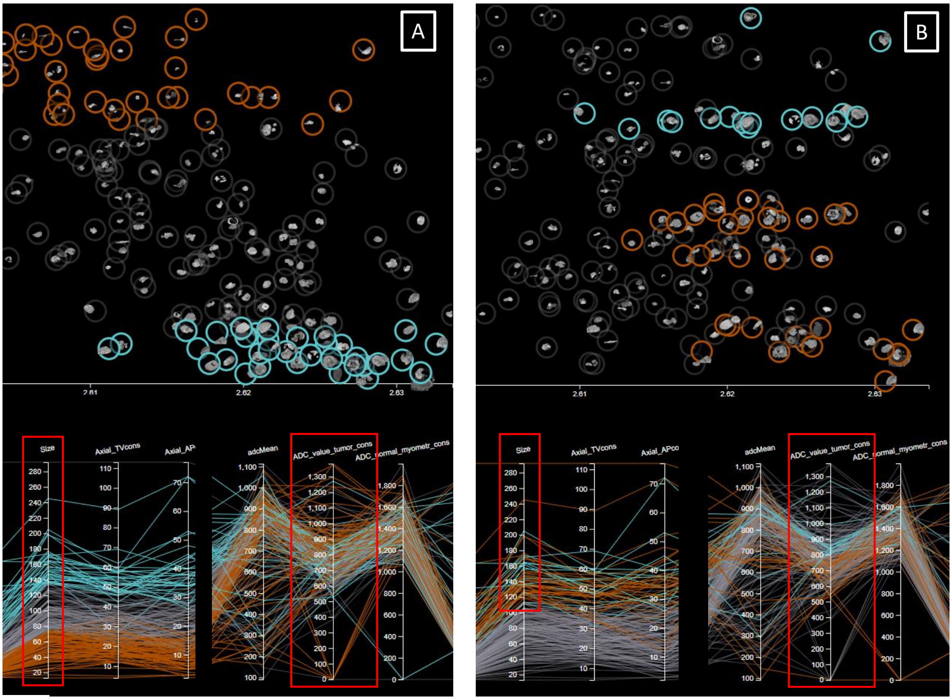

Figure B.5: A: t-SNE calcualation only taking size into account. B: Both size and ADC_value_tumor_cons are taken into account in the 1D dimensionality reduction. Selecting large tumors and different ADC value ranges for both groups indicate that the ADC value in large tumors could correlate with aggressiveness. A pattern is visible revealing in the scatterplot that patients are separated by ADC value.

tumor features (T5). To this end, we display a 1D dimensionality reduction result using t-SNE [163]. During our development process we also used a 1D principal component analysis (PCA) which delivered less convincing results. Therefore we chose t-SNE dimensionality reduction for our specific scenario, but this choice might not be the best option for other problem domains. Our scatterplot layout delivers an overview of the interplay between imaging data and clinical parameters for the purpose of radiomic tumor profiling. One example is presented in Figure B.2(a). In Figure B.5A only the size influences the t-SNE and in Figure B.5B both size and the ADC value of the tumor are considered. As the y values of the dots change, different groupings are visible. In Figure B.5A only size related clusters can be found while in Figure B.5B the ADC value has an influence and new clusters are present. This interactive dimensionality reduction enables hypothesis generation relating imaging and clinical parameters.

B.5.2 Parameter Overview

In addition to the imaging data represented in the central scatterplot view, clinical parameters are also a focal point in cohort analysis. These parameters consist of numerical, ordinal, and categorical data. Visualizing multiple data types together in one visualization can cause problems, because not all data types are compatible with all visualization idioms. Therefore, we opted for splitting these into two different visual representation. For the numerical data, we employ a parallel coordinate plot and for the categorical and ordinal data we use unit charts. As our tool was collaboratively developed with domain experts, we received iterative feedback on their ability to understand and work with selected visualization techniques during development.

Feature Dimension View A parallel coordinate plot (PCP) is an effective tool to analyze correlations between different feature dimensions [64]. Every patient represents one line in the parallel coordinate plot and each axis shows one feature dimension. The decision which axes/dimensions to use in the PCP is very important [64]. They serve as visual anchor and allow for use of ticks and descriptions. The ordering is also important, because it is difficult to compare dimensions which are further apart in the plot. Therefore, we decided to put specific axes next to each other where the correlation serves a specific purpose, e.g., the size of the tumor measured by the experts in the VIBE modality and the amount of voxels derived from the tumor mask. If the correlation between these two measurements is suspicious for certain patients, there might be something wrong with the data (T2). The PCP dimensions can also be selected for validation purposes specifically. If the user is performing segmentation or registration validation, the mean values of all modalities within the tumor are visualized. This enables detection of outliers, which may be caused by a misaligned segmentation mask (T1, T2).

Clinical Parameter View Unit charts are one of the simplest visualization methods and have already been described by Neurath in the early 1930's [112]. More recent work by Park et al. [117] states that this type of visualization can provide information that matches the user's mental model and allows for novel interactions. The unit chart representation is used for all categorical and ordinal parameters. Each dot in each column of the visualization represents exactly one patient. The color of each dot represents the value of the parameter for that specific patient. In addition, we use a tooltip to present imaging data when hovering with the mouse over the marks. Missing values are at the bottom of the chart and colored in dark gray. The other dots follow a quantitative grayscale colormap. The values in the unit chart are ordered according to expected outcome severity, meaning that values that have a negative influence on the outcome, e.g., life expectancy and suspected quality of life after treatment, are positioned on top of the chart (T3-5).

B.5.3 Settings and Interaction Techniques

Our application features a group selection feature where the user can select two different groups (T3). User can select which group is active and if selected patients should be combined using an 'AND' or 'OR' function. This allows for a detailed group selection. Using the 'AND' option, the user can select patients that, e.g., have multiple clinical parameters in common. In contrast, the 'OR' connection allows for selecting patients that, have a large tumor and myometrial invasion, but do not necessarily need to have both properties. The group selection interaction is supported across application views. Selection operations can be performed in the scatterplot by brushing with a rectangular selection box, in the PCP by selecting along an axis, or in the clinical parameter view by clicking one of the dots representing a specific value. Selections can be reset using a clear function. The user is further able to change the x- and y-axis properties. When changing the settings, the scatterplot updates with an animation, to improve context preservation. This allows users to locate parameter ranges to maximize selected target group separation. In the settings, the user is also able to change the modality presented in the tooltip and to swap the PCP to the registration and segmentation validation view.

B.6 Implementation

Our web-based application is composed of modules. The main part of the preprocessing, namely the data extraction and feature generation, is developed in Python. We use the PyRadiomics, a library developed by Gillies et al. [45]. The library supports first-order statistical features such as voxel-intensity histogram-based features, e.g., the median, the standard deviation or the maximum and the minimum value. In addition, also second-order statistics are supported. These include, for example, features based on the gray level co-occurrence matrix or the gray level run length matrix [45]. We use the Visualization Toolkit from Kitware [138], to create the tumor icons and the tooltip information for further analysis. Numpy [114] is used for working with the high dimensional radiomic tumor feature arrays and data handling within Python. We handle csv data handling using the Python library OpenCV [13].

The web-based part of our application is implemented in Javascript. The scatterplot and the parallel coordinate plot are both implemented using D3 [12]. Our implementation of the unit chart visualization is based on the approach by Park et al.[117]. For our dimensionality reduction, we use the t-SNE implementation TSNEJS provided by Karpathy [76], which is based on the original work from van der Maaten [163].

B.7 Case Studies

The RadEx application is visible in Figure B.3 and consists of multiple components as described in the Section B.5. To demonstrate the utility of our application, we showcase its functionality in three case studies, developed in close collaboration with our collaborators. We identified three major areas of application for our tool, namely the exploration of tumor characteristics, the co-registration validation functionality, and the machine learning segmentation check. Data of 330 patients is provided by two of our co-authors. Before including them in the application 12 patients have been excluded because the co-registration did not work, due to imaging quality problems. 97 patients were excluded because the machine learning based segmentation mask was completely misplaced or much too small. In the end 221 patients are included for further analysis. For 92 patients a manual created segmentation mask is available and for 129 patients a machine learning created segmentation mask is used. For every patient seven MRI sequences created by either a 1.5 Tesla or 3 Tesla MRI scanner, ten clinical parameters, and five measures from radiologists are available.

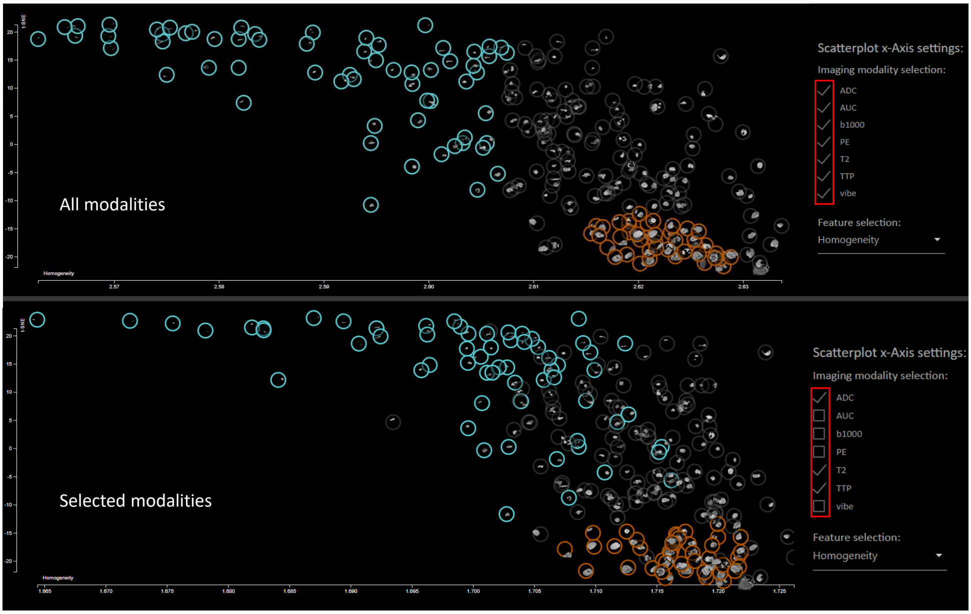

Scatterplot x-Axis settings:
Imaging modality selection:
☑ ADC
☑ AUC
☑ b1000
☑ PE
☑ T2
☑ TTP
☑ vibe
Feature selection:
Homogeneity

Scatterplot x-Axis settings:
Imaging modality selection:
☑ ADC
☐ AUC
☐ b1000
☐ PE
☑ T2
☑ TTP
☐ vibe
Feature selection:
Homogeneity

Figure B.6: Textural tumor features include several homogeneity measurements believed to correlate with tumor aggressiveness. We support a selection of homogeneity measurements and combine up to seven MRI sequences resulting in one measurement presented on the x-axis. Changes in parameter selection result in animated transitions.

B.7.1 Explorative Radiomic Tumor Profiling

Radiomic tumor profiling involves the calculation of high-dimensional feature vectors that need to be analyzed to discover tumor characteristics that are marker for possible outcome or to evaluate existing ones. Typical use cases in tumor profiling include homogeneity analysis of the imaging data and analyzing the association between various radiomics features with respect to different clinical parameters. In our application clinical experts can select two groups, one with low aggressiveness and one with higher aggressiveness. Aggressiveness can be measured, e.g., by presence of metastases or if the patient has already died from disease or had recurrence. Also, the time between the surgical removal of the tumor and a possible recurrence is an indicator for aggressiveness.

During the exploration of different settings for the x-axis, the group selection stays the same and gives a clear picture if the group separation improves or not, Figure B.6 reveals two possible group separations using different x-axis values. Allowing the user to try out different homogeneity measurements enables interactive exploration and assessment of the differences between each of the measures for this patient cohort. While exploring different measures, e.g., the homogeneity measurement based on the normalized inverse difference moment as described by Gillies et al. [45], we found a separation of a group of patients that features a low aggressiveness while having a low homogeneity.

Preliminary exploration revealed that the derived textural homogeneity can separate patients with high-risk disease from these with low-risk disease, demonstrating the usefulness of this tool for identification of imaging markers to be further explored. After selecting different feature combinations for dimensionality reduction, we are able to find different patient clusters that have different distinctive clinical parameters, such as for example the cluster shown in Figure B.6 in orange. This cluster features patients presenting with large tumors, which is known to be associated with increased risk of metastases and death from disease. This association is already well known, and our application can show this link. The status for these patients tells that they are either alive with active disease or dead from disease. The presence of metastases which evolve from the primary tumor is also an indicator for an aggressive tumor phenotype.

B.7.2 Registration Validation

Co-registration of multi-parametric images is an essential part of the data processing step needed for our application. The result of a successful co-registration is a set of perfectly aligned volumes. When dealing with high number of patients where most of the co-registration works well and which fails only in a small portion of cases, it is important to support visualization of likely outliers or error cases. To support such a filtering, the dimension selection visualized in the feature dimension view in the lower center of our application can be used to support the registration check. When doing so, the chart visualizes the median values of the tumor in each of the seven imaging parameters. This view allows for a quick check of outliers in the graph which might highlight segmentation masks that do not mask the tumor adequately due to misaligned volumes. The segmentation is then also misaligned because the segmentation mask is only available for one modality in our case. One example for such an operation is shown in Figure B.2(a), where outliers in the ADC value range revealed potential errors in the co-registration.

B.7.3 Segmentation Validation

Tumor segmentation masks crated by machine learning algorithms need to be validated by experts before using them in the data analysis step. Before validating all segmentation masks, experts could first find outliers where the algorithm did a bad job and those could be used to further improve the algorithm. There are multiple ways to find these outliers. One way is to validate if the number of voxels in the tumor mask align with the size measured by the clinical experts. Both parameters are present in the feature dimension view in

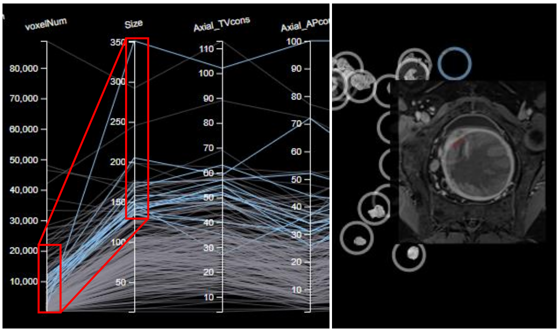

Figure B.7: One possible measure to find incorrect segmentation masks is the size, which in our case is measured in our case by radiologists and our application preprocessing phase based on the VIBE image and on the segmentation mask. These measures have a natural correlation and unexpected relations might indicate incorrect segmentation.

neighboring positions. Selecting a small number of voxels and a larger size measured by the experts or vice versa points to potential error cases. To validate if the segmentation is deficient, the user can hover over highlighted glyphs in the scatterplot and slice through the imaging volume. In the view presented in Figure B.4, we use a semi-transparent red overlay of the tumor segmentation mask on top of the various imaging sequences. Another possibility to find potential misaligned tumor segmentation masks is to select measured ADC mean values from our radiomics approach and compare them to the representative ADC value measured by the clinical experts. The experts indicate one representative region within the tumor and one region in the healthy endometrium to compare these throughout patients. If there is a major discrepancy, these cases should be investigated more closely. One such case is shown in Figures B.2(a) and B.2(b). A third method to inspect segmentation quality is to inspect the segmented tumor symbols in the scatterplot. When selecting very large tumors and seeing very small symbols on the scatterplot dots, it is an indication that the segmentation may be incorrect. Similarly, any discrepancy between selected features and visible segmentation symbols would lead to detection of questionable segmentation quality. Figure B.7 reveals such a case. The segmentation validation feature of our tool is of major interest for the machine learning experts working on implementing automatic segmentation algorithms since it allows for quick and intuitive validation that spurs on further development of automated methods.

B.8 Evaluation

To further evaluate the utility of our interactive exploration and analysis platform, we invited six gynecological cancer imaging research experts to validate our tool. E1 is a professor in radiology and expert in gynecologic and abdominal radiology. She has over 14 years experience in this field and is one of the co-authors on this paper. E2 is a radiologist since 2006 and holds a medical doctors degree. She is currently a PhD student in gynecological imaging, has over 10 years of experience in MRI reading, and is a co-author of the paper. E3 is a medical physicist in radiology and is a PhD student in medical physics with over 13 years of experience. E4 holds a masters degree in cell biology since 2010 and a PhD in neuro-oncology since 2015. E5 has 4 years of experience in pelvic imaging and holds a medical doctors degree. E6 has 5 years experience in MRI reading of gynecological cancer, holds a medical doctors degree, and is currently a PhD student.

In the beginning of the evaluation we demonstrated the application to the experts worked through the different use cases. Our application works with data provided by E1 and another co-author which was not part of the evaluation. Afterwards, we invited them to try out the tool themselves. During the evaluation we asked the experts to discuss their experience and to talk about benefits and disadvantages of the system compared to their current workflow. After this phase, which took roughly 40 minutes we asked the experts to fill out a questionnaire with 34 questions discussion different aspects of our application. The questions are structured in the following groups: general (G1-7), tumor visualization (V1-5), group selection (S1-7), homogeneity (H1-5), dimensionality reduction (D1-5), and segmentation and co-registration validation (C1-5). In addition to our evaluation form, the experts filled out the system usability scale

Table B.1: Response of the experts on a 5-point Liker scale. The meaning of the values on the scale are: 1: strongly disagree, 2: disagree, 3: neither agree nor disagree, 4: agree and 5: strongly agree. Statements marked with a star were rephrased to present the positive form in this table and the scores have been inverted. On the right end of the table the average value over all experts is presented and in the last row the result of the system usability scale questionnaire is presented. [1]paper co-authors.

	Statements:	$E1^{1}$	$E2^{1}$	E3	E4	E5	E6	Avg.
G1	The linked interactions between the scatterplot and the parallel coordinate plot are well established and intuitive	4	5	5	5	5	4	4,67
G2	The linked interactions between the scatterplot and the unit chart are well established and intuitive*	5	5	5	5	5	4	4,83
G3	The selection interactions between the unit chart and the parallel coordinate plot are well established and intuitive	5	5	5	5	5	4	4,83
G4	I see myself using RadEx in the future	3	5	5	4	5	4	4,33
G5	I would like to contribute in the future development of the application*	5	5	5	5	5	5	5,00
G6	I would like to use RadEx for exploring clinical cohort data	4	5	5	5	5	4	4,67
G7	The export funcionality helps me to further analyze the group selections in my statistics tool of choice*	4	5	5	5	4	4	4,50
V1	The small tumor icon enables a quick comparison between the tumors of different patients	5	4	4	5	2	4	4,00
V2	The tumor icons give me more information than only the size of it*	5	4	2	5	3	4	3,83
V3	The tooltip allows me to analyze the imaging data and the tumor segmentation*	5	3	5	4	5	4	4,33
V4	The extended tooltip is helpful to validate the tumor segmentation	5	3	5	4	5	4	4,33
V5	The extended tooltip is helpful to validate the co-registration*	5	3	5	5	4	4	4,33
S1	The group selection in the scatterplot view is easy to understand and to carry out	5	5	5	5	4	5	4,83
S2	I can select specific patients an add them to an existing group in the scatterplot*	5	5	5	4	5	5	4,83
S3	Selecting specific properties in the Unit chart view is easy to understand and to carry out	4	5	5	4	4	4	4,33
S4	I can select patients having specific states in different clinical parameters*	5	5	5	5	5	4	4,83
S5	Specifying a patient group including multiple clinical parameter manifestations is easy	5	4	5	5	5	4	4,67
S6	Selecting a group in the parallel coordinate plot is easy to understand and carry out*	4	4	5	5	5	4	4,50
S7	The applications makes it easy to select two different groups	5	4	5	5	4	4	4,50
H1	The homogeneity imaging modality selection in combination with the group selection helps me to identify important modalities*	5	4	2	4	4	4	3,83
H2	Trying different homogeneity measurements is easy and fast*	4	5	5	5	5	4	4,67
H3	The animation of the data when changing settings helps me to track the changes	5	5	4	5	5	4	4,67
H4	Havin the important measure homogeneity on the x-Axis of the scatterplot makes interpretation of the visualization easy	4	4	4	5	4	4	4,17
H5	I can imagine using this application to formulate hypothesis for future studies about homogeneity*	5	5	4	5	4	5	4,67
D1	The y-Axis in the scatterplot shows me interesting clusters of patients	5	4	5	5	5	4	4,67
D2	The dimensionality reduction allows me to analyze multiple clinical parameters*	5	5	5	5	5	4	4,83
D3	The selection of dimensions taking into account for the y-Axis allows me to explore my clinical data*	5	5	5	5	5	4	4,83
D4	I can imagine using this application to formulate hypothesis for future studies	4	5	5	5	4	4	4,50
D5	Exploring patients that are clustered by the t-SNE is interesting and potentially valuable for further investigation*	5	5	5	4	5	4	4,67
C1	I can select machine learning performed segmentations and validate their correctness	5	1	5	5	4	4	4,00
C2	Selecting potentially wrong segmentations is possible	5	3	5	4	4	4	4,17
C3	Exporting wrongly segmented patients is possible*	5	3	5	5	5	4	4,50
C4	The co-registration view enables me to spot potential wrong co-registrations*	4	5	5	5	4	4	4,50
C5	The tooltip view helps me to validate segmentations and co-registration results	5	3	5	5	4	4	4,33
SUS	System usability scale result	75,00	90,00	92,50	95,00	87,50	77,50	86,25

(SUS) provided by Brook et al. [17]. All statements are evaluated based on a 5-point Likert scale. We also included negatively formulated questions.

B.8.1 Evaluation Results

The result of the evaluation is shown in Table B.1. Questions marked with a star have originally been negatively formulated and here we present them in their positive form. The results for these questions are also inverted. In general, the application got positive feedback overall. All experts would like to contribute to the future development of the application and 5 out of 6 experts would like to use the application in the future. The tumor icons and tooltips received overall a good feedback. One expert (E2) mentioned that the tooltip pictures could be enlarged. E3 mentioned that the size of the tumor icons made it difficult to perceive shape, however, other participants agree that shape is also visible in our design. E5 also mentioned that the size of the tumor icons makes it difficult to compare based on icons alone.

All experts are in favor of the group selection. All questions, except one have an average value of at least 4,5. Only selecting specific properties in the Unit Chart or clinical parameter view has a value of 4,33. Regarding this point, we received the feedback that the dots used in the chart are challenging to click on. The homogeneity view also received strongly positive feedback overall. E3 mentioned that it is difficult to prove if a modality is important or not and mentioned that the question is formulated too narrowly. However, she is still in favor of the functionality. Our dimensional reduction got the most positive feedback with all average values over 4,5. All experts could imagine using the 1D dimensionality reduction to analyze the cohort data to spot potential groupings of patients. The co-registration and segmentation validation also received positive feedback overall. One expert (E2) mentioned that the tooltip images could be larger to make the validation easier. Another possibility for a more detailed validation could be to use a second screen to show the imaging data of specific patients in a view more like what radiologists are used to (E2, E4-6).

System Usability Scale Scores Our SUS scores are presented at the end of the evaluation result Table B.1. The results range from 77,5 to 95. In average our application reached a SUS score of 86,2. Bangor et al. [5], introduced different ways of interpreting SUS scores including the acceptability range, a grade scale (like in education), and an adjective rating scale. Our acceptance rate is: Acceptable (best score), grade scale: A (best score), and an adjective rating of Excellent (best score).

B.8.2 Evaluation Conclusion

We conclude from our results that our application is valuable for experts in gynecological cancer imaging research. All statement groups received positive feedback and the experts think the features are useful. E1 already thinks about using our application in a research setup to further evaluate machine learning-based segmentation masks and to train radiologists to perform segmentations and compare them to segmentation masks created by experts. E4 can also imagine using the application to validate results with new imaging series and E6 would like to use the application for his cohort data. Overall, we can say that the application has substantial potential in gynecological cancer imaging research.

B.9 Discussion

Experimental group selections performed in our application revealed that there is potential in further analysis of different homogeneity measures to separate patients with high-risk disease from those with low risk. Our application is also able to show well known coherence patterns, e.g., between the size of the tumor and clinical phenotype. This shows us that there is a potential to influence future analysis steps in gynecological cancer research and that our application may have an impact in the targeted domain.

For the group comparison we do not offer a feature that determines statistical significance values to prevent p-value significance fishing. Any hypothesis formulated using our application should be validated using an independent study cohort. After performing segmentation validation with our tool, we were able to spot tumor segmentation masks that did not meet the acceptance criteria, e.g., due to the presence of multiple tumors within the same region. The co-registration validation also highlights cases which would not have been suspicious at first sight. The involved machine learning expert therefore also sees potential in working with our application to further refine his machine learning algorithms to deliver even better results. The 1D dimensionality reduction is in the current version only supported by t-SNE but could also be performed using other dimensionality reduction methods such as PCA.

B.10 Conclusion and Future Work

We present RadEx, an interactive analysis platform and workflow for medical researchers which supports integrated exploration of radiomics and clinical features. By using multiple linked views, interactive group selection methods and custom feature selections, we empower researchers to explore new and validate existing tumor biomarkers. Close collaboration with gynecological cancer imaging and machine learning research experts resulted in several case studies for our application. The benefits brought by our work range from validation of machine learning results to validating and exploring new tumor characteristics. Being able to selectively include and exclude different sequences from the radiomic tumor exploration makes the analysis transparent and intuitive for the medical researchers. Our experimental findings of different patient groups allowed visualization of e.g. well-established association between large tumor size and high-risk disease but also showed potentially interesting new associations. The evaluation of our application revealed a positive response from the target audience reflected both by the qualitative evaluation and SUS score.

Potential future developments include adding an intuitive way to perform a landmark based registration. This could help to easily correct the registration for patients where the automatic co-registration fails. In addition, we plan to widen the scope of the application to different tumor types including cervical cancer. The goal of our efforts is that every single patient should benefit from the findings from cohort analysis to get one step further in the direction of personalized medicine.

Acknowledgments

This research was funded by the Trond Mohn Foundation (Grant ’811255’, ’813558’, and ’BFS2018TMT06’), the Western Norway Regional Health Authority (Grant #912263), the Norwegian Cancer Society (Grant #190202), the Norwegian Research Council (Grant #273280), and supported by the Norwegian Research Centre (NORCE).

Paper C

ICEVis: Interactive clustering exploration for tumor sub-region analysis in multiparametric cancer imaging

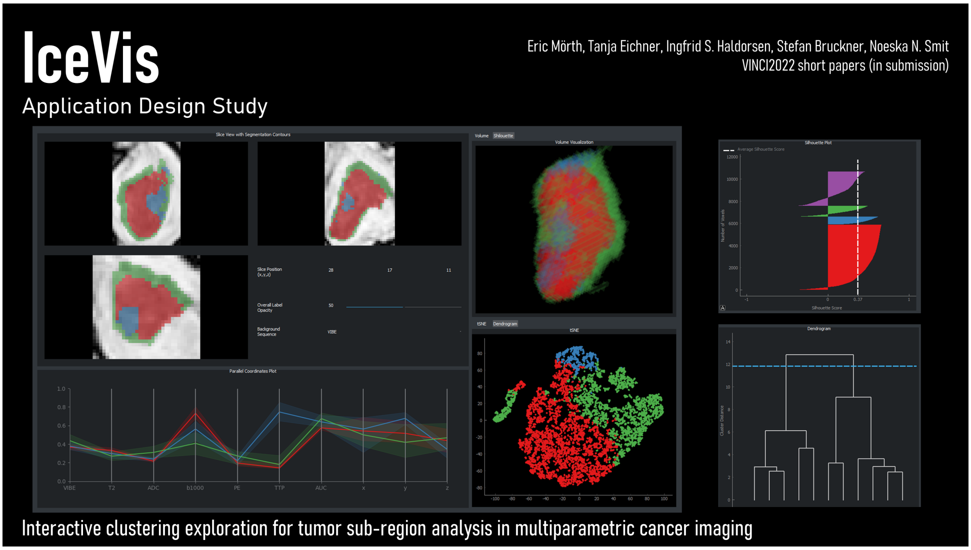

IceVis

Application Design Study

Eric Mörth, Tanja Eichner, Ingfrid S. Haldorsen, Stefan Bruckner, Noeska N. Smit
VINCI2022 short papers (in submission)

Slice Position (x,y,z) 28 17 11
Overall Label Opacity 50
Background Sequence VIBE

Interactive clustering exploration for tumor sub-region analysis in multiparametric cancer imaging

E. Mörth[1,2], Tanja Eichner[2,4], I. S. Haldorsen[2,3], S. Bruckner[1,2] and N. N. Smit[1,2]

[1]Department of Informatics, University of Bergen, Norway
[2]Mohn Medical Imaging and Visualization Centre,
Department of Radiology, Haukeland University Hospital, Norway
[3]Department of Clinical Medicine, University of Bergen, Norway
[4]Faculty of Informatics, Technical University Vienna, Austria

Abstract

Tumor tissue characteristics derived from imaging data are gaining importance in clinical research. Tumor sub-regions may play an important role in defining tumor types and may hold essential information about tumor aggressiveness. Depending on the location of the tumor within the body, such sub-regions can be easily identified and defined by physiology, but for others these sub-regions are not readily visible. Currently, exploration of regions within a tumor is performed via comparison of the image sequences and analyzing the tissue heterogeneity present. To improve the exploration of such tumor sub-regions, we propose a visual analytics tool called ICEVis. ICEVis supports the identification of tumor sub-regions and corresponding features combined with cluster visualizations highlighting cluster validity. At times, the ground truth in terms of the optimal number of clusters is not available. For such cases, we provide interaction possibilities to determine the optimal number of clusters, supported by various statistical measures and interactive exploration of the results. We evaluated our tool with three clinical researchers to show the potential of our approach.

C.1 Introduction

Cancer is globally the second leading cause of death according to a report published by the World Health Organization in 2018 [15]. Different tumor types exist, ranging from low-grade tumors with a favorable prognosis to high-grade tumors associated with high risk of relapse and cancer spread. Surgery is the most common treatment type. Early diagnosis and treatment planning are crucial for optimizing the therapy of endometrial cancer [2]. A challenge that can arise when dealing with endometrial tumors, as well as other types of cancer is tumor heterogeneity [2]. The number and composition of tumor sub-regions may vary between different patients, but their identification may be an important factor to consider for diagnosis, tissue sample (biopsy) selection, and/or therapy planning. At present, tissue heterogeneity is examined by the acquisition of histopathological data from tissue biopsies and, in recent research, by analyzing the tumor imaging data as a whole using radiomic tumor profiling [41, 62, 111]. Feature detection from medical imaging data alone could potentially replace invasive biopsies and improve the expected quality of life of the patients after treatment [128, 171]. The clinical imaging routine for endometrial cancer detection and therapy includes standard MRI sequences. In some cases, additional sequences sequences that might lead to a better distinction of tumor patterns are acquired for research purposes. This enables the consideration of different tumor characteristics that are visible in various imaging sequences [95]. To enable clinical researchers to explore and evaluate tumor sub-regions in cancer research, we propose ICEVis. ICEVis is part of an ongoing collaboration with two clinical experts and enables endometrial cancer researchers to interactively explore hierarchical clustering results based on their multiparametric medical imaging data. The main contribution is a visual analytics approach that supports the assessment

---

Submitted to the International Symposium on Visual Information Communication and Interaction (VINCI 2022)

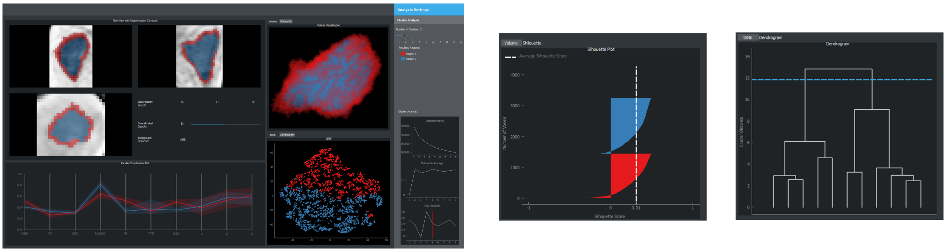

Figure C.1: ICEVis consists of a 3D tumor view and calculated cluster information. A parallel coordinate plot, a dendrogram, a silhouette plot, a t-SNE representation of the tumor segmentation, and three optimal cluster number plots are embedded to support the clustering exploration process.

of individual distinct tumor sub-regions and relevant features in single patients across imaging sequences. Our approach consists of a hierarchical clustering approach embedded in a visual analytics platform which allows for detailed tumor sub-part analysis. To confirm the utility of our approach, we conducted a qualitative evaluation including three experts in machine learning and gynecological cancer imaging research.

C.2 Related Work

In the medical domain, several clustering approaches deliver remarkable results in tissue classification. Juan-Albarracín et al. [71, 72] analyzed malignant brain tumors and managed to improve results obtained by supervised methods. They evaluated five different unsupervised voxel classification methods including K-means clustering, fuzzy K-means, Gaussian Mixture Models (GMM) as a generalization of K-means, and the Gaussian Hidden Markov Random Field (GHMRF) model. Ng et al. [113] used K-Means for medical image segmentation and additionally combined it with an improved watershed algorithm. Raidou et al. [128] applied dimensionality reduction and K-means clustering for cohort analysis of different tumors. Informed by these approaches, we analyzed our data by applying the same techniques but found the results to be unsatisfactory with our data. After careful consideration, we decided on using a hierarchical clustering approach instead.

Hierarchical Clustering: Selvan et al. [140] discussed how hierarchical clustering could aid diagnostic imaging data analysis. They outline examples such as the interpretation of x-ray mammography and multi-parametric prostate MRI. In the implementation, they chose the appropriate number of clusters by minimizing the dissimilarity within each cluster. In contrast to their work, we use several statistical parameters to define the optimal amount of clusters and support the user in the cluster number decision.

Visualization for Tumor Tissue Exploration: In 2019, Qu et al. [127] published a review of visual analytics tools for genomic and cancer data, comparing traditional and new methods, among them scatter plots, cluster visualization and networks. Existing systems for visual analytics in oncologic tissue exploration focus either on the exploration of perfusion data, or on the distinction between tissues belonging to the tumor and unaffected, healthy regions [128]. Yu et al. [172] presented an approach using a heatmap in combination with a dendrogram to analyze the validity of the calculated clusters. In contrast to their approach, we analyze tumor sub-regions instead of the whole tumor volume. Our work is inspired by the approach of Raidou et al. [128]. They propose a visualization tool for visual examination of the feature space resulting from pre-defined imaging-derived tissue characteristics. Similar to a related approach [129], they support the assessment of the effect on the parameter space due to different model-based variations in DCE-MRI. A central dimensionality-reduced scatter plot in combination with a density plot guides the interactive exploration. In comparison to their work, we enable clustering of the tumor imaging data before conducting a dimensionality reduction, which potentially leads to more accurate clustering results.

C.3 Medical Background

One of the main research goals of our collaborators is the exploration of tumor imaging biomarkers to improve predictions and personalized treatment strategies. Ongoing research aims to analyze relevant imaging features such as texture as well as to develop automatic methods for segmentation and classification [40, 57, 60, 171]. For this reason, the tumor is segmented prior to the analysis phase, which is currently done mainly manually, taking only one of the several sequences into account. Our collaborators acquired multi-parametric scans of 92 patients with endometrial cancer and performed manual segmentation. AvaiThe The seven available sequences for this work are the apparent diffusion coefficient (ADC), b1000, peak enhancement (PE), time to peak (TTP), area under the peak enhancement curve (AUC), T2 and volumetric interpolated breath-hold examination (VIBE). These multiple sequences provide complementary information and they are all aligned via co-registration. Together with two involved medical collaborators, we analyzed the requirements for our application:

- R1: Find potential tumor sub-regions based on imaging characteristics per patient
- R2: Support the decision of how many tumor sub-regions are most likely present in the given patient's data
- R3: Enable spatial analysis of the clustering results
- R4: Enable exploration of the results with respect to the given imaging characteristics

C.4 ICEVis

Based on the requirements defined together with our medical collaborators, we propose ICEVis. ICEVis consists of a pre-processing pipeline and a visual analytics platform to calculate potential tumor sub-regions. Furthermore, we enable clinical researchers in defining the optimal number of sub-regions based on imaging characteristics. The pre-processing of ICEVis consists of a hierarchical clustering step which takes all available

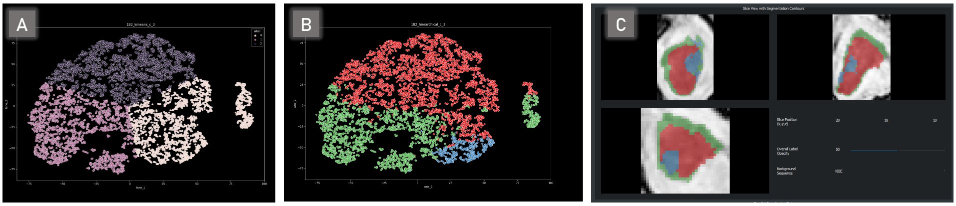

Figure C.2: A: a t-SNE result based on all available sequences with subsequent k-Means clustering (k=3). B: hierarchical clustering applied before dimensionality reduction and visualized in the same t-SNE result. C: The hierarchical clustering result in a slice-based view as an overlay on the medical imaging data. The last image shows a necrotic core within the tumor in blue which would no have been found by using k-Means on the t-SNE result.

sequences into account and is carried out using the full dataset without prior dimensionality reduction. This step fulfills requirement **R1**. Informed by prior work, we analyzed the results of different clustering techniques (k-Means on the full dataset, k-Means on the t-SNE result, and hierarchical clustering on the full dataset) and concluded that hierarchical clustering on the full dataset delivered the most promising results. One example of our analysis is presented in Figure C.2, where k-Means clustering on the t-SNE result is juxtaposed with the hierarchical clustering result before applying t-SNE. As an example for the benefits of our approach, a histopathologically relevant part of the tumor (a necrotic core) is depicted in blue – this important feature was not identified correctly when using t-SNE followed by clustering.

Hierarchical Clustering To perform hierarchical clustering, different sub-steps are necessary. We first pre-process the data and create feature vectors based on tumor segmentation voxels. Then, we calculate a distance measure by defining the criterion according to which the clusters are split. We normalized the data using min-max normalization. The voxel intensities of each of the seven sequences represent the feature vectors. To compute cluster proximity, different metrics are available. In this work, Ward's method [110] was used. It is an centroid-based approach which defines the proximity between two clusters by calculating the increase of the sum of squared errors (SSE) when potentially merging those two clusters. Similar to K-means, it minimizes the sum of the squared distances of data objects from their corresponding group centroid.

Determining the Number of Clusters To support the user in determining an optimal number of clusters and therefore to fulfill requirement **R2**, we present the user with the result of three different measures: the elbow method, the average silhouette method, and the gap statistic. The elbow method is a function of the within-cluster sum of squares. The optimal number of clusters here is the amount where adding another cluster would not improve the result [154]. Kaufman and Rousseeuw [77] introduced the average silhouette method, defining the optimal number of clusters as the one which maximizes the average silhouette score. The last method we support is the gap statistic method by Tibshirani et al. [156]. They propose the comparison of the within-cluster variation with the expected variation under a reference null distribution. The optimal value is

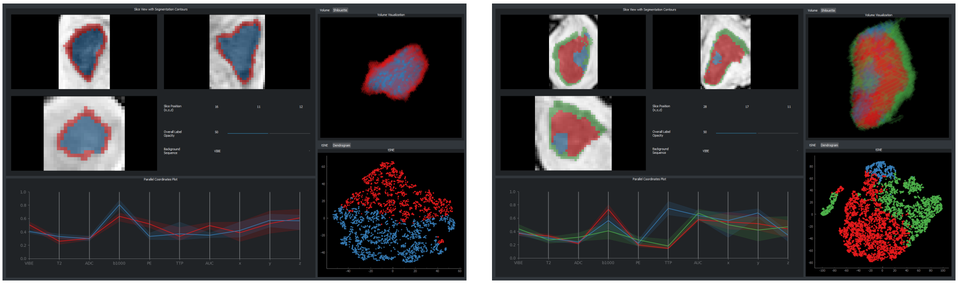

Slice Position (x,y,z) 16 11 12
Overall Label Opacity 50
Background Sequence VIBE

Slice Position (x,y,z) 28 17 11
Overall Label Opacity 50
Background Sequence VIBE

(a) In this particular example, the optimal amount of clusters is two. According to our collaborators, the red area represents an inflammatory region around the tumor.

(b) In this example, three clusters differ in ADC, b1000 and in TTP according to the parallel coordinate plot. The blue cluster highlights the necrotic core of the tumor, a region consisting of dead cells within the tumor.

Figure C.3: ICEVis use cases presenting interesting cases defined by our clinical collaborators. (a) is showing a necrotic core and (b) is presenting a potential inflammatory region of the tumor.

derived by the maximized gap statistic, meaning the cluster structure which has the biggest difference to a uniform distribution of points. As there is no definitive answer to the question of how many clusters is the optimal solution, we provide the expert with an ensemble of different methods to support the decision process. In Figure C.1 on the right side, of the first image all three methods are visualized with the optimal number of clusters highlighted. Note that there is not necessarily agreement among these three methods.

Visualizations Our tool consists of spatial and non-spatial visualizations enabling medical experts to find the optimal number of tumor sub-regions as well as inspecting how they change in the imaging data. In close collaboration with our clinical experts, we found these visualization idioms to be the most effective and meaningful to use within the proposed application. To enable the analysis of the resulting clusters with respect to the underlying imaging data, we employ three orthogonal slice views and a 3D visualization of the data. The clusters are presented as an adaptable semi-transparent overlay and the underlying medical data is presented in a grayscale colormap. Users can choose which sequence they would like to see. The 3D visualization reveals cluster information in the same colors used in the rest of the tool. This visualization can be used to analyze the spatial relationship between the clusters. These two visualizations fulfill requirement **R3**. To fulfill requirement **R4** and therefore to characterize clusters by their imaging data content, we add a parallel coordinates plot (PCP). Each feature vector is represented as a polyline and the axes correspond to the individual features. The PCP is further enriched inspired by box plots [104]. We chose this design over drawing a line for each voxel within the cluster to avoid visual clutter. For each region, the median value of voxel intensities is drawn on the axis of the single modalities. Their connection forms a line and is visualized in an opaque manner. The more transparent bands around the median lines represent the inter-quartile area to give a better impression of the value range within the clusters.

Cluster Representation To ensure a meaningful representation of each clustered region, we assign a specific color which is consistent throughout the different views within the application. We chose the colors according to a recommendation from the commonly used *ColorBrewer* [58] tool. To visualize the hierarchical clustering, we created a dendrogram, as shown in Figure C.1 on the right side. It shows the hierarchical structure as a tree with distance-based connections between sub-clusters. The length of the vertical lines represents the distance between those clusters, measured by using the Euclidean distance [155]. Furthermore, a cutting line is included based on the selected number of clusters. The vertical lines that are crossed by the cutting line represent the associated clusters. In addition, the clustering result is presented on top of a t-SNE embedding of the imaging data within the tumor. The feature vectors for the embedding are the seven image sequences which are available for exploration. This allows the medical experts to see if there is a clear difference in clustered voxels within the tumor.

C.5 Implementation

Our tool was implemented in Python. For data handling and working with high-dimensional feature arrays, we use the *NumPy* library and for hierarchical clustering we employ the computationally efficient library *fastcluster* [106]. The dendrogram and label extraction are calculated using the *SciPy* [69] cluster hierarchy package. The application was developed with *PyQt5* [28] using the *matplotlib* [63] library and the *PyQtGraph* [21] package for visualizations. Furthermore, the 3D visualization uses *pyvista* [152].

C.6 Results

Figure C.1 depicts the complete user interface for our tool. In the analysis settings on the right sidebar, the number of clusters can be selected via a slider. In this case, two clusters have been selected. Below, detected regions with their assigned colors are listed. A slice view gives the possibility to scroll through the slices with a segmentation mask overlay in x-, y- and z-direction. The slice position is shown numerically in addition and updated when scrolling through the volume. To examine the regions or the imaging data in more detail, the opacity of the segmentation can be adjusted with a slider or by entering a percentage in the corresponding text field. The background sequence can be changed to one of the seven available sequences. The dendrogram, silhouette plot, 3D view, PCP and t-SNE view, as introduced in Section C.4, assist in the exploration and update when changing the number of clusters.

Case Studies and Expert Feedback We describe two different case studies for our tool based on expert feedback. The first interesting case we found is shown in Figure C.3(a). In this case we found a shell around the tumor. Our experts hypothesized that this region may comprise an inflamed region on the outside of the tumor. This result could have an impact on the treatment decision and may lead to an earlier surgery to avoid negative effects when waiting too long. In the second case study presented in Figure C.3(b), our medical experts found an optimal number of three clusters and one of them comprises a necrotic core, which is a distinct pattern within the tumor composed of dead tumor cells. For the evaluation of our application, we invited two radiologists involved in gynecological cancer imaging research and one expert in machine learning for medical imaging data to give us feedback on our approach. All invited experts are not co-authors of this paper but are part of ongoing collaborative research. In general, we got positive feedback on our application and the experts found interesting and compelling use cases for our application. The two radiologists stated that they can recognize histo-pathological properties which would have to be confirmed by correlating our results with corresponding surgical findings. When this pattern is confirmed, our application could have an essential impact as such histo-pathological properties could then be taking into account at the time of imaging and treatment decision. Furthermore, the radiologists would like to use our tool to find areas of the tumor which seem to be the most aggressive parts based on imaging data to consider these for further analysis, e.g., targeted biopsies. The machine learning expert sees a potential use case in further analyzing the results of our approach in a radiomic tumor profiling setup to further correlate to clinical variables, e.g., aggressiveness of the tumor. Overall, all three experts find the tool useful for different purposes and would like to use it for analyzing their data.

C.7 Conclusion

We proposed ICEVis, an interactive clustering exploration tool consisting of a hierarchical clustering approach embedded in a visual analytics application. The tool includes visualizations for the exploration of clusters and their validity, as well as the characterization of cluster content. We developed our application in close collaboration with cancer imaging experts and conducted a qualitative evaluation including three experts from different fields. Our results show two interesting cases where the clinical experts where able to find histo-pathological interesting regions within the tumor which they were not able to see in the sequences before. Furthermore, they would like to evaluate the results by comparing the visual results with the histo-pathological results after the surgery of the tumor. In a future continuation of this work, we would like to enable the clinical experts to analyze the results in an even more interactive way by improving the linking and brushing capabilities of ICEVis. Finally, this work represents a further step towards patient-specific tumor sub-region exploration for endometrial cancer, and could, in the future, lead to opportunities for more personalized treatment.

Paper D

ScrollyVis:
Interactive visual authoring of guided dynamic narratives for scientific scrollytelling

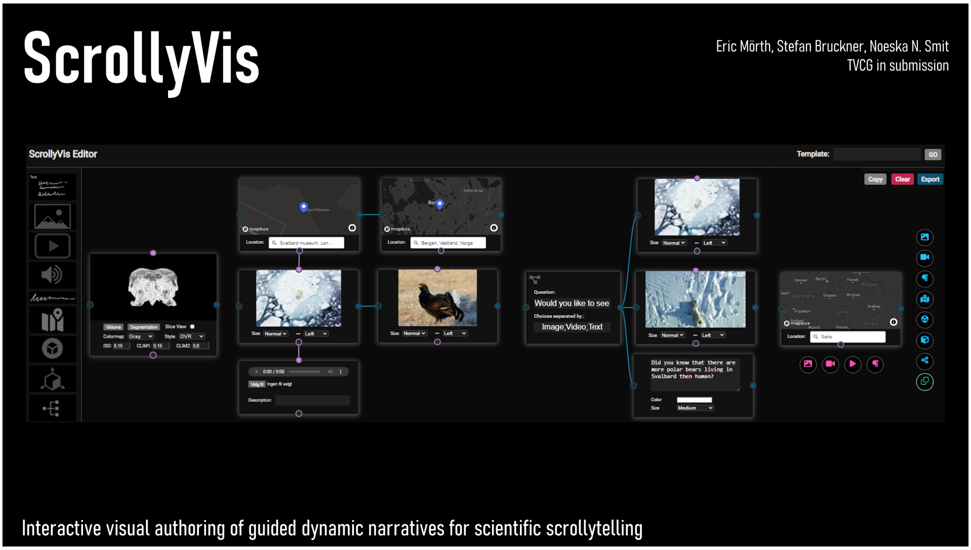
ScrollyVis

Eric Mörth, Stefan Bruckner, Noeska N. Smit
TVCG in submission

Question:
Would you like to see
Choices separated by ;
Image;Video;Text

Did you know that there are more polar bears living in Svalbard then human?

Interactive visual authoring of guided dynamic narratives for scientific scrollytelling

E. Mörth[1,2], S. Bruckner[1,2] and N. N. Smit[1,2]

[1]Department of Informatics, University of Bergen, Norway
[2]Mohn Medical Imaging and Visualization Centre, Department of Radiology, Haukeland University Hospital, Norway

---

This article is submitted to *TVCG*

Abstract

Visual stories are an effective and powerful tool to convey specific information to a diverse public. Scrollytelling is a recent visual storytelling technique extensively used on the web, where content appears or changes as users scroll up or down a page. By employing the familiar gesture of scrolling as its primary interaction mechanism, it provides users with a sense of control, exploration and discoverability while still offering a simple and intuitive interface. In this paper, we present a novel approach for authoring, editing, and presenting data-driven scientific narratives using scrollytelling. Our method flexibly integrates common sources such as images, text, and video, but also supports more specialized visualization techniques such as interactive maps as well as scalar field and mesh data visualizations. We show that scrolling navigation can be used to traverse dynamic narratives and demonstrate how it can be combined with interactive parameter exploration. The resulting system consists of an extensible web-based authoring tool capable of exporting stand-alone stories that can be hosted on any web server. We demonstrate the power and utility of our approach with case studies from several of diverse scientific fields and with a user study including 12 participants of diverse professional backgrounds. Furthermore, an expert in creating interactive articles assessed the usefulness of our approach and the quality of the created stories.

D.1 Introduction

Storytelling is deeply embedded in our society. Its purpose ranges from informing and recording to explaining and entertaining. Stories play a major role in understanding the world, cultural identity and to trigger and explain emotions. Stories can be told in many different ways including passive and interactive forms as well as linear and nonlinear methods [80]. Filmmakers explore a vast range of methods to convey a story in the best possible way. Quentin Tarantino is an example of a film director who is known for exploring new ways to present a story and therefore his movies are frequently discussed [80]. Stories are an important aspect not only of movies but also in the context of video games, books, and articles. Newspaper article authors engage the viewer by including graphics and images but on the web they can even add interaction [135].

With the rise of the internet, storytelling became a part of conveying information in online media. In more recent times, storytelling on the web was re-imagined into so-called scrollytelling [141]. Viewers interact with presented information not by clicking through a web slideshow but by simply scrolling through a website. The author can plan the flow of information and is able to guide the reader through the experience. Scrollytelling is extensively used by news outlets such as the New York Times [119, 121] as it can be an engaging and effective way to present information. On mobile devices in particular, scrollytelling is one of the standard interactions [119] and heavily used by social media platforms and various online media.

Storytelling can be a powerful tool to convey scientific information [89]. Science has to compete with other storytellers, many of whom are not bound to scientific evidence [32]. Blastland et al. [10] referenced the philosopher Onora O'Neill who said: "inform but not persuade" and when sharing evidence, suggested to strive "to be accessible, comprehensible, usable and assessable". Scientific communication must preserve its credibility, but also needs to engage the audience with compelling communication formats. Dahlstrom [31] noted that storytelling often has a bad reputation in science and that there is even a mantra saying "The plural of anecdote is not data". He proposed the following adaption to this statement: "The plural of anecdote is engaging science communication". Scientific storytelling can leverage data visualization to convey outcomes and facilitate reasoning about scientific results, serving goals such as the communication with peers in the field or engaging a wider community.

Hohman et al. [61] demonstrated that interactive articles can be used to present the latest progress in various research fields and to make the findings accessible and understandable to a broad audience. The challenge in doing so is that often there is no clear incentive structure as well as little funding for research dissemination and communication. Interactive and engaging articles are most viable on the web [61]. One challenge is that not all scientists possess web-development skills. There is a variety of editors and content management systems to create basic websites, but interactive and engaging storytelling is not supported sufficiently in these. As demonstrated by the user evaluation of Seyser et al. [141], a whole team including authors, designers, and developers is needed to create such rich experiences. In terms of dynamic narratives, authoring tools facilitating storyboarding are currently available, but scrollytelling support is lacking. At present, there is a missing link between dynamic narratives and scrollytelling presentation on the web.

Our approach aims to fill this gap. With this paper, we present ScrollyVis, an extensible web-based authoring tool for creating guided dynamic narratives with a particular focus on scientific narratives. ScrollyVis offers a processing pipeline which exports authored stories such that they are ready for deployment on any web server. We support dynamic as well as static narratives. Our approach enables users of all technical skill levels to create scrollytelling web experiences with ease. We allow for the integration of a broad range of visual media, such as images, videos, and map views, but also accommodate more advanced visualization techniques such as direct volume rendering, slice-based visualization and 3D surface-based visualizations. To verify the utility of our approach, we present a user evaluation including 12 participants with diverse professional backgrounds. Furthermore, we present four case studies in collaboration with experts from a range of scientific disciplines and a qualitative evaluation of our approach and the resulting stories by an expert in interactive storytelling on the web. Our main contributions in this paper are:

- We introduce a system which allows for efficient authoring, generation, and presentation of dynamic media-rich scrollytelling experiences on the web for scientific communication.
- We present ScrollyVis, a prototype storyboard-based editor realizing the system description enabling users to author and publish dynamic narratives on the web without requiring prior web development skills.
- We demonstrate the power and utility of the approach through case studies and evaluate the usability of the editor and potential through user feedback.

D.2 Related Work

In this section we describe work related to our approach, such as storytelling in visualization and visual exploration. Furthermore, we reflect on scrollytelling as a web-based long-form article, and storytelling editors.

Storytelling in Visualization Storytelling is a focus in visualization research over a longer period of time. Wohlfart et al. [170] combined storytelling with interactive volume visualization to enable a better understanding of the underlying data and information. They introduced an authoring and a storytelling step to separate exploration and presentation of a story. Their approach employed Shneiderman's [142] information seeking mantra and proposed a taxonomy for interaction with the user. Furthermore, they presented a story model consisting of story nodes and story transitions. Building up on their efforts, we introduce story segments as a higher-level abstraction of semantically similar story nodes.

Kosara et al. [89] argued that stories are a good way to present data as they package important information and knowledge in an easily understandable way. They highlight that interaction is one of the most important aspects of visualization, including altering the pace and direction of the story. Ma et al. [100] reflected on the question of what "good pacing" means when it comes to scientific storytelling using visualizations. Every spectator has their own pacing preference and there is always a compromise when introducing a fixed pace to tell a story. The authors furthermore emphasized the importance of user domain knowledge and relevance of the story to the users. The paper is especially relevant for our approach as it guides our design in pacing and adapting stories to a target audience. Tong et al. [158] describe common visualization types used in storytelling. Based on their analysis, we designed our system to support all of these visualization types in order to make them readily available for scientific communication. In the following, we present visualization approaches which specifically aim to tackle challenges related to scrollytelling.

Scrollytelling According to Pimbaud [121], scrollytelling has been around since 2010 and is linked to the success of social media. He further speculates that scrolling might be the easiest user interaction possible. Seyser et al. [141] claim that scrollytelling is the web equivalent to long-form articles used in journalism. Scrollytelling presentations frequently consist of multimedia content and information visualization in particular. The authors stated that the narrative structure of scrollytelling articles is either linear or elastic. The latter enables the user to dive deeper into the story on demand. Scrollytelling articles often use at least three different multimedia elements including photography, videos, and visualizations [50, 162]. Pettersen [119] proposed that storytelling is not only about the words but presenting information in the most interactive and exciting way. Furthermore, she mentions that scrollytelling is beneficial to engage and actively keep the viewer's attention during story consumption. Scrollytelling gives a sense of control, exploration, and discoverability.

Unfortunately, to the best of our knowledge, there are no publications on concrete guidelines for best practices in scrollytelling. However, practical recommendations are available from various blogs. For example, McKinley [120], principal engineer at Etsy, shared that endless scrolling is unfavorable and resulted in lower sales at Etsy. Furthermore, Seyser et al. concluded that the Gestalt laws should be applied to structure visual elements in order to make them understandable to the viewers. Kosara [88] provides an extensive overview of what he considers bad scrollytelling design features and demonstrates that many aspects can easily annoy and distract viewers of scrollytelling websites. Kosara mentions that the user should know upfront how long the story is going to be and that direct access to different parts of the story is favorable [88]. Scrollytelling content needs to be carefully designed. Transforming existing websites by simply adding scrolling and unveiling content on demand often ends up in an unsatisfactory user experience. Bostock [11] proposed scrolling implementation guidelines, which equip us with five rules to follow in order to create an effective scrollytelling website, namely: *(1) Prefer scrolling to clicking, (2) allow rapid, incremental, reversible scrolling, (3) provide instantaneous consistent feedback, (4) avoid unwanted disruptions and (5) support standard keyboard controls.*

Narrative Authoring and Storytelling Editors According to Hohman et al. [61], creating interactive articles today is still difficult. It is often closer to building a website than to writing a blog post or article. It also takes considerably more time than writing a static article or even a scientific publication. According to Conlen et al. [29], media such as the New York Times, Washington Post, the Guardian, and FiveThirtyEight, provide high quality multimedia narratives often referred to as interactives. The authors noted that the data visualization community suggested research opportunities in creating tools for driving the production of such interactive narratives. Creating or recreating interactives as presented by high quality newspaper providers is complex and involves several experts [29].

Tableau [42] implemented their story feature by using story points. Story authors can present interactive visualizations created in Tableau in a slideshow manner. While the visualizations within one story point are interactive and allow for exploration, the transition between points is not dynamic. In comparison to their approach, we focus on smoothly animated transitions that change visual elemements in an incremental and reversible manner. Furthermore, we allow for 3D data visualization, such as volume and surface visualizations. Kouřil et al. [90] presented a novel way to prepare story structures and automatically create concrete narratives for molecular documentaries. They present a technique called *story graph foraging* and techniques for real-time narrative synthesis. In contrast to their approach, we provide our stories on the web and not as video and have a broader focus than molecular visualization. *VizFlow* by Sultanum et al. [153] instead focused on data-driven articles. The authors used text-chart linking strategies to create scrollytelling experiences and evaluated their approach with 12 authoring and 24 reading participants. Compared to *VizFlow*, we do not focus on data visualization exclusively and allow for more extensive authoring opportunities. In addition to narrative videos and text enriched with data visualization, interactive data comics are another approach for presenting scientific insights. Wang et al. [167] presented a lightweight specification language entitled *Comic Script* to create interactive and dynamic data comics. The approach supports branching, change of perspective, and details-on-demand, as interaction methods for the viewer. The approach has similar dynamic capabilities as our approach, but leverages a slideshow-like format, employing point and click user interactions. Furthermore, we enable the visualization of more complex data types such as maps, surfaces and volumes.

Two well known software applications for creating dynamic stories are *Stornaway* [151] and *Twine* [161]. *Stornaway* features a node-link editor which allows for the creation of dynamic interactive videos. *Twine* is an open-source tool for telling interactive, nonlinear stories. *Twine*'s editor is based on a node-link diagram and is designed for the creation of interactive fiction. *Twine* exports the story directly to HTML. In comparison to these two applications, our approach enables the creation of stories including various additional media types and creates a scrollytelling website rather than a click-based website or video-based presentation, as is the case with *Twine* or *Stornaway*, respectively. Furthermore, our editor includes a preview of all different media content and introduces a hybrid approach between node-link diagrams and storyboards. In contrast to these approaches based on graphical user interfaces, Satyanarayan et al.[135] introduced a system which combines a domain specific language (DSL), Ellipsis, with a graphical user interface-based story authoring tool. The authors contributed a model for narrative visualization which helps story authors who may not be familiar with web development to convey their stories as websites. They evaluated their approach with a qualitative user study with feedback from journalists. The journalists where positive overall but mentioned that a node-link interface would be a good way to author a story. Furthermore, the journalists asked for an easier way to present non-linear stories. In contrast to this related work, we use scrollytelling to guide the user through the story. Furthermore, we provide a node-link authoring tool and support guided dynamic narratives.

The closest previous work to our approach is *Idyll* [29, 30] which consists of a markup language for authoring and publishing interactive articles on the web. This approach is based on a DSL designed for authoring interactive narratives combining a markup language and in-line JavaScript components. *Idyll Studio* [30] has a graphical user interface that lowers the threshold for non-experts to create interactive articles. Compared to *Idyll*, we do not focus on creating a DSL and rather concentrate on editor functionality and flexible support for different visualizations methods. Furthermore, in addition to using parameter-based interaction we allow for dynamic narratives where the viewers can choose pre-defined narrative paths.

D.3 Scrollytelling and Narratives

Storytelling is an effective way of conveying information and knowledge [96]. According to Joubert et al. [70], storytelling is the soul of science communication. Storytelling and especially scrollytelling is increasingly used on the web. This trend is also reflected in news outlets such as The Economist, the BBC, the New York Times and German science magazine Substanz. Two examples of such stories are Unearthing the Truth [37] and Genexpressionen [98]. In addition to stories created by news agencies, Apple research recently started to present their latest research in form of interactive scrollytelling based web pages, e.g., a story about interpretable adaptive optimization [148]. All of these stories are a joint effort between authors, designers and programmers to enable media-rich and interactive scrollytelling experiences.

Across these sample stories, we have identified several common patterns. All of the

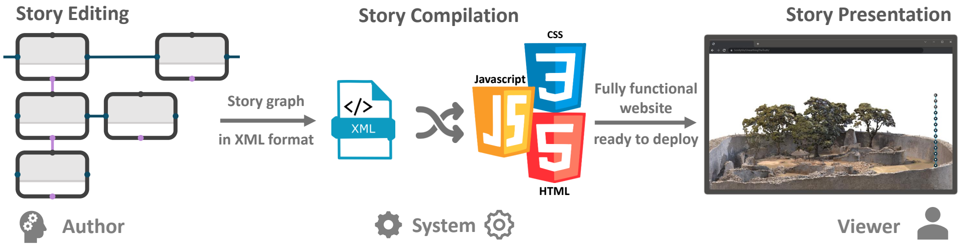

Figure D.1: The ScrollyVis system consists of three main parts: story editing, story compilation and story Presentation. The editing phase allows authors to create complex dynamic narratives in the ScrollyVis editor. The story is exported in an XML format and is compiled into a fully functional website within the ScrollyVis compilation phase. The website can subsequently be presented in any standard web browser while the ScrollyVis created code handles the dynamic narrative handling.

stories include textual information, in most cases combined with media such as images, videos and audio. In some cases, more complex visualizations are presented, such as the surface visualization of photogrammetry data in the story by the Economist [37]. The stories can be further categorized into partial or full scrollytelling websites. One example of a full scrollytelling website is the the Genexpression story [98] , while the Economist story [37] or the story created by Apple Research [148] employ it partially. Creating such stories is associated with considerable costs and frequently involves a team of programmers, web developers, authors, designers, directors and content consultants, as exemplified by the the Genexpression story [98]. One of our main goals is to reduce this cost and efforts in order to enable much smaller teams or even single individuals to create immersive and impactful stories about their work or other topics of interest.

Stories can be told in many ways, influencing how the story is perceived by the listeners, viewers or readers. A narrative specifies the order in which events are told as opposed to the order the events actually happened [44]. Narratives can be in chronological order, telling the story in the order it happened, but it is also possible to present events out of order in a nonlinear way [80]. In contrast to the story itself, which focuses on the content, narrative is the expression of a story [157]. Narratives can be used to maintain a sense of mystery by, e.g., withholding information in order to keep tension high and the audience engaged [80]. User engagement is a very important aspect in storytelling and to this end user-directed paths can be employed [158]. We differentiate in this way between passive and interactive stories. To prevent ambiguity by overloading the term "interactive" in a visualization context, we will refer to such stories as "dynamic" in our approach.

Authoring media-rich scrollytelling experiences mainly consists of two parts: defining the content to present and specifying the relationships between content items. This can be achieved in many different ways, as existing editors like Stornaway [151] and Twine [161] demonstrate. Based on the discussions in the literature as well as our analysis of scrollytelling content, we identify the following requirements for a system enabling authors to create such stories:

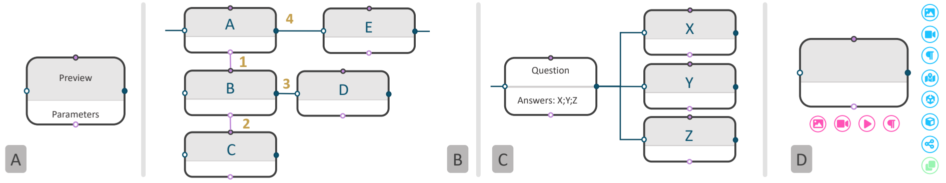

Figure D.2: A: presents the abstract unit story node which consists of a preview, a properties section and connection nodes on all four sides. In B story transitions are depicted, in this example the layering possibilities in ScrollyVis are shown. C presents a decision node where the viewer can later in the story presentation choose which path they would like to pursue and in D the context menu is depicted which allows for fast and accurate story creation.

1. The system should enable efficient definition of story content and content item relations.
2. To support multiple narrative structures, there needs to be support for story branching.
3. The system should ideally support all narrative structures proposed by Munday [107].
4. The system should enable authors to specify how content items representing different or same media types should be combined.
5. Complex content layering combinations should be possible in the system.
6. The system should be easy to use for users at various skill levels and allow for efficient story creation.
7. The created stories should be viewable in a standard web browser.
8. The story should support rapid incremental and reversible scrolling interaction.

D.4 ScrollyVis

To support the creation of media rich scrollytelling experiences, we introduce ScrollyVis, a web based story authoring system which follows a no-code paradigm to create media-rich and dynamic narratives. In the following, we present the individual concepts and components of the authoring system. The system has three main components responsible for story editing, compilation, and presentation (see Figure D.1). In the following, we outline the main concepts in these three components.

D.4.1 Story Editing

The first component of our system is the story editor. This part of the system is responsible for the specification of story content and the definition of relations between the content items.

Wohlfart and Hauser [170] employed the concepts of story nodes and story transitions in their visual storytelling approach. Story nodes are stages within a story where content is presented. They are connected by transitions and, similar to the work by Wohlfart et al. [170], these transitions ensure a gradual progression between story nodes and fulfill requirement 8 from Section D.3. Instead of clicking through a slide show, we use scrolling as the only interaction control of the story progress which follows rule 1 of Bostock [11]. We abstract story nodes to consist of three different parts as visualized in Figure D.2 A: a node preview, node parameters, and node connections. In our system a story node is used as a content item and fulfills requirement 1 as outlined in Section D.3. This design was motivated by storyboards which are widely used for the purpose of pre-visualizing a movie, animation, or interactive media sequence. Depending on the node type, the preview can be static or interactive. In the current version of ScrollyVis, nodes can consist of the following media types: text, image, video, audio, map, 3D volume visualization and 3D surface visualization. Nodes which allow for camera control such as volume or surface visualizations provide corresponding interaction facilities within the preview window. Static previews are used for images, videos, text, and audio files. The second part within a story node are the parameters. The node parameter set is dependent on the type of visualization. While images, for instance, have parameters specifying the position within the website and the size of the presentation, complex nodes like volume visualizations feature a larger set of parameters. These include the definition of which volume visualization method shall be used, e.g., ISO value, maximum intensity projection or direct volume rendering with a transfer function, and associated parameters. To support rapid story authoring, we initialize every story node with a pre-defined set of sensible default parameters which can be altered as needed. Story nodes feature connections ports, which are used to control the story flow. In our system, every node is associated with a pre-defined code segment which is instantiated with the given parameters. This instance is used when the story author places the node onto the canvas of the editor. This allows for easy extension as new node types simply have to be defined as a new template code in order to be usable in the editor. Story nodes can be connected from left to right and from top to bottom. Every node therefore features an output port on the right and on the bottom side and an input port on the left and top.

Layers: To create complex combinations of story nodes we allow authors to combine different nodes in a layered manner. Authors have to be able to specify that one story node is presented while another story node is still shown to support media layering. There are different ways to depict such a behavior in an editor. Inspired by layered tracks used in common video editing software, we propose a sub-path feature as visible in Figure D.2 B. This sub-path allows authors to link story-nodes not only going from left to right but also from bottom to top. Story nodes connected at the bottom of a previous node are added as another layer in the final story result. Story nodes attached on the right side of previous node replace the previous one. With this simple design, different media types can be combined easily. In the example in Figure D.2 B, the story elements are presented as follows. Item A is presented first, and item B is presented while A stays visible. C is presented while A and B stay visible. D is presented while A stays visible, and B and C disappear. Finally, E gets shown as a new content item. This fulfils requirements 4 and 5 discussed in Section D.3. We refer to a node together with its sub-path as a *story segment*. Story segments encapsulate parts of the story that belong together and are used as the primary navigational unit in the presentation of a story.

Dynamic Narratives: Scrollytelling on the web is characterized by incrementally revealing information based on scrolling interaction by the viewer. In scrollytelling, viewers traverse a webpage along one axis in a linear fashion. The interaction is constrained to a single degree of freedom, i.e., scrolling up or down. When introducing dynamic narratives this pattern is no longer sufficient. Dynamic narratives are a powerful story telling tool which boosts viewer engagement as story immersion increases [107]. To support dynamic narratives and therefore requirement 2, we introduce decision nodes and allow for multiple story endpoints in our node-link storyboard as depicted in Figure D.2 C. At decision nodes, the author determines a point in the story where the story branches out and the viewer is able to decide the next story segment. The author defines several story path options and defines the available branches. By introducing such decision nodes we also fulfill requirement 3 introduced in Section D.3.

Interface: Our prototype web-based editor was designed with efficiency and ease of use in mind, according to requirement 6 from Section D.3. We allow users to build a scrollytelling website by visual programming, i.e., dragging and dropping story nodes onto a storyboard editor or by using a context menu shown and establishing links. We introduce several usability features to allow for efficient story editing. First, every connection port on a story node features a context menu which shows the possible new node type connections on demand. The context menu is presented in Figure D.2 D. Via the bottom connection menu, an image viewer, video player, audio player or a text node can be created. When using the context menu on the main path, every node type is available. As soon as the user selects one of the possible node types, the new node is automatically placed on the canvas at a fixed offset to the current node and on the same horizontal or vertical position, depending on whether it is a sub-path or main path connection. Furthermore, editor canvas automatically scrolls to focus on the new node. This allows for fast and accurate story creation with minimal interaction, while the story graph layout is optimized by default.

In the context menu along the main path, we also provide a copy functionality as present in Figure D.2 D in green. This feature allows the story author to copy the current node with all its properties. The copy is automatically connected to the current node. This feature is mainly interesting for node types which feature complex interaction methods like camera control (e.g., volume or surface visualizations). By copying the current node with all its settings, animated camera transitions can be authored in a rapid manner. In the copied node, the author just has to move adjust the settings and in the final web page the camera will zoom, pan, and rotate to transition between the two views.

The story graph is serialized into an XML format. The XML document stores all nodes and their parameters as well as all node connections. The resulting files are human readable and can also be adjusted manually. Furthermore, this simple intermediate format also enables the future exploration of additional higher-level interfaces such as wizards which may serve as a starting point for customization. After the author finishes story editing, they can export the story as a web page ready to deploy it on a web server. We call this process story compilation and describe it in the following subsection.

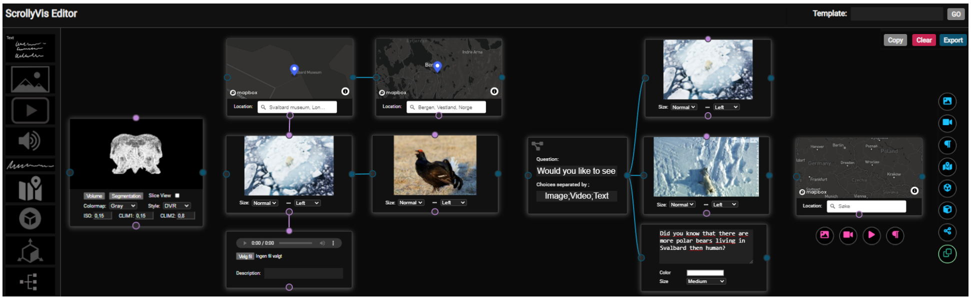

Colormap: Gray
Style: DVR
ISO: 0,15
CLIM1: 0,15
CLIM2: 0,8

Location: Svalbard museum, Lon...

Location: Bergen, Vestland, Norge

Size: Normal
Left

0:00 / 0:00
Velg fil Ingen fil valgt
Description:

Question: Would you like to see
Choices separated by ; Image;Video;Text

Did you know that there are more polar bears living in Svalbard then human?
Color
Size: Medium

Location: Søke

Figure D.3: Depiction of the ScrollyVis editor representing the same functionalities as presented in the abstract system depiction in Figure D.2. Stories can be created either by dragging nodes from the left onto the canvas or by using the context menu as presented on the right side of the Figure. All nodes where applicable feature a preview window.

D.4.2 Story Compilation

At any point during the editing phase, the user can trigger the compilation of the story into a ready-to-deploy website. The input of the story compiler is the story graph in XML format together with all media content, e.g., images and videos. Depending on whether decision nodes are present, a story can be presented as a linear sequence of story nodes or as a tree where the narrative structure branches out after every decision node. Every story node represents different types of content, e.g., text or an image, that has to be translated from story logic in XML to HTML, CSS and JavaScript code for the final presentation in every standard web browser which fulfills requirement 7 presented in Section D.3. The compilation phase is divided into two steps: first the story content nodes are created and then the node transitions are set up.

Story Contents: First, the story graph delivered from the ScrollyVis editor is traversed to define the story tree and to create the content for the resulting website. Each story node is represented as pre-defined HTML and CSS code, dependent on the node media type. Less complex media like text, images, videos, and audio can be converted directly to HTML code. More complex media, e.g., map views, volume visualization, and surface visualization, need content loading code in addition. On traversal of the story graph, a story tree is created. Except for the root, leaf, and decision nodes, every node has exactly one predecessor node and one successor node. In this step the sub-tree of every node connected via the sub-path port is flattened. The sub-path feature is especially important in defining the transition between story nodes. A decision node is the only node which has multiple successor nodes to create branching points in the story tree. In a second step, we traverse the generated story tree and take main- and sub-path information into account.

Story Transitions: In the second step of the story compilation, we focus on story node transitions. The behavior of every node in the story is defined by its state before, during, and after node traversal. Before traversal, the node the content is not visible. This state is defined in the JavaScript code by setting the opacity to 0. Node transition handling depends on the existence of a sub-path. If there is no sub-path, the node will be blended in and out within the scrolling extent of the current node. In our current version we a scroll extent of 1000 pixels per story node by default, but this value can be adjusted by the author. If there is a sub-path, the current node will be blended in and then faded out only when the next main path node is traversed. In addition to blending node opacity, further transitions such as camera movement and dynamic parameter changes are also handled within the node scrolling extent. After the node has been traversed, it indicates its successor node. The only exception to this is the decision node. At a decision node, successor selection is dependent on viewer interaction. As soon as the viewer decides upon a story branch and scrolls, the selected successor node is hooked to the decision node and the story viewer can carry on viewing the story. When scrolling back up the successor is unhooked again to enable different story path traversals. The compiled story includes node content in HTML and CSS as well as node content loading and transition code in JavaScript. In addition, it includes static code handling volume visualization, the Sketchfab integration for surface rendering, Mapbox code for map views, and static CSS styling code. The compiled story is put together as a complete website that is ready to deploy on a web server.

D.4.3 Story Presentation

After compilation, the story is ready for presentation. Stories created in the ScrollyVis editor are exported ready to be deployed on a webserver after compilation. Viewers can simply open the scrollytelling story in a web browser of their choice and interact with the story by scrolling. One important aspect of smooth story presentation are the transitions between different story nodes. In our case, the transition between the nodes is defined automatically, based on the current node content and content of the story segment displayed previously. These transitions allow for rapid, incremental, reversible scrolling and to provide consistent and instantaneous feedback following rule 2 and rule 3 from Bostock [11]. In addition, we introduced a rudimentary keyboard control option to progress the story to fulfill the last rule of Bostock [11] which could be extended on in a future version of ScrollyVis.

Transitions

Node transitions depend on the predecessor and successor node content and parameters. Furthermore, behavor differs according to whether they are part of main- or sub-path traversal. In general, as sub-paths are traversed, transitions are made by blending the opacity of the current and previous node content. Within the main path more elaborate transition methods may be used. Between media nodes of a different type, e.g., picture to video or volume to surface visualization, the transition works via opacity change. More elaborate transition methods are available for map, volume, slice, and surface nodes. Every transition is linked to viewer scroll interaction and works in both scroll directions without using triggered animations. One exception is the embedding of videos, where video playback is triggered by viewer scrolling interaction.

Map View: When first entering a map view from any other node type, the map will be introduced by an opacity change combined with a zooming transition from far away to close to the target location. Between two map views at different locations, the story viewer will fly from the initial location to the next target location. On this flight, the zoom level is adjusted to replicate a parabolic flight or jump [164].

Direct Volume Rendering and Slice Views: When entering a volume visualization node, the zoom level is adjusted to the zoom level set by the story author in the volume visualization node. When transitioning from one volume visualization node to another with the same input data, multiple parameters are linearly blended. Camera translation, rotation, and zoom level are updated based on user scrolling interaction. In addition, there are volume-specific parameters available for blending such as the lower and upper intensity value range limit and the ISO value. The lower and upper intensity value range limits are used to change the contrast for the volume rendering by limiting the applied color map to a reduced value range. The ISO value defines the value which is used in the ISO-Surface volume rendering method and defines which voxels are set to be within and which ones are outside of the surface of interest. If the volume is shown as a slice view, the slice index is also altered from the previous node slice index to the current one.

3D Surface Visualization: When entering a 3D surface visualization node, the opacity of the visualization is altered based on viewer scrolling interaction. Between two 3D surface visualization nodes representing the same data set, similarly to volume rendering, camera position, camera rotation, and camera zoom level are updated through scrolling. These transitions allow for natural camera movements and animations by simply defining start and end camera positions in two story nodes within the ScrollyVis editor.

Story Length

Story length is determined by the overall number of nodes included in the story. In order to ensure a stable story viewer experience, we do not alter the length of the story transition based on the content. The story nodes and transitions have a consistent scroll length. As transitions at times contain complex camera movement and transitions, they are equally important as image and video presentation. Story pacing is solely determined by the scrolling speed of the story viewer. The overall story length can be influenced by the overall website height. We provide a pre-defined website height where we allocate 3000 pixels for each story node, including in- and out-transitions. This value can be changed on demand by the story author. As we also allow for dynamic narratives, the visible scroll bar extent loses importance in the story presentation as shorter paths in a story result in the scroll bar indicator not reaching the end of the scroll bar. To remedy this, we introduce a concept that replaces the scroll bar in our story presentation, similar to how the Genexpressionen story handles scrolling indication [98].

Dynamic Narratives

To enable exploration of dynamic narratives on the web with reversible decisions, we introduce a real-time on demand story node linkage method. In the exported web page, the story is created on the fly. Every story node is linked on demand to the next one and the story transitions are defined in real-time in order to enable dynamic narratives. If a decision node is encountered by a viewer, they can select a path and the selected branch is traversed. If the viewer scrolls back previous traversed story nodes are "unhooked" allowing for rapid story transitions in both directions in real-time. This flexible approach allows for reversible dynamic narratives on the fly while allowing for smooth and uninterrupted scrollytelling viewer experiences.

Story Tree View

According to Kosara [88], it is important for scrollytelling viewers to get a sense of the overall story length and where they are currently positioned within the overall story. As visible in the example story in Figure D.5 on the right-hand side, we present a story overview. The viewer can see how many segments the story consists of, if there are decision steps, and which media types are present. In addition, they can see what story point they are currently viewing, how many steps have been traversed, and how many there are still to explore. The story tree view consists of nodes representing story segments depicting node type as an icon, where the root node is visualized at the top and subsequent nodes are placed underneath predecessor nodes. Furthermore, decision nodes show up as splits into story branches. The viewer can always track what part of the story they are currently exploring, and to which point they have to scroll back to explore another decision path. In Figure D.5, an example story and the accompanying story tree view are visible. While scrolling through the website, the outer ring of the story nodes the viewer has already visited will turn gray from the original blue color one by one. This keeps track of previous decisions and allows viewers to scroll back up to select other unexplored paths.

D.5 Implementation

In Figure D.4, the overall structure of exported websites is presented. In *ScrollyVis*, we build up upon a variety of different JavaScript libraries. *Drawflow* [147] delivers the base functionality of dragging and dropping nodes onto a canvas for further linking up with other nodes. It provides standard interactions for placing and linking nodes. Furthermore, we customize the existing XML export functionality for the serialization of our stories. In order to fulfill our requirements, we added several features and altered specific behaviors of the library. First, we added previews to all nodes in order to enable storytellers to preview the final result of the website while editing. We also implemented support for drop-downs and check-boxes within nodes for setting specific parameters and the ability to customize connection ports in order to realize our layering approach. Another major addition is the context menu that is depicted in Figure D.3 on the right. It can be customized for each node type and features only nodes which are compatible with the associated node port.

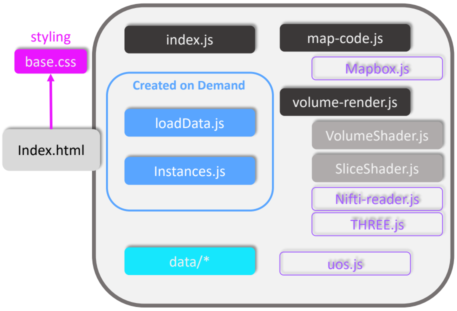

Figure D.4: Exported website structure: index.html is the starting point for the scrollytelling website. In blue, we present nodes where code is generated on the fly and in gray we present nodes which have been developed by us, but are story-independent. Violet nodes reflect imported external libraries.

The volumetric visualization is based on WebGL and we leverage Three.js [20] as a

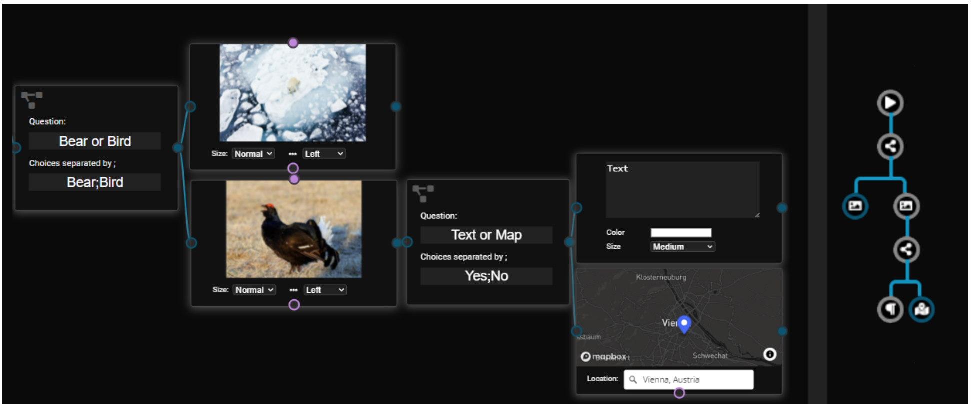

Question:
Bear or Bird
Choices separated by ;
Bear;Bird

Size: Normal ··· Left

Size: Normal ··· Left

Question:
Text or Map
Choices separated by ;
Yes;No

Text
Color
Size: Medium

Location: Vienna, Austria

Figure D.5: An example story is visible on the left with the associated story tree visualization on the right.

basis for both direct volume rendering and slice views. We use custom vertex and fragment shaders to enable direct volume rendering in Three.js, based on code provided by Valentin Demeusy [34]. To allow for volume rendering in Three.js [20], we added shaders for volume and slice rendering to the library using WebGL. For the volume shader, we currently support maximum intensity projection, isosurface visualization and direct volume rendering. When including map views, we use *Mapbox* and *OpenStreetMap*. We use JSZIP [86] for dynamic website packing making the website ready for downloading. For reading NIfTI files, we use NIFTI-READER-JS developed by Jack L. Lancaster and Michael J. Martinez [92–94]. We utilize uos.js to support scrollytelling in our approach, which is provided by Colin van Eenige [38]. To allow for efficient 3D surface visualization we integrate *Sketchfab* [143] by using the Sketchfab API. We plan to make our ScrollyVis editor freely available so that interested parties can generate their own scrollytelling web pages for scientific communication and outreach activities.

D.6 Case studies

We invited researchers from three different scientific fields to create stories about their work together with us. For further details, we refer to the additional materials for high resolution images of the networks created and video versions of the websites. Interactive additional materials, including the case studies as websites, are available at the following link: ScrollyVis use cases . For our scenarios, we invited an osteology expert, a meteorology visualization expert and a PhD student in anatomy education. Two of the experts have a doctoral degree in their respective fields and several years of experience and one is a PhD student. For our case studies, we created a guided dynamic narrative about their research and gathered their impression of our prototype, the presentation-, and interaction-style. In addition, we invited them to share their views on the potential of scrollytelling as a tool for scientific communication, education, and outreach to

---

https://ericmoerthuib.github.io/ScrollyVis/

a more general audience. As a first case study we present a recreation of a story by The Economist.

D.6.1 The Economist: Unearthing the Truth

In order to exemplify that our approach is capable of reproducing professionally developed scrollytelling content, we use the story mentioned in section D.3 by The Economist. It employs 3D models of the great enclosure and the hill complex, created by the Zamani Project of the University of Cape Town, South Africa. Creating such websites involves extensive web development skills and takes a substantial amount of time to create. To show the utility of our editor, one of our authors created a similar story with our editor in 15 minutes using models available on Sketchfab (see Figure D.6). The story is included in the additional materials and available at the following link: Our Story, the original story is available under the following link: The Economist [37].

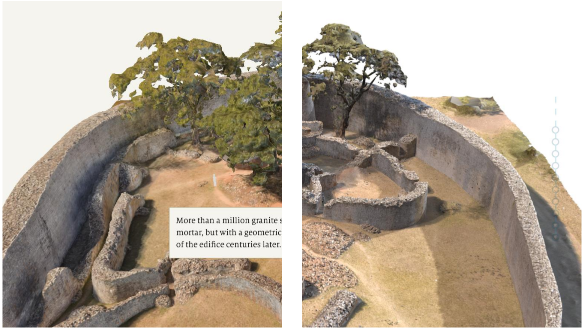
More than a million granite
mortar, but with a geometric
of the edifice centuries later.

Figure D.6: An archaeological story by The Economist including 3D models of the Great Enclosure and the Hill Complex by the Zamani Project of the University of Cape Town. On the left a screenshot of the original story and on the right side our version.

D.6.2 Scientific Outreach: Osteology Research

Together with the first expert, an Associate Professor at the University Museum of Bergen and the curator of the modern osteological collections, we have formulated a case study focused on outreach activities. Part of the story content is highlighted in Figure D.7. The volumetric data used in this story were acquired by a Computed Tomography (CT) scanner. The story primarily focuses on polar bears in Norway, inspired by a blog post on this topic explaining characteristics of their skulls in relation to their habitat. In the story, we include several decision nodes to support viewer engagement and educational goals. Furthermore, we import custom geo-spatial data, consisting of polar bear sightings in Norway around Svalbard. The goal of the story is to showcase the potential of using our tool as a scientific outreach tool and for creating interactive web-based museum exhibits. First, we introduced our ScrollyVis editor to the osteology expert and invited her to create a story completely on her own including various media types like a map view or a volume visualization. The expert has some experience in working with volume visualization and gave us the feedback that our interaction methods are as simple as the other tools she uses in her work. Furthermore, she created a volume animation consisting of several steps and told us that she has never created

---

https://ericmoerthuib.github.io/ScrollyVis/UnearthingTheTruth such an animation before. After this phase of the evaluation, we created a story together with the expert and encouraged her to speak openly about potential advantages and disadvantages of our approach.

She thinks that the story created with ScrollyVis has a high potential to excite and engage visitors of the University Museum. According to her, the questions presented in the story can engage the viewers. During our discussion, she added that it would be beneficial to make the story mobile phone-friendly such that museum visitors can re-experience or share their experience within the museum with friends and family. She thinks enhancing the exhibition with on-demand mobile- or touchscreen-based scrollytelling content has high potential to enrich the museum attending experience. In comparison to specifically designed museum exhibits that are developed together with visualization researchers, e.g., *Living Liquid* by Ma et al. [99] or *Sea of Genes* by Dasu et al. [33], creating a story with ScrollyVis is feasible without collaborating with external researchers or paying for professional services. She reflected that our editor is nicely designed and looks user-friendly, though she expects that an initial learning phase together with us might be helpful. All in all, she is excited about the potential of ScrollyVis for scientific outreach activities and would like to further explore the opportunities in her blog posts and potentially for an upcoming exhibition in the University Museum.

Figure D.7: Direct iso-surface volume rendering and slice volume rendering of scanned animal skulls in the osteology story co-created with an expert.

D.6.3 Meteorological Visualization

The second story is authored together with a visualization expert. She is specialized in environmental visualization and her latest publication features geospatial and meteorological data visualization. First, she created a story including map views and decision nodes, which is essential for communication of her work, on her own. Finally, we jointly created a story presenting one of her papers [35]. The paper presents *Hornero*, a visual analytics tool for the detec-

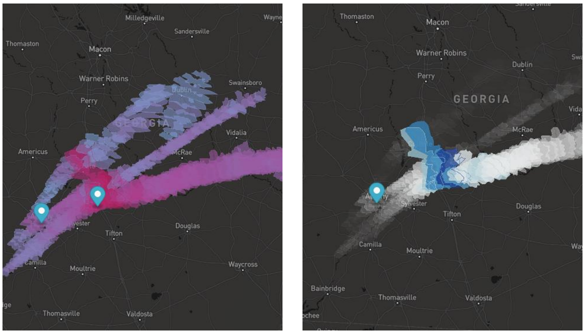

Figure D.8: A meteorological story including a storm strength and hail visualization using custom code provided by the paper authors.

tion and characterization of hazardous thunderstorms. Meteorological visualization is a very powerful tool to analyze the potential effects of hazardous weather phenomena. Presenting this information to a targeted audience or the general public is crucial to limit life-threatening risks. In addition to providing analytics tools for meteorologists and expert forecasters, it is equally important to present the results in an easy and comprehensible way to a more general audience. Selected story elements from the Hornero story are shown in Figure D.8.

In general, this expert liked our ScrollyVis editor for authoring the story about her work, but she thought it would be interesting to actually show the Hornero-based interactive visualizations in the story. To this end, we included custom code provided by the authors of Hornero to integrate their custom geo-spatial visualizations. She thinks that ScrollyVis is an exciting, accessible and easy way to create scrollytelling stories for the web. She has experience in creating websites and thinks that our approach greatly simplifies the process of producing high-quality and easily accessible content. In a recent project, she has used Figma to create mock ups and she thinks that one advantage of our approach is also that the full source code of the resulting website is available. In the future, she considers using our approach to explore the implications of interaction methods on communicating severe weather conditions to the general public. Similar to the first expert, she also requested improved support for mobile devices, as we currently mainly focus on desktop viewers. By simply altering the style sheet, we could already make the websites mobile friendly. Furthermore, she mentioned that our story tree view can be valuable for the story viewers but she would like to add keywords to the different nodes and would like them to be interactive and clickable. In conclusion, this expert thinks that ScrollyVis could be a part of many aspects of meteorological visualizations and can be used to communicate weather events in an exciting and engaging way.

D.6.4 Anatomy Education

The third expert we invited to use ScrollyVis is a PhD student who studies different means of anatomy education for medical students. She investigates different presentation media to analyze the effectiveness and learning outcomes. In their current setup, Sketchfab is used to show anatomical surface scans of bones which are made available within the online learning platforms the students are used to. The expert would like to use ScrollyVis to explore the effectiveness of guided dynamic narratives as teaching method compared to their existing content. She was invited to use ScrollyVis and was then asked to take part in an interview with us. The researcher was able to create stories on her

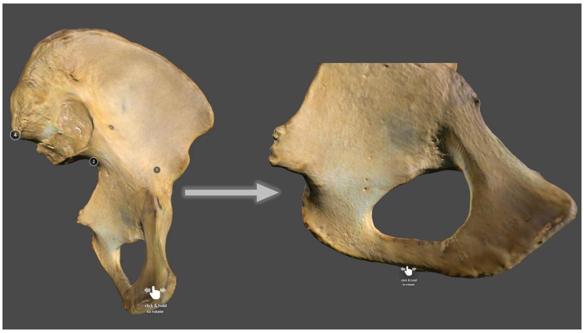

Figure D.9: ScrollyVis case exploring anatomy education potential based on specimen surface scans. The automatic camera transitions defined by starting and end point in the editor is particularly interesting in this case.

D

own after watching the introduction video included in the additional material, without needing any support from us. She sees great potential in the editor and thinks the integration of Sketchfab enables them to create stories which might increase the learning outcomes of the medical students. Her advisor, a professor teaching anatomy and practicing orthopedic surgeon, appreciates the support for including audio files to add spoken commentary to interesting spots of the bone surface scans. This is a feature that is not available in the standard interface of Sketchfab. Furthermore, the expert would wish for a separate quiz mode where the students are not able to change answers after making a decision in a branching node.

Table D.1: Response of the participants on a 5-point Liker scale, where 1: strongly disagree, 2: disagree, 3: neither agree nor disagree, 4: agree and 5: strongly agree. Statements marked with a star are rephrased to the positive form in this table with inverted scores for presentation purposes. In the rightmost column, average values are presented. The second to last row reveals the results of the SUS questionnaire. The last row shows the time it took the participants to create the second evaluation story.

	Statements	S1	S2	S3	S4	S5	S6	P1	P2	P3	P4	O1	O2	Avg.
G1	I would like to use the ScrollyVis editor for creating Scrollytelling websites	4	5	4	5	5	5	4	5	4	5	5	4	4,58
G2	Interacting with the editor is straightforward *	5	4	4	5	4	4	5	5	4	5	4	5	4,50
G3	I don't need any web development pre-knowledge to create a website with the editor *	5	5	4	5	5	5	5	5	5	5	5	5	4,92
G4	Creating a story based on the nodes and links in the editor is straightforward	4	5	4	5	3	5	4	5	4	4	4	4	4,25
G5	I don't need help using the editor to create a Scrollytelling website in addition to the provided tutorial	5	5	4	4	4	5	4	5	2	5	3	4	4,17
N1	The node preview (see image) helps me to find the right visualization settings for each file	5	5	4	5	4	4	4	5	3	5	5	5	4,50
N2	Adding a map view with a specified location was easy for me *	5	5	5	5	5	5	5	5	3	5	5	4	4,75
N3	Using ScrollyVis, I don't need pre-knowledge about shaders and WebGL to create a volume visualization on the web	4	5	3	5	5	4	5	5	3	5	5	5	4,50
N4	I can easily integrate a video in my story	5	5	5	5	5	5	5	5	5	5	5	5	5,00
N5	Combining different node types (e.g., map, text, image, . . .) is easy *	4	5	3	5	5	5	4	5	5	5	5	5	4,67
I1	The preview window in each node helps me to imagine the resulting website while editing it	3	4	4	5	4	3	4	4	5	4	5	4	4,08
I2	The clone interaction (see image) helps me to create an animation of volume data	5	5	5	5	4	4	4	5	5	4	5	4	4,58
I3	The main path and sub path feature helps me to create more complex stories	5	5	5	5	4	5	4	5	3	4	5	3	4,42
I4	The context menu (see image) helps me to create stories more efficiently *	5	5	5	5	5	5	5	5	5	5	5	5	5,00
I5	Linking the nodes to create the story I would like to tell is easy and self-explanatory	3	5	3	5	4	4	4	5	5	5	4	4	4,25
I6	Creating a story with the editor does not need pre-existing knowledge about designing storyboards *	5	5	4	5	4	5	2	5	5	5	5	5	4,58
I7	Adding questions for the viewer to decide which path the story shall go is easy and intuitive	4	4	4	5	4	4	4	5	2	5	5	3	4,08
I8	The interaction methods used to define the volume visualization are intuitive and easy to use	4	4	5	2	3	3	4	5	4	5	4	4	3,92
R1	The tree view on the final website helps me to know in which path of the story I am.	4	5	4	5	5	3	4	5	3	4	5	5	4,33
R2	The resulting webpage reflects the intention I had when designing the story in the editor *	5	5	5	5	4	5	4	5	4	5	5	5	4,75
R3	The viewer interaction (decisions) is nicely integrated on the website	5	5	4	5	5	3	4	5	4	5	4	3	4,33
SUS	System Usability Scale	90	97,5	82,5	100	87,5	82,5	80	100	90	90	92,5	95	90,63
T	Time to create the second story	04:36	05:39	03:36	03:50	02:56	02:46	04:42	09:30	04:18	05:17	07:20	08:21	05:14

D.7 Evaluation

To evaluate the usability of our web-based authoring tool, we invited twelve people with diverse professional backgrounds. E1-E6 are masters students in computer science, E7-9 are PhD candidates in the field of visualization, E10 is a PhD candidate in the field of machine learning, and E11 holds a PhD degree in cell biology and is a post doc in the field of cancer imaging since 2018. Finally, E12 holds a PhD degree in medical imaging physics. None of the participants were involved in the development of our tool.

In the beginning of the evaluation, we invited the participants to watch a ten minute long ScrollyVis introduction and a tutorial video that is included in the additional materials. After that, participants were invited two create two different stories with the editor. The first story had the goal to familiarize the study participants with the different node types and interactions the editor provides by asking them to construct a story about a black grouse. In this story, the participants could still ask for help if they got stuck or had questions regarding the editor. The requirements for this story were that it has to include the following node types: decision, image, video, text and volume visualization. If these requirements were fulfilled, the first story was finished. The second story authoring task was more specific and included map views, image views, a decision node and text views. The participants also had to define both main paths as well as sub-paths to fulfill the task requirements. For this part of the evaluation, we only monitored the work of the participants but did not help them to reach their goal. Furthermore, we measured the time it took the attendants to fulfill the requirements. We used a think-aloud protocol where we asked participants to vocalize their thoughts and share their experience while they created the stories. After the study, the participants were invited to answer a questionnaire with 21 questions regarding different aspects of ScrollyVis. In addition to our evaluation form, the participants filled out the system usability scale (SUS) designed by Brook et al. [17]. All statements are evaluated based on a 5-point Likert scale with some statements negatively formulated.

D.7.1 Evaluation Results

The results of the evaluation are shown in Table D.1. All questions which are marked with a star were negatively formulated and we present them here in their positive form with inverted responses for ease of interpretation. Overall, the study participants provided positive feedback about the editor. The lowest average value in our study is 3.92 out of 5 regarding the volume visualization methods and the two most positive responses with 5 out of 5 are about the integration of videos in a story and the usefulness of the context menu. The study participants were positive about the preview windows we integrated in the editor, but some wished for a more accurate depiction of how things will be aligned on the final website. The current functionality of placing text in horizontal and vertical alignment presets was sufficient for the participants, but some wished to see where the text box will be displayed in the node preview. One participant mentioned that he thinks some help was needed during the creation of the website, but overall, the score for that point is 4.17 out of 5. One evaluation participant suggested that some previous knowledge on working with storyboards is helpful. Another participant thought that the decision node usage could be improved by highlighting the options along the paths exiting the decision node. All participants generally agreed that they do not need prior web development knowledge (4.92/5) and most participants would like to use the editor in the future (4.58/5). We further received positive feedback about the clone feature and that the resulting website reflects the intention of the story creators (4.75/5).

System Usability Scale Scores Our SUS scores are presented in the last row of the evaluation result in Table D.1. The results range from 82,5 to 100. On average, our application reached a SUS score of 90.63. Bangor et al. [5] introduced different ways of interpreting SUS scores, including the acceptability range, a grade scale, and an adjective rating scale. ScrollyVis achieved the highest score possible in all three categories: the acceptance rate is Acceptable, the grade scale score is A, and the application received an adjective rating of Excellent.

Story Creation Time The user evaluation included a section where participants created a well-defined story on their own after familiarizing themselves with the editor. For this part of the evaluation, we report the average and individual story creation time. The average story creation time was 5 minutes and 14 seconds. The shortest time was 2 minutes and 46 seconds and the longest was 9 minutes and 30 seconds. In our results, the master students of computer science were the quickest, followed by the PhD students. The two participants with backgrounds in medical physics and cell biology took slightly longer to create their stories.

We conclude from this initial study that ScrollyVis was considered useful by students and researchers of various backgrounds. In general, ScrollyVis received positive feedback from all participants and the most of them would like to use the editor in the future for various tasks. For example, some stated they would like to create a review of their achievements in the last year or to present the latest research results. The participants reported that the editor features all necessary and relevant features to use it for creating scrollytelling websites. Still, the participants had feature requests to improve the user experience even more. However, none of the evaluation participants felt that features were currently missing in order to effectively work with the editor.

D.7.2 Expert Feedback

To gain further insights into the utility of our approach, we conducted an interview with an expert in scientific storytelling via interactive articles. He holds a PhD in Computer Science and Engineering and worked collaboratively with designers, developers, and scientists at Apple, Microsoft Research, and the NASA Jet Propulsion Lab. We presented the expert our editor and all capabilities it provides and afterwards we discussed the stories created in our case studies. In general, the expert found that current tools capable of creating stories comparable to our results are cumbersome and involve extensive knowledge of web development skills. One similar approach called *Idyll* [29] focuses mainly on parameter exploration and does not support as many media types out of the box as our editor does. According to the expert, the support of immersive media like 3D volumetric and surface data is unique and makes our editor stand out compared to other approaches. Furthermore, the non-linear path support makes the stories more engaging and interesting to explore. The expert thought that the stories we presented were of high quality and was impressed by the ease of use of our editor. Normally, developing such scrollytelling experiences involves different skills and may include teams of 6 to 12 people, whereas with our editor it is feasible for one person alone to create a similar story if suitable assets are already available. Furthermore, the expert thought that scrollytelling websites are an important aspect of scientific communication and he believes that our editor could make significant contributions to the developments in this field.

D.8 Discussion

Our approach can generate guided interactive narratives in a scrollytelling environment. In our prototype authoring tool, authors without programming experience can create such interactive narratives with simple drag, drop, and linking interactions. During our case studies with experts from different fields, we noted high engagement both in viewing and authoring stories. Our approach has potential for outreach to the general public and as well as for more specialized applications such as medical education. For example, scientific communication goals could be achieved by including ScrollyVis-authored stories on the web to advertise museum exhibitions or to communicate research results. Other potential use cases for ScrollyVis are the creation of additional materials for scientific papers and anatomy education based on scanned anatomical specimens. ScrollyVis could be a part of paper submission materials as well as 'science in plain English' presentations of scientific results. While scrollytelling in general might not be the best fit for all visualization goals, in particular for tasks that are of a more exploratory or analytical nature, we believe that in combination with guided dynamic narratives it can be an effective way to present information. As demonstrated by our use cases, all experts were excited about the potential ScrollyVis offers and have concrete plans for use of ScrollyVis in the future.

One potential avenue of improvement frequently mentioned by the experts is limited support for mobile devices in the current prototype implementation. One solution would be to use server-based rendering for some of the more computationally heavy visualization techniques. In addition, the style sheet for the exported websites is currently not designed for mobile devices. Furthermore, we would like to explore the possibilities of including a Latex and Python interpreter to our editor in able to allow researchers to present their mathematical formulas in the same way as in their papers. A Python interpreter would be beneficial as it would simplify the creation of information visualization on the fly in a form many researchers are used to. Currently, we support all dynamic narrative structures introduced by Munday [107], but some improvements could be made in the context of concentric narratives, where it may be convenient for viewers to automatically return to the initial decision node after having reached the end of a story segment.

D.9 Conclusion and Future Work

With ScrollyVis, we introduce an authoring approach for realizing guided dynamic narratives as scrollytelling websites. We designed a extensible web-based story authoring tool that exports results ready for deployment on a web server. Our approach utilized a hybrid node-link storyboard editor which allows storytellers to get a good understanding of the resulting story during the authoring phase. Furthermore, a story tree view is available during story viewing that shows the story extent at a glance and where in the story viewers currently are. Finally, our story nodes support a variety of different media types including images, videos, audio, interactive maps, direct volume rendering and surface visualizations. Our system was designed with extensibility in mind, allowing for the easy integration of additional content.

We present a quantitative user evaluation with twelve independent participants with various professional backgrounds. Overall, ScrollyVis got a positive response from the study participants. The System Usability Scale was on average at the best possible grade level. The participants would like to use the editor in the future and think all relevant features are included to effectively work with the editor. Furthermore, we present four case studies, three of them collaboratively authored with experts from three different scientific disciplines. Overall, all experts were highly engaged in both authoring and viewing the stories. They expressed interest in using our approach in future projects, ranging from blog posts to creating additional publication materials. Furthermore, we invited an expert in the creation of interactive online articles to qualitatively evaluate the ScrollyVis editor and the quality of the created stories. The expert thinks that our approach has great potential and fills a gap in the scrollytelling editor landscape. The quality of the stories matches those created by big news agencies, but the creation does not require a large team of web developers. In general, the expert is convinced that ScrollyVis is valuable and useful for scientific outreach.

With this work, we demonstrate the potential for our approach to create immersive guided dynamic scrollytelling web experiences without having to write any code. We are confident that our authoring tool is a basis for further research in guided dynamic narrative structures and their effectiveness for scientific communication and educational purposes. Viewer engagement is a top priority for storytellers and with ScrollyVis we empower authors to create dynamic stories and explore their effect on viewers. In the future, we would like to extend our authoring tool and provide default templates to further lower the barrier to getting started with ScrollyVis. These templates could provide general structures used in different story types and the authors can then bring in custom content. We also envision a collaborative editing process where multiple authors can work on the same story. Furthermore, we would like to enhance the tree view with interactive navigation actions to directly access various story elements and to add further node transition possibilities between nodes of different types. With our work, we aim to inspire people to share their stories as we are confident that everyone has a story worth telling.

Acknowledgments

This research was funded by the Trond Mohn Foundation (grant numbers ’811255’ and ’813558’).

Bibliography

[1] Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., and Zheng, X. Tensorflow: A system for large-scale machine learning. In *Proceedings of the USENIX Conference on Operating Systems Design and Implementation* (2016), OSDI'16, doi: 10.5555/3026877.3026899, p. 265–283. B.4.3

[2] Amant, F., Moerman, P., Neven, P., Timmerman, D., Van Limbergen, E., and Vergote, I. Endometrial cancer. *The Lancet 366*, 9484 (08 2005), 491–505, doi: 10.1016/S0140-6736(05)67063-8. B.1, C.1

[3] Angelelli, P., Oeltze, S., Haasz, J., Turkay, C., Hodneland, E., Lundervold, A., Lundervold, A. J., Preim, B., and Hauser, H. Interactive visual analysis of heterogeneous cohort-study data. *IEEE Computer Graphics and Applications 34*, 5 (2014), 70–82, doi: 10.1109/MCG.2014.40. 2, 2.3, B.3

[4] Bakr, S., Gevaert, O., Echegaray, S., Ayers, K., Zhou, M., Shafiq, M., Zheng, H., Benson, J. A., Zhang, W., Leung, A. N. C., Kadoch, M., Hoang, C. D., Shrager, J., Quon, A., Rubin, D. L., Plevritis, S. K., and Napel, S. A radiogenomic dataset of non-small cell lung cancer. *Scientific Data 5*, 1 (2018), 180202, doi: 10.1038/sdata.2018.202. B.1

[5] Bangor, A., Kortum, P., and Miller, J. Determining what individual SUS scores mean: Adding an adjective rating scale. *Journal of Usability Studies 4*, 3 (May 2009), 114–123. A.6, B.8.1, D.7.1

[6] Becker, R. A., and William, S. C. Brushing scatterplots. *Technometrics 29*, 2 (1987), 127–142, doi: 10.2307/1269768. 2, 2.1

[7] Berg, A., Fasmer, K. E., Mauland, K. K., Ytre-Hauge, S., Hoivik, E. A., Husby, J. A., Tangen, I. L., Trovik, J., Halle, M. K., Woie, K., Bjørge, L., Bjørnerud, A., Salvesen, H. B., Werner, H. M. J., Krakstad, C., and Haldorsen, I. S. Tissue and imaging biomarkers for hypoxia predict poor outcome in endometrial cancer. *Oncotarget 7*, 43 (2016), 69844–69856, doi: 10.18632/oncotarget.12004. A.1, B.1

[8] Bernard, J., Sessler, D., May, T., Schlomm, T., Pehrke, D., and Kohlhammer, J. A visual-interactive system for prostate cancer cohort anal-

ysis. *IEEE Computer Graphics and Applications 35*, 3 (2015), 44–55, doi: 10.1109/MCG.2015.49. 2, 2.4, 2.3, B.3

[9] Bhatla, N., Berek, J. S., Cuello Fredes, M., Denny, L. A., Grenman, S., Karunaratne, K., Kehoe, S. T., Konishi, I., Olawaiye, A. B., Prat, J., and Sankaranarayanan, R. Revised FIGO staging for carcinoma of the cervix uteri. *International Journal of Gynecology & Obstetrics 145*, 1 (2019), 129–135, doi: 10.1002/ijgo.12749. B.4.1

[10] Blastland, M., Freeman, A. L. J., van der Linden, S., Marteau, T. M., and Spiegelhalter, D. Five rules for evidence communication. *Nature 587*, 7834 (2020), 362–364, doi: 10.1038/d41586-020-03189-1. D.1

[11] Bostock, M. How To Scroll. `https://bost.ocks.org/mike/scroll/`, 2014. [Online; accessed 31-March-2021]. D.2, D.4.1, D.4.3

[12] Bostock, M., Ogievetsky, V., and Heer, J. D3 Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (dec 2011), 2301–2309, doi: 10.1109/TVCG.2011.185. B.6

[13] Bradski, G. The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000). B.6

[14] Brancati, G. E., Brekke, N., Bartsch, H., Sørhaug, O. J. E., Ousdal, O. T., Hammar, Å., Schuster, P. M., Oedegaard, K. J., Kessler, U., and Oltedal, L. Short and long-term effects of single and multiple sessions of electroconvulsive therapy on brain gray matter volumes. *Brain Stimulation 14*, 5 (Sept. 2021), 1330–1339, doi: 10.1016/j.brs.2021.08.018. 4.4

[15] Bray, F., Ferlay, J., Soerjomataram, I., Siegel, R. L., Torre, L. A., and Jemal, A. Global cancer statistics 2018: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: A Cancer Journal for Clinicians 68*, 6 (2018), 394–424, doi: 10.3322/caac.21492. B.1, C.1

[16] Brehmer, M., and Munzner, T. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (Dec. 2013), 2376–2385, doi: 10.1109/TVCG.2013.124. 2, 2.1, B.4.4

[17] Brooke, J. SUS-a quick and dirty usability scale. usability evaluation in industry. In *Usability Evaluation In Industry*, P. Jordan, B. Thomas, I. McClelland, and B. Weerdmeester, Eds. CRC Press, 1996, ch. 10, pp. 266–290. A.1, A.6, B.1, B.8, D.7

[18] Bruckner, S., Solteszova, V., Gröller, E., Hladuvka, J., Bühler, K., Yu, J., and Dickson, B. Braingazer - visual queries for neurobiology research. *IEEE transactions on visualization and computer graphics 15* (11 2009), 1497–504, doi: 10.1109/TVCG.2009.121. A.2

[19] Buja, A., McDonald, J., Michalak, J., and Stuetzle, W. Interactive data visualization using focusing and linking. In *Proceeding Visualization '91* (1991), doi: 10.1109/VISUAL.1991.175794, pp. 156–163. 2, 2.1

[20] CABELLO, R. Three.js. `https://github.com/mrdoob/three.js`, 2021. [Online; accessed 31-March-2021]. D.5

[21] CAMPAGNOLA, L. PyQtGraph. `http://www.pyqtgraph.org/`, 2020. [Online; accessed 05-October-2020; Version: 0.11.0]. C.5

[22] CARD, S., MACKINLAY, J., AND SHNEIDERMAN, B. *Readings in Information Visualization: Using Vision To Think*. Morgan Kaufmann, 01 1999. 2

[23] CENEDA, D., GSCHWANDTNER, T., MAY, T., MIKSCH, S., SCHULZ, H., STREIT, M., AND TOMINSKI, C. Characterizing guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2017), 111–120, doi: 10.1109/TVCG.2016.2598468. B.3

[24] CENEDA, D., GSCHWANDTNER, T., AND MIKSCH, S. A review of guidance approaches in visual data analysis: A multifocal perspective. *Computer Graphics Forum 38*, 3 (2019), 861–879, doi: 10.1111/cgf.13730. B.3

[25] ÇIÇEK, Ö., ABDULKADIR, A., LIENKAMP, S. S., BROX, T., AND RONNEBERGER, O. 3d u-net: Learning dense volumetric segmentation from sparse annotation. *CoRR abs/1606.06650* (2016). B.4.3

[26] CLARK, K., VENDT, B., SMITH, K., FREYMANN, J., KIRBY, J., KOPPEL, P., MOORE, S., PHILLIPS, S., MAFFITT, D., PRINGLE, M., TARBOX, L., AND PRIOR, F. The cancer imaging archive (tcia): Maintaining and operating a public information repository. *Journal of Digital Imaging 26*, 6 (Dec 2013), 1045–1057, doi: 10.1007/s10278-013-9622-7. A.5, A.6

[27] CLEVELAND, W. S., AND SCHMIEG, G. M. The elements of graphing data. *American Journal of Physics 55*, 8 (1987), 767–767, doi: 10.1119/1.15021. B.5.1, B.5.1

[28] COMPUTING, R. PyQt5. `https://www.riverbankcomputing.com/software/pyqt/`, 2020. [Online; accessed 05-October-2020; Version: 5.15.0]. C.5

[29] CONLEN, M., AND HEER, J. Idyll. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Oct. 2018), ACM, doi: 10.1145/3242587.3242600. 2, 2.4, D.2, D.7.2

[30] CONLEN, M., VO, M., TAN, A., AND HEER, J. Idyll studio: A structured editor for authoring interactive data-driven articles. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Oct. 2021), ACM, doi: 10.1145/3472749.3474731. 2, 2.4, 2.9, D.2

[31] DAHLSTROM, M. F. Using narratives and storytelling to communicate science with nonexpert audiences. *Proceedings of the National Academy of Sciences 111*, 4 (2014), 13614–13620, doi: 10.1073/pnas.1320645111. 2, 2.4, D.1

[32] DAHLSTROM, M. F., AND SCHEUFELE, D. A. (escaping) the paradox of scientific storytelling. *PLOS Biology 16*, 10 (2018), e2006720, doi: 10.1371/journal.pbio.2006720. D.1

[33] Dasu, K., Ma, K.-L., Ma, J., and Frazier, J. Sea of genes: A reflection on visualising metagenomic data for museums. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (2021), 935–945, doi: 10.1109/TVCG.2020.3030412. D.6.2

[34] Demeusy, V. Volume Rendering in Three.js. `https://github.com/stity`, 2021. [Online; accessed 31-March-2021]. D.5

[35] Diehl, A., Pelorosso, R., Ruiz, J., Pajarola, R., Gröller, M. E., and Bruckner, S. Hornero: Thunderstorms characterization using visual analytics. *Computer Graphics Forum 40*, 3 (2021), 299–310, doi: 10.1111/cgf.14308. D.6.3

[36] Eckelt, K., Adelberger, P., Zichner, T., Wernitznig, A., and Streit, M. Tourdino: A support view for confirming patterns in tabular data. *EuroVis Workshop on Visual Analytics (EuroVA '19)* (2019), doi: 10.2312/eurova.20191117. B.3

[37] Economist, T. Unearthing The Truth: A Zimbabwean archaeologist reimagines the story of a momentous African civilisation. `https://www.economist.com/interactive/christmas-specials/2021/12/18/great-zimbabwe-archaeology`, 2021. [Online; accessed 24-March-2022]. D.3, D.6.1

[38] Eenige, C. v. Update On Scroll (uos). `https://github.com/vaneenige/uos`, 2019. [Online; accessed 31-March-2021]. D.5

[39] Ellis, G., and Dix, A. An explorative analysis of user evaluation studies in information visualisation. In *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization* (New York, NY, USA, 2006), BELIV '06, Association for Computing Machinery, doi: 10.1145/1168149.1168152, p. 1–7. 2, 2.5

[40] Fasmer, K. E., Bjørnerud, A., Ytre-Hauge, S., Grüner, R., et al. Preoperative quantitative dynamic contrast-enhanced MRI and diffusion-weighted imaging predict aggressive disease in endometrial cancer. *Acta Radiologica 59*, 8 (2018), 1010–1017, doi: 10.1177/0284185117740932. A.1, B.1, B.2, B.4.3, B.4.3, B.4.4, C.3

[41] Fasmer, K. E., Hodneland, E., Dybvik, J. A., Wagner-Larsen, K., Trovik, J., Salvesen, , Krakstad, C., and Haldorsen, I. H. Whole-volume tumor MRI radiomics for prognostic modeling in endometrial cancer. *Journal of Magnetic Resonance Imaging 53*, 3 (2021), 928–937, doi: 10.1002/jmri.27444. 3.4, 4.3, 5, C.1

[42] Fields, E. Build narratives in your dashboards with Story Points in Tableau 8.2 . `https://www.tableau.com/about/blog/2014/5/82-preview-tell-story-your-data-story-points-30761`, 2014. [Online; accessed 24-March-2022]. 2, 2.4, D.2

[43] FRIENDLY, M. A.-M. Guerry's Moral Statistics of France: Challenges for multivariable spatial analysis. *Statistical Science 22*, 3 (2007), 368–399, doi: 10.1214/07-STS241. A.2

[44] GENETTE, G., LEWIN, J. E., AND CULLER, J. Narrative discourse : an essay in method. *Comparative Literature 32* (1980), 413, doi: 10.2307/2906507. D.3

[45] GILLIES, R. J., KINAHAN, P. E., AND HRICAK, H. Radiomics: Images are more than pictures, they are data. *Radiology 278*, 2 (2016), 563–577, doi: 10.1148/radiol.2015151169. B.1, B.2, B.6, B.7.1

[46] GILLMANN, C., SMIT, N. N., GROLLER, E., PREIM, B., VILANOVA, A., AND WISCHGOLL, T. Ten open challenges in medical visualization. *IEEE Computer Graphics and Applications 41*, 05 (sep 2021), 7–15, doi: 10.1109/MCG.2021.3094858. 1, 2, 2.1, 2.2

[47] GLASSER, S., PREIM, U., TÖNNIES, K., AND PREIM, B. A visual analytics approach to diagnosis of breast dce-mri data. *Computers & Graphics 34*, 5 (2010), 602–611, doi: 10.1016/j.cag.2010.05.016. 2, 2.2, 2.3

[48] GLASSER, S., SAALFELD, P., BERG, P., MERTEN, N., AND PREIM, B. How to Evaluate Medical Visualizations on the Example of 3D Aneurysm Surfaces. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2016), S. Bruckner, B. Preim, A. Vilanova, H. Hauser, A. Hennemuth, and A. Lundervold, Eds., The Eurographics Association, doi: 10.2312/vcbm.20161283. 2, 2.5

[49] GLEICHER, M., ALBERS, D., WALKER, R., JUSUFI, I., HANSEN, C. D., AND ROBERTS, J. C. Visual comparison for information visualization. *Information Visualization 10*, 4 (2011), 289–309, doi: 10.1177/1473871611416549. A.2

[50] GODULLA, A., AND WOLF, C. *Digitale Langformen im Journalismus und Corporate Publishing*. Springer Fachmedien Wiesbaden, 2017. D.2

[51] GORDON, Y., PARTOVI, S., MÜLLER-ESCHNER, M., AMARTEIFIO, E., BÄUERLE, T., WEBER, M.-A., KAUCZOR, H.-U., AND RENGIER, F. Dynamic contrast-enhanced magnetic resonance imaging: fundamentals and application to the evaluation of the peripheral perfusion. *Cardiovascular Diagnosis and Therapy 4*, 2 (2014), doi: j.issn.2223-3652.2014.03.01. B.4.2

[52] GREEN, S. B., AND SALKIND, N. J. *Using SPSS for Windows and Macintosh: Analyzing and Understanding Data*. Prentice Hall Press, 2010. B.3

[53] GULLI, A., AND PAL, S. *Deep learning with Keras*. Packt Publishing Ltd, 2017. B.4.3

[54] HAEHN, D., KAYNIG, V., TOMPKIN, J., LICHTMAN, J. W., AND PFISTER, H. Guided Proofreading of Automatic Segmentations for Connectomics. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2018), 9319–9328, doi: 10.1109/CVPR.2018.00971. B.3

[55] HALDORSEN, I. S., GRÜNER, R., HUSBY, J. A., MAGNUSSEN, I. J., WERNER, H. M. J., SALVESEN, Ø. O., BJØRGE, L., STEFANSSON, I., AKSLEN, L. A., TROVIK, J., TAXT, T., AND SALVESEN, H. B. Dynamic contrast-enhanced MRI in endometrial carcinoma identifies patients at increased risk of recurrence. *European Radiology 23*, 10 (Oct 2013), 2916–2925, doi: 10.1007/s00330-013-2901-3. A.1

[56] HALDORSEN, I. S., AND SALVESEN, H. B. What is the best preoperative imaging for endometrial cancer? *Current Oncology Reports 18*, 4 (Feb 2016), 25, doi: 10.1007/s11912-016-0506-0. B.1

[57] HALDORSEN, I. S., STEFANSSON, I., GRÜNER, R., HUSBY, J., MAGNUSSEN, I. J., ET AL. Increased microvascular proliferation is negatively correlated to tumour blood flow and is associated with unfavourable outcome in endometrial carcinomas. *British Journal of Cancer 110*, 1 (2014), 107–114, doi: 10.1038/bjc.2013.694. A.1, B.1, B.2, B.4.3, C.3

[58] HARROWER, M., AND BREWER, C. A. Colorbrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal 40*, 1 (2003), 27–37, doi: 10.1179/000870403235002042. C.4

[59] HASTREITER, P., AND ERTL, T. Integrated registration and visualization of medical image data. *Proceedings of Computer Graphics International* (1998), 78–85, doi: 10.1109/CGI.1998.694253. B.3

[60] HODNELAND, E., DYBVIK, J. A., WAGNER-LARSEN, K. S., ŠOLTÉSZOVÁ, V., MUNTHE-KAAS, A. Z., FASMER, K. E., KRAKSTAD, C., ET AL. Automated segmentation of endometrial cancer on MR images using deep learning. *Scientific Reports 11*, 1 (Jan. 2021), doi: 10.1038/s41598-020-80068-9. 5, C.3

[61] HOHMAN, F., CONLEN, M., HEER, J., AND CHAU, D. H. P. Communicating with interactive articles. *Distill* (2020), doi: 10.23915/distill.00028. https://distill.pub/2020/communicating-with-interactive-articles. 2, 2.4, 2.7, D.1, D.2

[62] HOIVIK, E. A., HODNELAND, E., DYBVIK, J. A., WAGNER-LARSEN, K. S., FASMER, K. E., BERG, H. F., HALLE, M. K., HALDORSEN, I. S., AND KRAKSTAD, C. A radiogenomics application for prognostic profiling of endometrial cancer. *Communications Biology 4*, 1 (Dec. 2021), doi: 10.1038/s42003-021-02894-5. 4.3, 5, C.1

[63] HUNTER, J. D. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering 9*, 3 (2007), 90–95, doi: 10.1109/MCSE.2007.55. C.5

[64] INSELBERG, A., AND DIMSDALE, B. Parallel coordinates: A tool for visualizing multi-dimensional geometry. In *Proceedings of IEEE Visualization* (Oct 1990), doi: 10.1109/VISUAL.1990.146402, p. 361–378. B.5.2

[65] ISENBERG, T., ISENBERG, P., CHEN, J., SEDLMAIR, M., AND MÖLLER, T. A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2818–2827, doi: 10.1109/TVCG.2013.126. 2, 2.5

[66] JÄCKLE, D., FUCHS, J., AND KEIM, D. A. Star glyph insets for overview preservation of multivariate data. *IS and T International Symposium on Electronic Imaging Science and Technology* (2016), 1–9, doi: 10.2352/issn.2470-1173.2016.1.vda-506. 2, 2.2, A.2

[67] JENKINSON, M., BANNISTER, P., BRADY, M., AND SMITH, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage 17*, 2 (2002), 825–841, doi: 10.1006/nimg.2002.1132. B.3, B.4.3

[68] JENKINSON, M., AND SMITH, S. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis 5*, 2 (2001), 143–156, doi: 10.1016/S1361-8415(01)00036-6. B.4.3

[69] JONES, E., OLIPHANT, T., PETERSON, P., ET AL. SciPy: Open source scientific tools for python. `https://www.scipy.org`, 2001. [Online; accessed 07-April-2022]. C.5

[70] JOUBERT, M., DAVIS, L., AND METCALFE, J. Storytelling: the soul of science communication. *Journal of Science Communication 18*, 05 (Oct. 2019), E, doi: 10.22323/2.18050501. 2, 2.4, D.3

[71] JUAN-ALBARRACÍN, J., FUSTER-GARCIA, E., GARCÍA-FERRANDO, G. A., AND GARCÍA-GÓMEZ, J. M. Oncohabitats: A system for glioblastoma heterogeneity assessment through MRI. *International Journal of Medical Informatics 128* (2019), 53–61, doi: 10.1016/j.ijmedinf.2019.05.002. C.2

[72] JUAN-ALBARRACÍN, J., FUSTER-GARCIA, E., PÉREZ-GIRBÉS, A., APARICI-ROBLES, F., ET AL. Glioblastoma: vascular habitats detected at preoperative dynamic susceptibility-weighted contrast-enhanced perfusion MR imaging predict survival. *Radiology 287*, 3 (2018), 944–954, doi: 10.1148/radiol.2017170845. C.2

[73] JÖNSSON, D., BERGSTRÖM, A., FORSELL, C., SIMON, R., ENGSTRÖM, M., WALTER, S., YNNERMAN, A., AND HOTZ, I. Visualneuro: A hypothesis formation and reasoning application for multi-variate brain cohort study data. *Computer Graphics Forum 39*, 6 (2020), 392–407, doi: 10.1111/cgf.14045. 2, 2.5, 2.3

[74] 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 multivariate clinical data. In *Proceedings of Eurographics VCBM* (2019), B. Kozlíková, L. Linsen, P.-P. Vázquez, K. Lawonn, and R. G. Raidou, Eds., The Eurographics Association, doi: 10.2312/vcbm.20191232. A.2, B.3

[75] KARIMOV, A., MISTELBAUER, G., AUZINGER, T., AND BRUCKNER, S. Guided Volume Editing based on Histogram Dissimilarity. *Computer Graphics Forum 34*, 3 (2015), 91–100, doi: 10.1111/cgf.12621. B.3

[76] KARPATHY, A. tSNEJS. `https://github.com/karpathy/tsnejs`, 2016. B.6

[77] KAUFMAN, L., AND ROUSSEEUW, P. J., Eds. *Finding Groups in Data*. John Wiley & Sons, Inc., Mar. 1990. C.4

[78] KEHRER, J., AND HAUSER, H. Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Transactions on Visualization and Computer Graphics 19*, 3 (March 2013), 495–513, doi: 10.1109/TVCG.2012.110. 2, 2.2

[79] KEIM, D. A., MANSMANN, F., SCHNEIDEWIND, J., AND ZIEGLER, H. Challenges in visual data analysis. In *Tenth International Conference on Information Visualisation (IV'06)* (2006), IEEE, doi: 10.1109/IV.2006.31, pp. 9–16. 2.1

[80] KIM, N. W., BACH, B., IM, H., SCHRIBER, S., GROSS, M., AND PFISTER, H. Visualizing nonlinear narratives with story curves. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2018), 595–604, doi: 10.1109/TVCG.2017.2744118. D.1, D.3

[81] KINGSLEY, P. B., AND MONAHAN, W. G. Selection of the optimum b-factor for diffusion-weighted magnetic resonance imaging assessment of ischemic stroke. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine 51*, 5 (2004), 996–1001, doi: 10.1002/mrm.20059. B.4.2

[82] KLEIN, S., STARING, M., MURPHY, K., VIERGEVER, M. A., AND PLUIM, J. P. W. Elastix: A toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging 29*, 1 (Jan 2010), 196–205, doi: 10.1109/TMI.2009.2035616. A.4.1

[83] KLEMM, P., OELTZE-JAFRA, S., LAWONN, K., HEGENSCHEID, K., VÖLZKE, H., AND PREIM, B. Interactive visual analysis of image-centric cohort study data. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 1673–1682, doi: 10.1109/TVCG.2014.2346591. 2, 2.3, B.3

[84] KLIPPEL, A., HARDISTY, F., LI, R., AND WEAVER, C. Colour-enhanced star plot glyphs: Can salient shape characteristics be overcome? *Cartographica 44*, 3 (2009), 217–232, doi: 10.3138/carto.44.3.217. 2, 2.2

[85] KLIPPEL, A., HARDISTY, F., AND WEAVER, C. Star plots: How shape characteristics influence classification tasks. *Cartography and Geographic Information Science 36*, 2 (2009), 149–163, doi: 10.1559/152304009788188808. 2, 2.2, A.4.2

[86] KNIGHTLEY, S. jszip. `https://github.com/Stuk/jszip`, 2021. [Online; accessed 31-March-2021]. D.5

[87] Koh, D., and Padhani, A. Diffusion-weighted mri: a new functional clinical technique for tumour imaging. *The British Journal of Radiology 79*, 944 (2006), 633–635, doi: 10.1259/bjr/29739265. B.4.2

[88] Kosara, R. The scrollytelling scourge. `https://eagereyes.org/blog/2016/the-scrollytelling-scourge/`, 2016. [Online; accessed 31-March-2021]. 3.3.2, D.2, D.4.3

[89] Kosara, R., and Mackinlay, J. Storytelling: The next step for visualization. *Computer 46*, 5 (2013), 44–50, doi: 10.1109/MC.2013.36. 2, 2.4, D.1, D.2

[90] Kouril, D., Strnad, O., Mindek, P., Halladjian, S., Isenberg, T., Groeller, E., and Viola, I. Molecumentary: Adaptable narrated documentaries using molecular visualization. *IEEE Transactions on Visualization and Computer Graphics* (2021), 1–1, doi: 10.1109/TVCG.2021.3130670. D.2

[91] Krause, J., Perer, A., and Stavropoulos, H. Supporting iterative cohort construction with visual temporal queries. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2016), 91–100, doi: 10.1109/TVCG.2015.2467622. B.3

[92] Lancaster, J. L., Cykowski, M. D., McKay, D. R., Kochunov, P. V., Fox, P. T., Rogers, W., Toga, A. W., Zilles, K., Amunts, K., and Mazziotta, J. Anatomical global spatial normalization. *Neuroinformatics 8*, 3 (2010), 171–182, doi: 10.1007/s12021-010-9074-x. D.5

[93] Lancaster, J. L., Laird, A. R., Eickhoff, S. B., Martinez, M. J., Mickle Fox, P., and Fox, P. T. Automated regional behavioral analysis for human brain images. *Frontiers in Neuroinformatics 6*, JULY 2012 (2012), doi: 10.3389/fninf.2012.00023.

[94] Lancaster, J. L., McKay, D. R., Cykowski, M. D., Martinez, M. J., Tan, X., Valaparla, S., Zhang, Y., and Fox, P. T. Automated analysis of fundamental features of brain structures. *Neuroinformatics 9*, 4 (2011), 371–380, doi: 10.1007/s12021-011-9108-z. D.5

[95] Lawonn, K., Smit, N., Bühler, K., and Preim, B. A survey on multimodal medical data visualization. *Computer Graphics Forum 37*, 1 (2018), 413–438, doi: 10.1111/cgf.13306. 2, 2.1, 2.2, A.2, C.1

[96] Lidal, E. M., Hauser, H., and Viola, I. Geological storytelling: Graphically exploring and communicating geological sketches. In *Proceedings of the International Symposium on Sketch-Based Interfaces and Modeling* (2012), SBIM '12, Eurographics Association, p. 11–20. D.3

[97] Ltd., F. M. Texrad (texture + radiology. https://fbkmed.com/texrad-landing-2/, 2020. [Online; accessed 24-March-2022]. B.4.3

[98] LÖLL, C. Mission Jurassic: Searching for dinosaur bones. `https://www.substanzmagazin.de/krebsgentest`, 2015. [Online; accessed 24-March-2022]. D.3, D.4.3

[99] MA, J., LIAO, I., MA, K.-L., AND FRAZIER, J. Living liquid: Design and evaluation of an exploratory visualization tool for museum visitors. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2799–2808, doi: 10.1109/TVCG.2012.244. 2.4, 2.8, D.6.2

[100] MA, K., LIAO, I., FRAZIER, J., HAUSER, H., AND KOSTIS, H. Scientific storytelling using visualization. *IEEE Computer Graphics and Applications 32*, 1 (2012), 12–19, doi: 10.1109/MCG.2012.24. 2, 2.6, 2.4, D.2

[101] MALIK, M. M., HEINZL, C., AND GRÖLLER, M. E. Comparative visualization for parameter studies of dataset series. *IEEE Transactions on Visualization and Computer Graphics 16*, 5 (Sep. 2010), 829–840, doi: 10.1109/TVCG.2010.20. A.2

[102] MEUSCHKE, M., GARRISON, L., SMIT, N. N., BRUCKNER, S., LAWONN, K., AND PREIM, B. Towards narrative medical visualization. *CoRR abs/2108.05462* (2021), doi: 10.48550/arXiv.2108.05462. 4.4

[103] MIKSCH, S., AND AIGNER, W. A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data. *Computers & Graphics 38* (2014), 286–290, doi: 10.1016/j.cag.2013.11.002. 2, 2.1

[104] MIRZARGAR, M., WHITAKER, R. T., AND KIRBY, R. M. Curve boxplot: Generalization of boxplot for ensembles of curves. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2654–2663, doi: 10.1109/TVCG.2014.2346455. C.4

[105] MLEJNEK, M., ERMES, P., VILANOVA, A., VAN DER RIJT, R., VAN DEN BOSCH, H., GERRITSEN, F., AND GRÖLLER, M. E. Profile flags: a novel metaphor for probing of t2 maps. In *Proceedings of IEEE Visualization 2005* (Oct. 2005), H. R. C. T. Silva, E. Gröller, Ed., IEEE CS, pp. 599–606. A.2

[106] MÜLLNER, D. fastcluster: Fast hierarchical, agglomerative clustering routines for r and python. *Journal of Statistical Software 53*, 9 (2013), 1–18, doi: 10.18637/jss.v053.i09. C.5

[107] MUNDAY, R. A guide to interactive documentary: Structures, tools & narrative. `https://directorsnotes.com/2016/08/08/interactive-documentary-guide/`, 2016. [Online; accessed 31-March-2021]. 3, D.4.1, D.8

[108] MUNZNER, T. A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (Nov. 2009), 921–928, doi: 10.1109/TVCG.2009.111. 2, 2.5, A.3

[109] MUNZNER, T. *Visualization Analysis and Design*. AK Peters Visualization Series. CRC Press, 2014. 2, 2.1

[110] MURTAGH, F., AND LEGENDRE, P. Ward's hierarchical agglomerative clustering method: Which algorithms implement ward's criterion? *Journal of Classification 31*, 3 (Oct. 2014), 274–295, doi: 10.1007/s00357-014-9161-z. C.4

[111] 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. *Computer Graphics Forum 39*, 7 (2020), 611–622, doi: 10.1111/cgf.14172. C.1

[112] NEURATH, O. *International Picture Language: The First Rules of Isotype*. London, K. Paul, Trench, Trubner & co., ltd., 1936. B.5.2

[113] NG, H., ONG, S., FOONG, K., GOH, P., AND NOWINSKI, W. Medical image segmentation using k-means clustering and improved watershed algorithm. In *2006 IEEE Southwest Symposium on Image Analysis and Interpretation* (2006), IEEE, doi: 10.1109/SSIAI.2006.1633722, pp. 61–65. C.2

[114] OLIPHANT, T. E. *A guide to NumPy*, vol. 1. Trelgol Publishing USA, 2006. B.6

[115] OPACH, T., POPELKA, S., DOLEZALOVA, J., AND RØD, J. K. Star and polyline glyphs in a grid plot and on a map display: which perform better? *Cartography and Geographic Information Science 45*, 5 (2018), 400–419, doi: 10.1080/15230406.2017.1364169. 2, 2.2, A.2, A.4.2

[116] OUSDAL, O. T., BRANCATI, G. E., KESSLER, U., ERCHINGER, V., DALE, A. M., ABBOTT, C., AND OLTEDAL, L. The neurobiological effects of electroconvulsive therapy studied through magnetic resonance: What have we learned, and where do we go? *Biological Psychiatry 91*, 6 (Mar. 2022), 540–549, doi: 10.1016/j.biopsych.2021.05.023. 4.4

[117] PARK, D., DRUCKER, S. M., FERNANDEZ, R., AND ELMQVIST, N. <sc>atom</sc>: A grammar for unit visualizations. *IEEE Transactions on Visualization and Computer Graphics 24*, 12 (2018), 3032–3043, doi: 10.1109/TVCG.2017.2785807. B.5.2, B.6

[118] PENG, W., WARD, M. O., AND RUNDENSTEINER, E. A. Clutter reduction in multi-dimensional data visualization using dimension reordering. *Proceedings - IEEE Symposium on Information Visualization, INFO VIS* (2004), 89–96, doi: 10.1109/INFVIS.2004.15. A.4.2

[119] PETTERSEN, M. What Is Scrollytelling? (+Its Impact on Digital Content): An Interview. `https://learn.g2.com/scrollytelling`, 2019. [Online; accessed 31-March-2021]. D.1, D.2

[120] PIMBAUD, B. Scrollytelling: Storytelling for The Next Decade. `https://mcfunley.com/continuous-experimentation/`, 2012. [Online; accessed 31-March-2021]. D.2

[121] PIMBAUD, B. Scrollytelling: Storytelling for The Next Decade. `https://bit.ly/3sdifUJ`, 2020. [Online; accessed 31-March-2021]. D.1, D.2

[122] PREIM, B., BAER, A., CUNNINGHAM, D., ISENBERG, T., AND ROPINSKI, T. A survey of perceptually motivated 3d visualization of medical image data. *Computer Graphic Forum 35*, 3 (jun 2016), 501–525, doi: 10.5555/3071534.3071587. 2, 2.5

[123] PREIM, B., AND BOTHA, C., Eds. *Visual Computing for Medicine: Theory, Algorithms, and Applications*, second edition ed. Morgan Kaufmann, Boston, 2014. 2, 2.1

[124] PREIM, B., AND LAWONN, K. A survey of visual analytics for public health. *Computer Graphics Forum 39*, 1 (2020), 543–580, doi: 10.1111/cgf.13891. 2, 2.3, B.3

[125] PREIM, B., OELTZE, S., MLEJNEK, M., GROLLER, E., HENNEMUTH, A., AND BEHRENS, S. Survey of the visual exploration and analysis of perfusion data. *IEEE Transactions on Visualization and Computer Graphics 15*, 2 (2009), 205–220, doi: 10.1109/TVCG.2008.95. 2, 2.2

[126] PREIM, B., ROPINSKI, T., AND ISENBERG, P. A Critical Analysis of the Evaluation Practice in Medical Visualization. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2018), A. Puig Puig, T. Schultz, A. Vilanova, I. Hotz, B. Kozlikova, and P.-P. Vázquez, Eds., The Eurographics Association, doi: 10.2312/vcbm.20181228. 2, 2.5

[127] QU, Z., LAU, C. W., NGUYEN, Q. V., ZHOU, Y., AND CATCHPOOLE, D. R. Visual analytics of genomic and cancer data: A systematic review. *Cancer Informatics 18* (Jan. 2019), doi: 10.1177/1176935119835546. C.2

[128] RAIDOU, R., VAN DER HEIDE, U., DINH, C., GHOBADI, G., KALLEHAUGE, J., BREEUWER, M., AND VILANOVA, A. Visual analytics for the exploration of tumor tissue characterization. *Computer Graphics Forum 34*, 3 (2015), 11–20, doi: 10.1111/cgf.12613. 2, 2.2, 2.2, B.3, C.1, C.2, C.2

[129] RAIDOU, R. G., BREEUWER, M., VILANOVA, A., HEIDE, U. A. V. D., AND HOUDT, P. J. V. The iCoCooN: Integration of Cobweb Charts with Parallel Coordinates for Visual Analysis of DCE-MRI Modeling Variations. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2014), I. Viola, K. Buehler, and T. Ropinski, Eds., The Eurographics Association, doi: 10.2312/vcbm.20141178. C.2

[130] RAIDOU, R. G., CASARES-MAGAZ, O., AMIRKHANOV, A., MOISEENKO, V., MUREN, L. P., EINCK, J. P., VILANOVA, A., AND GRÖLLER, M. E. Bladder runner: Visual analytics for the exploration of rt-induced bladder toxicity in a cohort study. *Computer Graphics Forum 37*, 3 (2018), 205–216, doi: 10.1111/cgf.13413. 2, 2.3, B.3

[131] RIZZO, S., BOTTA, F., RAIMONDI, S., ORIGGI, D., FANCIULLO, C., MORGANTI, A., AND BELLOMI, M. Radiomics: the facts and the challenges of image analysis. *European Radiology Experimental 2* (2018), 36, doi: 10.1186/s41747-018-0068-z. B.2

[132] ROFSKY, N. M., LEE, V. S., LAUB, G., POLLACK, M. A., KRINSKY, G. A., THOMASSON, D., AMBROSINO, M. M., AND WEINREB, J. C. Abdominal MR imaging with a volumetric interpolated breath-hold examination. *Radiology 212*, 3 (1999), 876–884, doi: 10.1148/radiology.212.3.r99se34876. B.4.2

[133] ROPINSKI, T., OELTZE, S., AND PREIM, B. Visual computing in biology and medicine: Survey of glyph-based visualization techniques for spatial multivariate medical data. *Comput. Graph. 35*, 2 (Apr. 2011), 392–401, doi: 10.1016/j.cag.2011.01.011. 2, 2.2, A.2

[134] RSTUDIO TEAM. *RStudio: Integrated Development Environment for R*. RStudio, PBC., Boston, MA, 2020. B.3

[135] SATYANARAYAN, A., AND HEER, J. Authoring narrative visualizations with ellipsis. *Computer Graphics Forum 33*, 3 (2014), 361–370, doi: 10.1111/cgf.12392. D.1, D.2

[136] SCHMAINDA, K., AND PRAH, M. Data from brain-tumor-progression, 2018. A.5, A.6

[137] SCHNABEL, J. A., TANNER, C., CASTELLANO-SMITH, A. D., DEGENHARD, A., LEACH, M. O., HOSE, D. R., HILL, D. L., AND HAWKES, D. J. Validation of nonrigid image registration using finite-element methods: Application to breast MR images. *IEEE Transactions on Medical Imaging 22*, 2 (2003), 238–247, doi: 10.1109/TMI.2002.808367. B.3

[138] SCHROEDER, W., KEN, M., AND LORENSEN, B. *Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, 4 ed. Kitware, Inc., 2006. B.6

[139] SEDLMAIR, M. Design study contributions come in different guises: Seven guiding scenarios. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), BELIV '16, Association for Computing Machinery, doi: 10.1145/2993901.2993913, p. 152–161. 2, 2.5

[140] SELVAN, A. N., COLE, L. M., SPACKMAN, L., NAYLOR, S., AND WRIGHT, C. *Hierarchical Cluster Analysis to Aid Diagnostic Image Data Visualization of MS and Other Medical Imaging Modalities*. Springer New York, New York, NY, 2017, pp. 95–123. C.2

[141] SEYSER, D., AND ZEILLER, M. Scrollytelling – an analysis of visual storytelling in online journalism. In *International Conference Information Visualisation* (2018), doi: 10.1109/iV.2018.00075, pp. 401–406. 2, 2.4, D.1, D.2

[142] SHNEIDERMAN, B. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages* (Sep. 1996), doi: 10.1109/VL.1996.545307, pp. 336–343. 2, 2.1, D.2

[143] SKETCHFAB, I. Sketchfab. `sketchfab.com`, 2021. [Online; accessed 13-March-2022]. D.5

[144] Smit, N., and Lawonn, K. An introduction to evaluation in medical visualization. In *Proceedings of the EuroVis Workshop on Reproducibility, Verification, and Validation in Visualization* (2016), EuroRv3 '16, Eurographics Association, doi: 10.5555/3057025.3057038, p. 41–43. 2

[145] Smit, N. N., Haneveld, B. K., Staring, M., Eisemann, E., Botha, C. P., and Vilanova, A. RegistrationShop: An Interactive 3D Medical Volume Registration System. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2014), I. Viola, K. Buehler, and T. Ropinski, Eds., The Eurographics Association, doi: 10.2312/vcbm.20141193. A.2

[146] Smit, N. N., Kraima, A. C., Jansma, D., Ruiter, M. C. d., and Botha, C. P. A Unified Representation for the Model-based Visualization of Heterogeneous Anatomy Data. In *EuroVis - Short Papers* (2012), M. Meyer and T. Weinkaufs, Eds., The Eurographics Association, doi: 10.2312/PE/EuroVisShort/EuroVisShort2012/085-089. A.2

[147] Soler, J. Drawflow. `https://github.com/jerosoler/Drawflow`, 2021. [Online; accessed 31-March-2021]. D.5

[148] Srivas, C., Maher, A., and Martin, J. Interpretable Adaptive Optimization. `https://machinelearning.apple.com/research/interpretable-adaptive-optimization`, 2021. [Online; accessed 24-March-2022]. D.3

[149] Steenwijk, M., Milles, J., van Buchem, M., Reiber, J., and Botha, C. Integrated visual analysis for heterogeneous datasets in cohort studies. *IEEE VisWeek Workshop on Visual Analytics in Health Care 3* (01 2010). B.3

[150] Stoppel, S., Hodneland, E., Hauser, H., and Bruckner, S. Graxels: Information rich primitives for the visualization of time-dependent spatial data. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (sep 2016), doi: 10.2312/vcbm.20161286, pp. 183–192. 2, 2.1, 2.2, A.2

[151] Stornaway.io. Stornaway.io. `https://www.stornaway.io/`, 2021. 2, 2.4, D.2, D.3

[152] Sullivan, C. B., and Kaszynski, A. PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *Journal of Open Source Software 4*, 37 (May 2019), 1450, doi: 10.21105/joss.01450. C.5

[153] Sultanum, N., Chevalier, F., Bylinskii, Z., and Liu, Z. Leveraging text-chart links to support authoring of data-driven articles with vizflow. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (2021), CHI '21, Association for Computing Machinery, doi: 10.1145/3411764.3445354. D.2

[154] Thorndike, R. L. Who belongs in the family? *Psychometrika 18*, 4 (Dec. 1953), 267–276, doi: 10.1007/bf02289263. C.4

[155] Thrun, M. C. *Approaches to Cluster Analysis*. Springer, 2018, pp. 21–31. C.4

[156] TIBSHIRANI, R., WALTHER, G., AND HASTIE, T. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63*, 2 (2001), 411–423, doi: 10.1111/1467-9868.00293. C.4

[157] TOMASULO, F. Story and discourse: Narrative structure in fiction and film by seymour chatman. *Journal of the University Film Association 32* (1980), 71–74, doi: 10.2307/20687541. D.3

[158] TONG, C., ROBERTS, R., BORGO, R., WALTON, S., LARAMEE, R. S., WEGBA, K., LU, A., WANG, Y., QU, H., LUO, Q., AND MA, X. Storytelling and visualization: An extended survey. *Information 9*, 3 (2018), doi: 10.3390/info9030065. 2, 2.4, D.2, D.3

[159] TORY, M., AND MOLLER, T. Rethinking visualization: A high-level taxonomy. In *IEEE Symposium on Information Visualization* (2004), doi: 10.1109/INFVIS.2004.59, pp. 151–158. 2, 2.1

[160] TURKAY, C., LEX, A., STREIT, M., PFISTER, H., AND HAUSER, H. Characterizing cancer subtypes using dual analysis in caleydo stratomex. *IEEE Computer Graphics and Applications 34*, 2 (2014), 38–47, doi: 10.1109/MCG.2014.1. B.3

[161] TWINE. Twine. `https://twinery.org/`, 2021. 2, 2.4, 3.3.2, D.2, D.3

[162] VAIH-BAUR, C. Scrollytelling. In *Textspiele in der Wirtschaftskommunikation*. Springer Fachmedien Wiesbaden, 2017, pp. 343–346. D.2

[163] VAN DER MAATEN, L., AND HINTON, G. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research 9*, nov (2008), 2579–2605. Pagination: 27. B.3, B.5.1, B.6

[164] VAN WIJK, J., AND NUIJ, W. Smooth and efficient zooming and panning. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No.03TH8714)* (2003), doi: 10.1109/INFVIS.2003.1249004, pp. 15–23. D.4.3

[165] VON LANDESBERGER, T., BREMM, S., KIRSCHNER, M., WESARG, S., AND KUIJPER, A. Visual Analytics for model-based medical image segmentation: Opportunities and challenges. *Expert Systems with Applications 40*, 12 (2013), 4934–4943, doi: 10.1016/j.eswa.2013.03.006. B.3

[166] WADE, B. S. C., HELLEMANN, G., ESPINOZA, R. T., WOODS, R. P., JOSHI, S. H., REDLICH, R., DANNLOWSKI, U., JORGENSEN, A., ABBOTT, C. C., OLTEDAL, L., AND NARR, K. L. Accounting for symptom heterogeneity can improve neuroimaging models of antidepressant response after electroconvulsive therapy. *Human Brain Mapping 42*, 16 (2021), 5322–5333, doi: 10.1002/hbm.25620. 4.4

[167] WANG, Z., ROMAT, H., CHEVALIER, F., RICHE, N., MURRAY-RUST, D., AND BACH, B. Interactive data comics. *IEEE Transactions on*

Visualization and Computer Graphics 28, 01 (jan 2022), 944–954, doi: 10.1109/TVCG.2021.3114849. D.2

[168] WANG BALDONADO, M. Q., WOODRUFF, A., AND KUCHINSKY, A. Guidelines for using multiple views in information visualization. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (New York, NY, USA, 2000), AVI '00, Association for Computing Machinery, doi: 10.1145/345513.345271, p. 110–119. 2, 2.1

[169] WICKHAM, H., HOFMANN, H., WICKHAM, C., AND COOK, D. Glyph-maps for visually exploring temporal patterns in climate data and models. *Environmetrics 23*, 5 (2012), 382–393, doi: 10.1002/env.2152. A.2

[170] WOHLFART, M., AND HAUSER, H. Story telling for presentation in volume visualization. In *Eurographics/ IEEE-VGTC Symposium on Visualization* (2007), The Eurographics Association, doi: 10.2312/VisSym/EuroVis07/091-098. D.2, D.4.1

[171] YTRE-HAUGE, S., DYBVIK, J. A., LUNDERVOLD, A., SALVESEN, Ø. O., KRAKSTAD, C., FASMER, K. E., WERNER, H. M., GANESHAN, B., HØIVIK, E., BJØRGE, L., ET AL. Preoperative tumor texture analysis on MRI predicts high-risk disease and reduced survival in endometrial cancer. *Journal of Magnetic Resonance Imaging 48*, 6 (2018), 1637–1647, doi: 10.1002/jmri.26184. 3.1.1, B.1, B.2, B.4.3, B.4.3, B.4.4, C.1, C.3

[172] YU, L., JIANG, H., YU, H., ZHANG, C., MCALLISTER, J., AND ZHENG, D. iVAR: Interactive visual analytics of radiomics features from large-scale medical images. In *2017 IEEE International Conference on Big Data* (2017), IEEE, doi: 10.1109/BigData.2017.8258398, pp. 3916–3923. C.2

Graphic design: Communication Division, UiB / Print: Skipnes Kommunikasjon AS

uib.no

ISBN: 9788230862193 (print)
9788230869369 (PDF)