

sMolBoxes: Dataflow Model for Molecular Dynamics Exploration

Pavol Ulbrich, Manuela Waldner, Katarína Furmanová, Sérgio M. Marques,
David Bednář, Barbora Kozlíková, and Jan Byška




Fig. 1. An overview of the sMolBoxes interface for molecular dynamics exploration, consisting of the 3D View, the Canvas with sMolBox Nodes, and the detailed bigBox View for navigation through the 3D animation.

Abstract—We present sMolBoxes, a dataflow representation for the exploration and analysis of long molecular dynamics (MD) simulations. When MD simulations reach millions of snapshots, a frame-by-frame observation is not feasible anymore. Thus, biochemists rely to a large extent only on quantitative analysis of geometric and physico-chemical properties. However, the usage of abstract methods to study inherently spatial data hinders the exploration and poses a considerable workload. sMolBoxes link quantitative analysis of a user-defined set of properties with interactive 3D visualizations. They enable visual explanations of molecular behaviors, which lead to an efficient discovery of biochemically significant parts of the MD simulation. sMolBoxes follow a node-based model for flexible definition, combination, and immediate evaluation of properties to be investigated. Progressive analytics enable fluid switching between multiple properties, which facilitates hypothesis generation. Each sMolBox provides quick insight to an observed property or function, available in more detail in the bigBox View. The case studies illustrate that even with relatively few sMolBoxes, it is possible to express complex analytical tasks, and their use in exploratory analysis is perceived as more efficient than traditional scripting-based methods.

Index Terms—Molecular dynamics, structure, node-based visualization, progressive analytics

1 INTRODUCTION

Studying molecular dynamics (MD) is one of the fundamental tasks in biochemical research, biotechnology, molecular medicine, or drug design. It helps scientists understand molecules' behavior and thus improve or modify their properties. In our research, we have been cooperating with a group of protein engineering experts for more than a decade, trying to support their workflows with dedicated visual repre-

sentations. Their main interest lies in designing mutations of protein structures in order to change protein properties, such as its stability [19] or activity towards other molecules [27]. In order to reveal candidate amino acids for mutation and evolution, biochemists need to explore several geometric and physico-chemical properties of the protein and their changes over time. For this purpose, molecular dynamics simulations are of the most powerful computational methods capturing molecular behavior.

Nowadays, molecular dynamics simulations span up to hundreds of thousands of time frames [12], as they try to capture important structural changes happening in the order of nanoseconds, microseconds, or even longer time span. The vast amount of snapshots prevents the experts from exploring and visually analyzing the simulations by traditional 3D animation techniques, commonly used for this purpose. For example, assuming the playback speed of 5 frames per second, users would have to spend approximately 22 hours observing the animation of a simulation consisting of 400,000 time frames. This is not feasible, and thus the exploration process has to be supported by other approaches. Instead of observing the whole simulation, protein engineers are trying to reveal interesting parts of the simulation, denoted as *events*. These can be derived from various descriptive measures or features, capturing geometric (e.g., the distance between atoms, angles, root-mean-square

- Pavol Ulbrich, Katarína Furmanová, and Barbora Kozlíková are with VisiTab, Faculty of Informatics, Masaryk University, Czech Republic.
- Jan Byška is with VisiTab, Faculty of Informatics, Masaryk University, Czech Republic, and University of Bergen, Bergen, Norway. E-mail: byska@mail.muni.cz.
- Manuela Waldner is with TU Wien, Vienna, Austria.
- Sérgio M. Marques and David Bednář are with Loschmidt Laboratories, Department of Experimental Biology and RECETOX, Faculty of Science, Masaryk University, Brno, Czech Republic, and International Clinical Research Center, St. Anne's University Hospital Brno, Brno, Czech Republic.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

deviation) or physico-chemical properties (e.g., electrostatic or van der Waals interaction energies, hydration) of the simulation. These measures can be visualized in simple 2D plots where the experts are trying to reveal outliers and patterns and, even more importantly, how these measures interplay or correlate. However, with MD simulations capturing proteins consisting of hundreds of amino acids and their interactions with multiple molecules, there are countless measures to explore. Extensive scripting would be required to generate plots for selected measures and to obtain the 3D animation sequences associated with observed patterns or events in these measures. Consequently, protein engineers usually restrict their analysis to a few structures, such as the amino acids responsible for a reaction, making it unlikely to discover unexpected interactions or events.

In our previously designed approach [5], we were trying to encourage exploration of the measures by providing the experts with a user interface where they could define aggregated measure functions representing the relevance of each time frame. Although the interface, showing superimposed functions, was well appreciated by the experts, we spotted the main drawbacks of this approach during evaluation. The interface offered only limited features for defining aggregated functions, which proved insufficient for describing some of the more complex events in simulations. Moreover, even with the simple aggregations, it was difficult for the experts to keep track of the combinations and the provenance of functions superimposed in a single view—they quickly lost the mental model of the exploration process. Considering the number of values and the multi-variate and time-series character of the measures the domain experts must scrutinize, our experience revealed that appropriate visual support for this process is crucial but currently lacking.

Therefore, our work aims to fill in this gap and facilitate the complex domain-specific problem of analysis of long MD simulations. We propose a visual analysis approach supporting the interactive exploration of multiple properties and their combinations. In order to simplify the discovery of unexpected *events*, we aim to enable analysts to easily query and navigate long simulation sequences through flexible selections of spatial structures, quick computation of measures associated with these structures, and linking between 2D plots of the measures and 3D visualizations. We call our novel solution sMolBoxes.

We based sMolBoxes on a dataflow network model that was already demonstrated to be intuitively understood by domain experts [14,28,37]. In the network, each node represents either an input, a measure, or combination of measures. By connecting nodes, users can create complex functions describing molecular behavior. Moreover, the network layout provides an overview of all functions, their settings, and aggregations at the same time.

The evaluation of sMolBoxes was performed with a protein engineering expert, who tested the proposed interface on two exemplary scenarios from his daily workflow (presented in Section 6). We demonstrate that with our tool, the domain expert was able to perform the analysis in a fast manner, and that the design of the system encouraged him to explore the data beyond his typical workflow steps. The main contributions of our work are thus the following:

- Easily adaptable definitions of complex aggregated MD simulation measures that can significantly speed up previous scripting-based approaches to MD simulation analysis and support hypothesis generation.
- A flexible network layout providing the users with a constant overview of the provenance of the molecular interaction analysis and thus reducing their mental workload when working with complex aggregations.
- Progressive analytics of large MD trajectories facilitating immediate previews of calculated measures and functions refined in an iterative manner.
- Direct linking of the constructed molecular interaction analysis network with a 3D representation of the protein, enabling the selection of inputs and qualitative inspection of outputs, which is unavailable when using standard statistical analysis.




Fig. 2. Illustration of input molecular structures associated with molecular dynamics simulations.

2 DATA—TASKS—REQUIREMENTS

In this section, we first describe the *data* we are working with and the related domain-specific terminology. Then we discuss the workflow traditionally followed by the experts when analyzing MD simulations and the main *tasks*, given by the domain experts, that we are aiming to support with this work. Finally, we discuss the design *requirements* for our system that we derived from the tasks.

2.1 Data and Terminology

As already stated, we are focusing on the analysis of molecular systems with thousands of atoms and molecular dynamics simulations of their movements, which can consist of hundreds of thousands of time frames. Due to the amount of data, binary formats which store only information about the position of atoms over time are often used. Therefore, some properties, like mutual distance or energy between atoms, have to be extracted afterward. For some of these properties, this may require nontrivial computation time. While our research is primarily driven by the needs of biochemists, the proposed solution can be applied to similar datasets consisting of positional information of particles over time.

Molecules can be further described as multi-scale structures where individual atoms form *amino acid residues* (often referred to as *residues*), which are then connected into linear structures called *chains* (see Figure 2). Each molecule in the simulation then consists of one or multiple of these chains. Additionally, the simulation may contain water molecules forming the *solvent* (outer environment of the protein) or other small molecules called *ligands*. These are playing important role in protein reactivity. In protein-ligand simulations, ligands bind to the *active site* of the protein (see Figure 2), and protein engineers analyze which characteristics and properties influence the success and speed of this binding. The simulation of binding can be augmented with additional time-series data computed with different tools, such as binding energies or various properties computed for the hollow cavities (e.g., *tunnels*) [3] inside the protein, often serving as access pathways for the ligand to the active site. The narrowest part of the tunnel is called the *tunnel bottleneck* and it is formed by multiple *bottleneck residues*. The radius of the tunnel bottleneck varies over time and its opening and closing has impact on the tunnel’s permeability.

2.2 Tasks

Our proposed work directly builds on lessons learned from our previous work [5], where we have identified several requirements for the analysis of long MD simulations in protein engineering based on 3D animations. However, the aforementioned work primarily addresses design considerations for spatio-temporal focus+context in the 3D animations and offers very limited flexibility for expressing the investigation focus.

Our previous requirements can therefore be considered orthogonal to the requirements of the present work.

Within the period of almost two years, we conducted numerous sessions with always at least six protein engineering experts on different levels of seniority (two senior researchers, two postdocs, and Ph.D. students). Of these experts, one senior researcher and one postdoc are co-authors of this paper. These sessions had two main goals: 1) shadowing, i.e., observing the experts in their daily routine with the existing workflows and tools; 2) discussing the benefits and limitations of the workflows and tools, which could be addressed by visualization.

Within these sessions, we specifically analyzed three works published by the protein engineering experts, which focused on different aspects of molecular behaviour [18, 20, 21]. Based on the analysis of these cases and our observations, we have identified the following tasks for our proposed solution:

- T1: Property exploration.** The biochemists typically start the analysis of MD simulations by exploring various properties that could navigate them to important snapshots of the simulation, such as distances or angles between residues [5]. Another important aspect is also the correlation between these properties. One example of such a case is the identification of gating residues, which control the access to the protein's active site [18]. Gating residues usually exhibit complex movement, often in concordance with other residues. In this case, the biochemists would explore the root-mean-square deviation (RMSD) of the entire protein or movable parts of the protein, the distance between various residues, their angles, and the correlation between these measures on multiple residues and substructures of the protein. In their traditional workflow, experts calculate the bulk or the time-evolution of the properties using in-house bash or python scripts, jupyter notebooks, or Markov state models in order to cluster the conformational states and estimate their properties. All these analyses are hypothesis-driven, and, due to the large amount of data that is produced, it is easy to overlook relevant details not considered beforehand. The number of properties and residues the biochemist usually explore is limited to a certain extent.
- T2: Identification of important snapshots.** Once the domain experts identify the most important properties, they want to use them to identify the relevant snapshots of the simulation, i.e., time-stamps carrying potential causes for the observed values or change. At this stage, they typically formulate a hypothesis about their observations. For example, RMSD changes indicate conformation changes of the protein, and through visual inspection, it may be possible to speculate what might have triggered them. Another example is finding a geometric arrangement of the protein in which a reaction could take place [21]. Such molecular conformation can be identified as a snapshot, in which the distance between two selected molecular structures is below a certain threshold, and the angle at which they approach is within an acceptable range.
- T3: Causality and conditions analysis.** Finally, when the important snapshots are identified, the biochemists study the conditions and causes of the given state in greater detail to validate their hypothesis. Here, they sometimes extend their analysis with new properties and structures, if they suspect these can play a role in the inspected event. Additionally, the retrospective analysis of the properties can help them identify possible modifications to the protein that could lead to change in its behavior. An example could be a discovery of residue that is preventing the binding of a ligand [18]. Such residue could then be experimentally replaced to increase the reactivity of the protein and enable the entrance of a ligand into the active site.

These tasks are typically performed iteratively. Through visual analysis of causality and events (T3), interesting structures and behaviors might be discovered and subsequently analyzed in more detail for their properties (T1). However, since tasks T1 and T2 are tedious to perform with traditional tools, T3 is done relatively rarely nowadays.

2.3 Requirements

Based on the stated tasks and the limitations we identified in the current workflows of protein engineers, we derived the following design requirements for our system:

- R1: Selection of spatial inputs on multiple levels of detail**, e.g., atoms, residues, or even entire molecules. To start the analysis (T1), the system should allow the users to flexibly select, modify, and combine molecular (sub)structures to be analyzed.
- R2: Predefined measure functions** that compute a selected metric on one or multiple spatial inputs for each snapshot without the need for scripting. This is important for initial exploration of the properties (T1), but the measure functions also serve as indicators of snapshot importance (T2). The system should support a range of measures typically associated with MD simulation analysis, such as the distance between two spatial inputs or RMSD.
- R3: Flexible aggregations** and combinations of the measure functions. The system should support the aggregation of multiple measures with aggregation operations, such as extrema or average, to combine the measure functions into a single output defining the investigation focus (T2).
- R4: Data provenance representation** to support analytical scenarios where the exploration might lead to complex aggregated functions that are composed of multiple inputs and many different measures. In such cases, it is impossible for the users to keep the mental model of the constitution of aggregated functions. Nevertheless, it is vital for the analysis of the causes and effects (T3) of MD simulation phenomena. The system should provide a representation of the function model and the effects individual components have on the resulting function.
- R5: Real-time performance** of the structure loading and function computation. These operations need to be performed without any significant latency to enable a fluid analysis of the desired molecular properties (T1, T2).
- R6: Visualisation of large time-series data** containing hundreds of thousands of snapshots. Due to their size, measures extracted from long MD simulations cannot always be displayed in their full resolution and require dedicated solutions for their visualization (T1-T3).
- R7: Linking of measures with the 3D representation** to support tight coupling between qualitative inspection and in-place analysis of interesting observations and in-depth quantitative analysis, since both are needed for investigation of causality in MD simulations (T3).

3 RELATED WORK

In this work, we present a system for the visual analysis and exploration of MD simulations. Therefore, this section will first discuss related visual analytics approaches from this domain area. Furthermore, in our presented work, we take inspiration from dataflow models, which have already been successfully used in several applications for analytics in various domains, including life sciences. Therefore, in the second part of this section, we discuss these approaches.

3.1 Visual Analysis of MD Simulations

An extensive study performed by Hensen et al. [11] proved that there is a strong correlation between the dynamics of molecular structure and its function. Therefore, without a proper understanding of protein dynamic behavior, we cannot fully understand its function. There are already many existing approaches aiming to provide insight into MD simulations [22].

One possible approach is to aggregate the desired information about the simulation and represent it in a static overview visualization. Patro et al. [26] measure the saliency in MD simulations which guides the selection of representative and anomalous time frames summarizing

the MDs. In their following research [25], they represent the MD simulation as a state transition graph, illustrating the flow of a biomolecule through the trajectory space. Bryden et al. [4] use arrow glyphs to depict the molecular flexibility calculated from normal mode analysis. First, they cluster atoms according to their common movement patterns. These clusters are then annotated with arc-shaped arrows in 3D representation of the molecule. When the simulation contains many spatially-distributed trajectories of molecules, such as lipids, the overview representation often utilizes streamline-based representations, such as in MolPathFinder [1] or in the approach proposed by Ertl et al. [9]. In the latter, the authors perform a spatial and temporal aggregation of ions transported through a porous structure, resulting in a velocity field visualization. Bidmon et al. [2] present an approach to visual abstraction of solvent paths traversing through a protein. Their aggregated visualization also encodes information about the direction and velocity of the solvent. The overview is reached by clustering similar pathlines and showing only principal paths. Vad et al. [35] also focus on studying solvent paths when they propose a set of visualizations for interactive filtering and exploration of solvent trajectories and observing the space inside protein occupied by solvent over time. Schatz et al. [29] aggregate information about molecular interactions in the whole MD simulation and display it in an enhanced amino-acid sequence diagram. However, in this case, the temporal aspect of the simulation is lost. In general, the overview visual representations of the MD simulations are very powerful and informative, giving the user the first insight without tedious frame-by-frame exploration. On the other hand, to the best of our knowledge, none of the existing approaches is suitable to facilitate a multifaceted interactive exploration of the MD simulation using traditional visual encodings for geometric and physico-chemical measurements.

There are also many existing approaches with the main focus on simulations of only a few molecules or particles and a more detailed exploration of changes in their properties over time. Furmanová et al. [10] present a visual analysis tool for detailed exploration of geometric and physico-chemical properties of individual ligand paths inside a protein. Although this approach could be useful in our case, the main limitation is the narrow focus of this tool—it can analyze only a single ligand path over time. Similarly, Vázquez et al. [36] study the path of a single ligand through the protein in a very abstracted way. Even though the proposed visualization is very space-saving and enables plotting many properties simultaneously, it only relies on abstract visualizations. It does not support tight coupling between abstract and 3D visualization for more exploratory analysis. Additionally, it does not intuitively allow combining these properties to detect important events in the MD simulation. This is partially solved in the approach by Duran et al. [7], which enables to explore a limited set of properties of more ligands at the same time. However, for each ligand, the tool creates an individual plot, and therefore, it is not extendable to more complex tasks, such as analyzing many ligands and solvents at the same time. Skånberg et al. [32] are following a similar idea and convey selected properties in 2D plots. Additionally, they show the density field depicting the spatial distribution of molecules within the simulation. However, the support for combining properties and thus deriving important events is missing. As we observed in our previous study [5], individual measures are not sufficient to describe complex events or behaviors, which we aim to address using sMolBoxes.

3.2 Navigation Systems for Visual Exploration

In our solution, we are building upon so-called modular visualization environments [34], where the user interacts with the system by directly manipulating objects in the scene. Combining these objects creates the mental map and resulting visual representation of the whole analysis process. Our work was inspired by visualization systems relying on dataflow diagrams to facilitate interactive visual exploration. One prominent example is the system by Waser et al. [37], who proposed a dataflow diagram for interactive exploration of the simulation setup of the flood emergency simulation system. The system consists of a set of nodes, in which each node type can contain a set of connectors, serving to establish a relationship with other nodes. However, the complexity of

the studied tasks often results in a complicated visual interface, as the nodes can also encode quite complex information, such as a 3D View. Dataflow diagrams were also employed by Yu and Silva [38]. They use them for specifying relations between components that process, filter, or visualize the data and can be combined by interacting with the visual interface. Similar examples are *DataMeadow* [8], *FlowSense* [39], or *ExPlates* [13]. These systems typically operate by connecting data source nodes with multiple data operators and visualization target nodes, which specify the visual mapping [23]. In contrast, sMolBoxes do not specify visualizations but rather serve as navigation aid through the long trajectories and for highlighting within the crowded molecular scene. The 3D visualization itself relies on conventional visual mappings for molecular data, such as balls-and-sticks or ribbon representation.

In medical imaging, the modular visualization environment was used in the MeVisLab tool [28]. Such an environment also forms an integral part of the system presented by Meyer-Spradow et al. [24]. Their Voreen framework consists of so-called processors—autonomous functional building blocks serving for specific volume rendering tasks. By combining these blocks into dataflow networks, the user can define the visual appearance of the rendered object. A similar idea was used already before in the VisTrails [31] toolkit.

The concept of modular environments is followed not only by researchers solving tasks using visual interfaces, but the same principle is used in visual programming. Here, instead of writing lines of codes, users are generating the program by visually connecting nodes with specific functions. The principles of visual programming are well implemented in several tools across multiple domains, such as VVVV [33] for graphics and media, LabVIEW [17] for engineering, and SimuLink [6] for simulations.

Recent development can be seen in the Inviwo tool [14] and its direct application into molecular exploration in the VIA-MD software [32]. However, these tools do not rely on a dataflow model in true sense. Instead, they are limited to displaying individual metrics and are not designed to describe rather complex semantic behaviors. In addition, they offer only limited linking between the 3D view and the 2D time-series plots.

In summary, to the best of our knowledge, dataflow models have not yet been used for the exploratory analysis of long MD simulations.

4 sMOLBOXES DESIGN

Our tool is conceptually similar to dataflow networks. We represent each node in the network by a *box* that we call sMolBox (see Figure 3). By analyzing the workflows of our collaborators (see Section 2), we concluded that we could support their complex analyses with just three types of nodes. The sMolBoxes act as input selectors, property measures, or aggregation functions. By forming connections between these boxes, users can easily and quickly explore various biochemical properties and their combinations (T1). The network also produces an *importance function* expressing the user-defined notion of relevant events over time. This function can be used to guide the analysis and navigate the users to the most interesting parts of the data (T2).

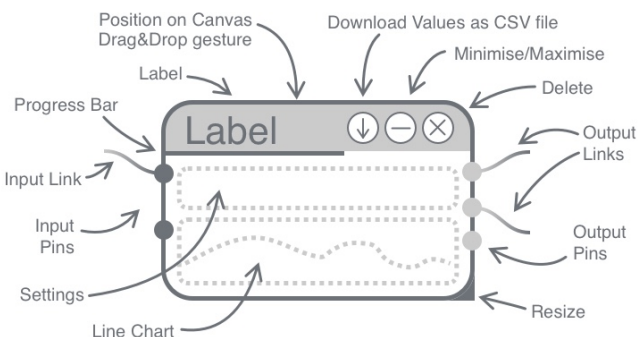


Fig. 3. Illustration of the general sMolBox layout and its components.

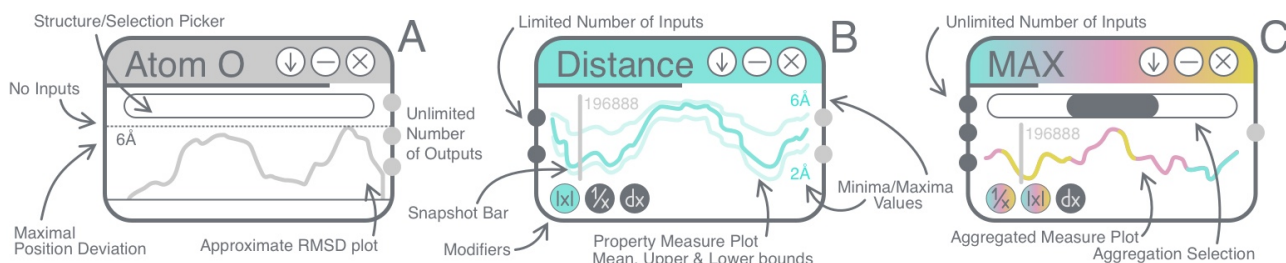


Fig. 4. Each sMolBox has different layout, based on its function: (left) Input sMolBox, (middle) Measure sMolBox, and (right) Aggregation sMolBox.




Fig. 5. Overview and categorization of the sMolBox instances.

Finally, since the network can be easily adjusted, additional measures can be added to study the possible causes of the identified events and observations in greater detail (T3).

The three types of nodes (Figure 4) in the sMolBox network have the following roles. We need *Input sMolBox* for specifying the spatial data to be analysed (R1), *Measure sMolBox* for calculating various properties of the selected data (R2), and *Aggregation sMolBox* for combining the properties (R3). The network layout of sMolBoxes, enhanced by visual cues, naturally provides an overview of data provenance (R4).

To provide users with a consistent visual design, all sMolBoxes have a similar layout with a preview line chart as the main element, as shown in Figure 3. To motivate the users to build the dataflow network from left to right so they utilize the available space most efficiently, we place the connection pins on the sides—input ones on the left and outputs on the right. The boxes may be linked by an intuitive drag&drop interaction with the pins. The system prevents faulty connections that would violate the correct flow of data, such as an Input sMolBox following a Measure sMolBox. In order to highlight the flow of data in complex networks, we color the links with a lighter color near the parent box and a darker color near the child box. All sMolBoxes are designed to be fully responsive, and all their components, including the time-series visualizations, are computed in real-time, in parallel, and are not blocking any other process (R5). We discuss the progressive data processing approach that facilitates this in Section 5.

The final application consists of three connected views (see Figure 1): Canvas, bigBox View, and 3D View. The *Canvas* is the analyst’s central workbench for adding, connecting, and managing individual sMolBoxes. When the analyst selects a sMolBox on the Canvas, the *bigBox* View shows the corresponding measure in detail. Finally, the *bigBox* or any sMolBox can be used to control the linked *3D View* for qualitative inspection of behaviors observed in the measures (R7). Below, we describe the three types of sMolBoxes and their interplay with the proposed views in more detail.

4.1 Input sMolBoxes

Protein engineers often start their analysis by investigating the behavior of ligands or relevant parts of the protein, such as protein residues directly involved in chemical reactions (e.g., *catalytic residues*). To allow a direct and quick selection of molecules or their parts on multiple levels of detail (R1), we provide Input sMolBoxes (see Figure 4 left).

Using Input sMolBox, the analyst can select a single atom, residue, or molecular chain by its respective ID or selecting it directly in the 3D View. To support the exploration of larger areas of interest, e.g., multiple catalytic residues at the same time, we also have to provide users with the means to create arbitrary spatial selections anytime during their exploration. Here we benefit from the existing functionality inside the CAVER Analyst software [15], where we implemented our prototype. Users can perform complex selections through direct interaction in the 3D View, by interacting with the Sequence View showing the color-coded list of amino acids forming the protein, or by typing commands in the built-in console. The selections can be then mapped to Input sMolBoxes. To enable users to easily relate Input sMolBoxes to the corresponding parts of the molecule later in the analysis process (R4), we highlight the corresponding parts in the 3D View every time the user clicks on an Input sMolBox.

Each Input sMolBox provides a preview line chart indicating to the users potentially interesting time frames (R2) with significant conformation changes already during the data selection process. Conformation changes indicate that the structure has changed its shape or position; they may therefore represent an interesting behavior to investigate in more detail. For Input sMolBox with a selected atom, this line chart corresponds to the RMSD of the atom with respect to the first snapshot of the simulation. For inputs consisting of multiple atoms, the RMSD is approximated (see Section 5.3) to provide a quick indication of the structure’s spatial development over time. To allow users to estimate the size and the relevance of these conformation changes, we display upper and lower bounds in Ångström units next to the line chart.

4.2 Measure sMolBoxes

Measure sMolBoxes compute geometric properties on the connected Input sMolBoxes for every snapshot. This way, the analyst can start extracting properties describing the behavior of molecules or their parts in time (R2). Measure sMolBoxes require one or multiple Input sMolBoxes depending on the function they are computing. For example, one Input sMolBox is enough to compute *RMSD*, but *distance* requires two inputs.

A major part of the Measure sMolBox space is occupied by the line chart depicting the computed measure over time (see Figure 4 center). We have opted for a line chart as it is easy to understand, provides a well-known representation of trends, and is also predominantly used by protein engineers in their traditional analyses. However, given the amount of data, it is not possible to accurately display all values, as it would require hundreds of thousands of horizontal pixels (R6). Therefore, we bin the function by the number of available pixels, and we visualize the average and the upper and lower bound of the values within each bin using polylines.

To enable qualitative inspection of behaviors observed in the line chart, we interactively linked it with the 3D View via the *Snapshot*




Fig. 6. Analysis of the release of ligands in an MD simulation with 400,000 snapshots. We selected an active site residue (1) and three ligands (2) in the 3D View (A), then we calculated their minimum distance (blue, green, and purple distance sMolBox) using MIN aggregation sMolBox. We visually compared the proximity to the active site with the RMSD (yellow sMolBox) of the bottleneck residues (3) and the bottleneck throughput loaded as an external file (red sMolBox, 4). After normalization, we were able to use the MAX sMolBox to evaluate the most important time-frames in the simulation with respect to all the measures. Subsequently, we utilized the MIN sMolBox to investigate the role of tunnel cost measure on the interesting snapshots. It led us to the part around snapshot 114,620, where we suspected a movement of ligands through the protein bottleneck. The detailed examination of the tunnel radius (5) using bigBox (C), followed by a more precise observation in the 3D View (A), has shown a sudden entry and exit of one of the ligands to the active site. This observation would not be feasible with a mere analysis of a distance graph.

shot Bar (R7). This vertical bar is situated over the line chart in the approximate position of the currently depicted snapshot in the 3D View. The bar position, and thus the snapshot depicted in the 3D View, can be adjusted by dragging the bar along the chart. For binned functions, repositioning of the bar selects the first snapshot from the corresponding bin. We depict the current snapshot number next to the bar for better orientation. If the Measure sMolBox is big enough, we also show the calculated value of the displayed measure.

Each Measure sMolBox is given a unique color, which is depicted on top of the sMolBox and in the line chart representation. This can be seen in Figure 5, which provides the overview of all sMolBox instances supported by our tool. We decided to use color for Measure sMolBoxes and not for Input sMolBoxes because the measured property gives a particular meaning to the input. For example, having two residues, we can measure their distance, binding energy, or other biochemical properties. If we colored respective Measure sMolBoxes by their input, they would all have the same color. However, by assigning a unique color to each measure, we ensure they are easily recognizable from each other in the network. Moreover, this way, we can also better communicate the provenance of combined measures (R4), as will be described in the next section.

4.3 Aggregation sMolBoxes

A major challenge when exploring MD simulations is the combined analysis of various measures. For example, protein engineers may want to detect when any of the ligands present in the simulation enters the protein active site. This process requires computing the distance between every ligand and the active site and consequent comparison of obtained measures (R3). To support such analyses, we designed Aggregation sMolBoxes. They accept an unlimited number of Measure sMolBoxes and Aggregation sMolBoxes as input and calculate a new function as a combination of all parent measures according to the selected type. Based on discussion with the domain experts, we have selected the following three aggregation functions that are most often used in their everyday workflow: *Minimal*, *Maximal*, and *Median* function. The type of the Aggregation sMolBox can be changed at any

time using a horizontal categorical selector (see Figure 4 right).

Aggregation sMolBoxes are grayscale at first, but they inherit the colors from their parent Measure sMolBoxes (or Aggregation sMolBoxes). To ensure the users can keep an overview of measures forming the aggregated function (R4), the inherited colors are used for the top part of the Aggregation sMolBox, connection links, and the final aggregated function depicted as a line chart. The colored links help with better navigation in the network representation. Furthermore, by coloring the function depicted in the Aggregation sMolBox, we convey to the user the time intervals where, for example, in the case of maximum aggregation, a particular measure was dominating (see Figure 6).

Instead of showing upper and lower bounds, the line chart in Aggregation sMolBox focuses solely on the average value of the aggregated bins. We opted for this solution because the maxima and minima for the aggregated intervals would not be connected into a smooth trend line and would increase clutter in the visualization. Similarly to the Measure sMolBox, also here, the users can utilize the Snapshot Bar for easy navigation through the time series.

4.4 Modifiers

For some analytical tasks, smaller values of the measured properties are more important (e.g., a small distance between protein and ligand residues implies an interaction). In other cases, larger values of the same measure are of interest (e.g., an increasing distance between the residues may imply the opening of the protein tunnel). To support these analytical tasks, both Measure sMolBoxes and Aggregation sMolBoxes offer a set of Modifiers—linear unary functions that take the computed importance function as a single input and adjust all the values according to the Modifier definition. We have implemented the following three Modifiers:

- **Normalization** can be used when the analyst is trying to compare seemingly incomparable measures, such as the distance between two atoms of the same residue, typically ranging in low numbers, and the distance between molecule and ligand that can be in the order of magnitudes higher. With normalized values, users can compare relative changes in the measures.




Fig. 7. The progressive calculation of RMSD for a ligand demonstrates the change in the line chart for 256 (A), 1024 (B), 65,536 (C), and 400,000 processed MD snapshots.

- **Inverse** can be used in case the analyst wants to compare a maximum of one property measure against a minimum of another, which can be easily flipped using this Modifier.
- **Derivation** can be used to identify rapid changes in the measure represented by the sMolBox.

Modifiers are represented by circular icons at the bottom left corner of sMolBoxes (see Figure 4 center and right). We opted for this solution because the Modifiers will always have only one input, and having them as separate nodes would unnecessarily complicate the final network. It also provides means for the users to quickly perform fast analysis of the measures, for example, by exploring the rate of change of the corresponding measure via derivatives.

In the beginning, all Modifiers are in a non-active state. They can be activated by clicking on their corresponding icons. Since not all Modifiers are commutative, their order might also be altered by easy drag&drop interaction.

4.5 bigBox

By connecting individual sMolBoxes, users can quickly explore various biochemical properties and their combinations (**R1-R3**). However, due to their restricted size, sMolBoxes can only provide a limited view of the investigated time series, the lengths of which can reach hundreds of thousands of time steps (**R6**). To support a detailed analysis of these more extensive measures, users can select any sMolBox as their current analysis focus. We then depict the corresponding measure in the linked bigBox View, which utilizes the whole width of a screen and thus shows the selected importance function in a better temporal resolution.

The bigBox also serves as the primary interface to control the 3D View. It thereby replaces the classic VRC-controls to play and pause the animation and the traditional navigation bar for scrubbing through the animation without any visual guidance. The bigBox supports scrubbing with respect to measure associated with the selected sMolBox. This way, the displayed measures can be qualitatively explained with the associated 3D visualization (**R7**).

4.6 3D View

As noted in our previous focus groups [5], 3D animations are useful to qualitatively characterize quantitative observations and to get a better intuition about spatial relations. For example, exploration of the RMSD for a molecule only indicates that the conformation changed, but it does not convey any information on *how* it changed. With the linked 3D View, it can be checked effortlessly (**R7**). However, dense molecular environments pose a significant challenge because the most interesting parts of the molecule may be invisible due to occlusions. sMolBoxes take liberty on several standard molecular representations with various levels of abstractions (e.g., van der Walls, ball-and-stick, cartoon, or backbone representations) implemented in CAVER Analyst software [15]. The user can choose different representations for the spatial selections represented by Input sMolBoxes. This way, users can visualize the selected structures in full detail while keeping the rest of the molecule abstracted as context.

Linking the Canvas with the 3D View enables a new workflow, where analysts can fluidly switch between rapid quantitative analysis of structures and functions and more qualitative exploration of molecular behaviors triggering these functions. The three main views—Canvas,

bigBox, and 3D View are connected through a set of interactions. Upon a close examination of the nuances of a calculated function, it is suitable to explore the 3D conformations in the snapshots of interest. Therefore, after clicking on a calculated function, the 3D View is set to display the corresponding snapshot.

5 REAL-TIME ANALYTICS

Exploratory analysis requires real-time selection and computation of measures to enable a fluid, uninterrupted workflow (**R5**). It poses a substantial challenge for long MD simulations consisting of tens of thousands of atoms and hundreds of thousands of snapshots. Exploratory analysis is, per definition, unpredictable, as even the analysts themselves do not initially know the objectives of their investigation. Pre-computation of all possible selections and their associated measures is therefore not feasible.

Dataflow models, by design, lend themselves to real-time analytics since their scope is clearly defined, and their associated actions, such as loading and computing, can be performed in parallel. Additionally, we introduced several measures to enable real-time interaction and analysis, including progressive loading, predictive caching, approximate computations, and selective pre-computations of very costly functions.

5.1 Progressive Loading and Computation

I/O operations are the most costly among all sMolBox types, as they have to load the positional vector data from the slow memory into the fast cache. To minimize the latency, atom positions are loaded in an iterative progressive manner, where each iteration adds twice the amount of data of the previous iteration (see Figure 7). Empirically, we found 256 snapshots uniformly distributed across the simulation to be a reasonable initial number of snapshots so that we can provide the first overview of the data in the sMolBoxes instantly. Since every calculation is being processed on a separate thread, the user can interact with the system without interruptions. The same principle applies also to the approximation of the importance function in Measure sMolBoxes. Every new iteration is processed as soon as it is available by the parent Input sMolBox. Every sMolBox also contains a progress bar (see Figure 3), indicating a ratio of dynamic snapshots that were already processed by it.

5.2 Predictive Caching

During the exploratory analysis process, there is relatively high probability that the user will also want to explore input one-step higher or lower in the structural hierarchy (i.e., whole residue for an observed atom and vice versa). For example, the analyst may inspect a residue and observe an unexpected side chain movement, which he or she subsequently wants to investigate in more detail. For this reason, we introduce the predictive caching approach based on the structure hierarchy, caching not only the selected structure, but also parent structures (i.e., atom, residue, and ultimately chain). Having the required information cached on different levels of detail allows users to freely traverse the data hierarchy without any significant latency. We applied the principles of singleton caching for the overlapping input structures, as is often the case with large selections or whole chains accessed by multiple sMolBoxes.

5.3 Approximate RMSD Computation

The main purpose of Input sMolBoxes is to represent the input for subsequent Measure sMolBoxes. In addition, Input sMolBoxes also contribute to the analysis by providing information about the temporal development of selected parts of the molecule. This approximate RMSD is computed as a central position for the input structure and indicates the center's deviation from its initial position in each snapshot. This trade-off between speed and accuracy of displayed representation should serve solely as an instant indication of current trends in the selected input data. Precise RMSD can be computed using dedicated Measure sMolBox.

5.4 Pre-Computation and Custom CSV Files

Despite our efforts, some measures are too complex to be computed on-the-fly, e.g., the detection of molecular tunnels and their properties, which may require hours of computation on a supercomputer. Therefore, we allow users to load CSV (Comma Separated Values) files containing time series generated by other dedicated applications or custom scripts. The CSV loader is designed in a simple manner: in the first line of the file it expects a list of measures to be loaded, then reads each remaining line as a snapshot index and the corresponding values for the measures. This way, our tool can be easily extended and used for various analysis tasks for which it was not designed initially.

5.5 Implementation

Implementation-wise, the application is written as a module in the CAVER Analyst visualization software [15], using the Java programming language. Therefore, it acts as a standalone, multiplatform application, mainly targeted to the use on personal computers. The rendering engine and optimised cache is developed in-house, with focus on performance and memory limitations of personal computers. The case studies, described in the following section, were performed on MacBook Pro with M1 ARM architecture and 16 GB of RAM.

6 CASE STUDIES

We performed two case studies with a senior protein engineer (10+ years of experience) using typical datasets from his daily work. Over a long-term collaboration, this domain expert was involved in discussions about the protein engineering workflow, but he was not involved in the sMolBoxes design prior to participating in the case studies. He is a co-author of the paper—a common collaboration approach for design studies [30]. For the first case study, the protein engineer revisited a previously published research based on two simulation ensembles [16]. In this study, he and his colleagues analyzed the binding of three molecules in the substrate 1-chlorohexane (CHA). They were initially placed in the solvent, to the haloalkane dehalogenase enzyme variant DhaA31, and the release of the respective dehalogenation products 1-hexanol (HAOL) and chloride ion (Cl^-) from the active site to the solvent. These processes were compared with the dynamics of the protein tunnel (i.e., opening and closing of its entrance and bottleneck). For the second case study, the protein engineer studied an unpublished unfolding process of two proteins—the wild type and its more stable mutant variant—to characterize their differences.

For the first case study, we used two simulation ensembles, each consisting of four 200 ns simulations with one snapshot per 2 ps (i.e., 400,000 snapshots in total). For the second case study, multiple parallel short simulations were run in adaptive rounds, for a total combined simulation time of 20 μs (200,000 snapshots spaced by 0.1 ns).

The focus of the case studies was twofold: 1) visual confirmation of the previous findings to validate the appropriateness of sMolBoxes for exploratory MD simulation analysis and 2) a follow-up qualitative analysis for in-depth explanations of observed events and potential hypotheses for follow-up research. The data sets were chosen because they are relatively small and therefore enable a fluid analysis even on a commodity laptop. At the same time, the simulations are sufficiently long so that watching the 3D MD simulation visualization is no longer feasible with any of the currently available tools, and scrubbing through the animation is far too imprecise. In addition, the investigation focus is very different for the two cases—analyzing the interaction between



Fig. 8. A setup of the network during the first study. The rightmost sMolBoxes show from the top: the RMSD of the residues forming the tunnel bottleneck, throughput of the tunnel bottleneck loaded from an external CSV file, and comparison of the minimal distance of three ligands (IDs 293, 294, 295) to the residue in the active site (ID 103).

structures vs. comparing conformational changes over time—to demonstrate the versatility of the sMolBoxes approach.

The studies lasted between one and two hours each. The protein engineer was asked to prepare a summary of the already performed analyses, as well as a script describing the previously performed analysis steps to be replicated using sMolBox for the validation part of the study. He was encouraged to think aloud while performing the investigation. We performed screen- and audio-recording and transcribed the audio track. The transcript was coded into five categories: 1) questions / hypotheses about the data, 2) explanation of the (traditional) workflow, 3) findings and insights, 4) positive aspects about the interface mentioned, and 5) problems with the interface or suggestions for improvement.

At the beginning of both case studies, the protein engineer tried to replicate previous analyses for **visual confirmation**. In the first case study, the first step was to measure the distances between any of the three ligands to a catalytic residue of the active site. Figure 8 shows the three distance Measure sMolBoxes, which were then combined using a minimum Aggregation sMolBox (bottom right). This way, he could quickly spot snapshots where at least one ligand was close to the active site. He particularly appreciated the color-coding of the aggregated function, which allowed him to immediately see which ligand was currently the closest one to the active site. The second confirmation step was to investigate if there is a correlation between the tunnel bottleneck and the ligands' distance to the active site. This would indicate that the presence of the ligand influenced the tunnel radius, which would be the expected behavior. By juxtaposing the minimum distance with the bottleneck radius, he could quickly confirm this expected correlation, which indicates that the presence of the ligand influences the tunnel radius (Figure 8). These analyses were very similar to their traditional workflow. However, the protein engineer found it more intuitive, mainly because he could easily exchange the input residues by directly selecting a structure in a potentially interesting region identified in the 3D view.

In the second case study, the first step of the traditional analysis pipeline has been the computation of a Markov state model to construct different conformational states based on the RMSD of the entire protein structures. Computing the RMSD for hundreds of thousands of snapshots for such large structures would violate the real-time analytics requirement of sMolBoxes (R5). Therefore, the protein engineer instead chose to compute distances from the start and end residues of the protein helices to a stable central residue of the protein. The helices and the central residue are known for this protein. Distance computations can be conducted on-the-fly, and the protein engineer confirmed that the shape of the curves resembles the RMSD curves: initially, the distances are low, and they increase over time.

For both case studies, the **qualitative follow-up investigation** yielded new insights. In the first case study, the protein engineer




Fig. 9. Sudden change in the tunnel bottleneck residue conformation near the snapshot 100,000. The top left shows the 3D View of the protein, highlighting those bottleneck residues (dark green) in a closed conformation. The top right shows the identical residues in the open conformation and the ligand (violet) approaching the tunnel entrance. The corresponding RMSD measure of the bottleneck residues is situated at the bottom.

investigated the influence of the ligands on the bottleneck residues. The bottleneck residues have been identified in prior analyses [16] and are therefore known for this protein. For the first two investigated bottleneck residues, he could quickly see a significant conformational change in longer time scale sequences of the simulations in the RMSD approximation curve. This was unexpected but highly relevant since conformational changes of bottleneck residues can impact the binding of molecules and the opening/closing of the tunnel. He confirmed this observation with a more precise RMSD function box (Figure 8 top right). By navigating to the snapshots where the conformation changed, he could observe a flipping of the two investigated bottleneck residues (Figure 9). He could confirm that one of the ligands was close to the tunnel mouth when the conformation change was registered. That indicates that the presence of the ligand induces the bottleneck re-orientation. For a more thorough follow-up analysis, he would investigate if multiple bottleneck residues open the tunnel in a concerted manner (synchronized) or if they move independently.

For the second case study, the protein engineer wanted to explore which helices initiate the unfolding process, which is not well understood yet. By juxtaposing the pair-wise distance computations between all selected helix residues and the central residue, he concluded that some residues of the cap domain probably initiate the unfolding. Other helices followed the unfolding later, but fluctuate even more in the end. Compared to Markov state models, the protein engineer found it easier to understand the time-wise mechanism behind the unfolding by observing the distances over time and short simulation sequences associated with peak distances in the 3D View. Their traditional Markov state analysis only delivers the states, but the temporal sequence gets lost.

7 DISCUSSION & FUTURE WORK

Overall, the protein engineer performing the case studies was intrigued by the potential of sMolBoxes. He particularly appreciated the rich options to perform on-the-fly analyses, such as quick comparisons between a variety of functions (R2) for any spatial selections (R1). He noted, for instance, that he would not have considered performing a systematic RMSD analysis of bottleneck or catalytic residues before. However, the approximated RMSD preview of Input sMolBoxes made him aware of several properties worth further investigation, such as an unexpected instability of the active site, which is unusual since active site residues normally do not shift very much. These discoveries were only made due to the instant computations and visualization within sMolBoxes (R5, R6).

To our surprise, our case study participant did not strictly follow the dataflow model from input to Aggregation sMolBox (R3) but rather opportunistically jumped between sMolBoxes for setting different investigation foci and the linked views (R7). Our prototype supported this behavior well. It illustrates the importance of having the option to quickly shift the focus forth and back—which would not be so easily possible in more linear analysis approaches, such as in Python notebooks.

We also did not expect the strong desire to analyze and compare the line charts directly within the Canvas and the importance of their spatial alignment. The protein engineer suggested providing more powerful methods to assess correlations—either through an easier superimposition of multiple function boxes or by a correlation aggregation that would compute the correlation of sliding time windows between two functions.

From the case studies, we could also observe that the analyst sometimes wished to perform a broader analysis comprising a larger number of structures, such as finding bottleneck, catalytic, or helix residues behaving in a similar way to hypothesize about entire functional units. As sMolBoxes operate on individual selections, such larger-scale comparisons are currently not supported. From a conceptual point of view, we could enable such a broader analysis by performing batch operations on every substructure of a selection, for instance computing the RMSD on all residues of a selected protein chain.

Furthermore, the dataflow interface currently provides information about data provenance (R4). However, in the future, we would also like to extend the system with user action provenance to provide information about any modifications of the dataflow during the analysis. We aim to use the provenance information combined with automated data analysis and event detection to train a recommendation system that could suggest potentially interesting data inputs, properties, and aggregations and guide the experts through the exploration process. We also plan to enable users to define their own sMolBoxes by combining multiple existing boxes into a tree that could be “collapsed” and stored as a single sMolBox for later use.

8 CONCLUSION

In this paper, we have presented sMolBoxes, a novel system for visual exploration of long sequences of molecular dynamics (MD) simulations. The proposed approach supports the workflow of computational chemists, enzymologists, and protein engineers and integrates quantitative analysis of molecular dynamics simulations into the visual exploratory analysis. Our work takes inspiration from dataflow diagrams and enables fast selection of molecular inputs and computation, aggregation, and visualization of their associated properties. Instead of obtaining measures for pre-defined structures through scripting, sMolBoxes provide a powerful and flexible method for the discovery and qualitative characterization of important aspects of the MD simulations.

Two case studies with a protein engineer indicate that sMolBoxes is flexible and responsive enough to support rich on-the-fly exploratory analysis to generate new hypotheses. In the studies, sMolBoxes enabled a fluid interplay between targeted MD simulation analyses using quantitative analysis with function plots and qualitative observations in a 3D visualization, even for simulations consisting of hundreds of thousands of snapshots.

ACKNOWLEDGMENTS

The authors wish to thank Sergej Stoppel, Robin Skånberg, and Mathieu Linares for their participation in the initial phases of the project. The authors would also like to express their gratitude to the Czech Ministry of Education (INBIO - CZ.02.1.01/0.0/0.0/16.026/0008451; ELIXIR - LM2018131; eINFRA - LM2018140) and the Czech Science Foundation (20-15915Y).

REFERENCES

- [1] N. Alharbi, R. S. Laramée, and M. Chavent. MolPathFinder: Interactive Multi-Dimensional Path Filtering of Molecular Dynamics Simulation Data. In C. Turkay and T. R. Wan, eds., *Computer Graphics and Visual Computing (CGVC)*. The Eurographics Association, 2016.

- [2] K. Bidmon, S. Grottel, F. Bös, J. Pleiss, and T. Ertl. Visual abstractions of solvent pathlines near protein cavities. *Computer Graphics Forum*, 27(3):935–942, 2008.
- [3] J. Brezovsky, B. Kozlíková, and J. Damborský. Computational analysis of protein tunnels and channels. In *Protein Engineering*, pp. 25–42. Springer, 2018.
- [4] A. Bryden, G. Phillips, and M. Gleicher. Automated illustration of molecular flexibility. *IEEE Transactions on Visualization and Computer Graphics*, 18(1):132–145, 2012.
- [5] J. Byška, T. Trautner, S. M. Marques, J. Damborský, B. Kozlíková, and M. Waldner. Analysis of long molecular dynamics simulations using interactive focus+ context visualization. In *Computer Graphics Forum*, vol. 38, pp. 441–453. Wiley Online Library, 2019.
- [6] J. B. Dabney and T. L. Harman. *Mastering simulink*, vol. 230. Pearson/Prentice Hall Upper Saddle River, 2004.
- [7] D. Duran, P. Hermosilla, T. Ropinski, B. Kozlíková, A. Vinacua, and P.-P. Vázquez. Visualization of large molecular trajectories. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):987–996, 2019.
- [8] N. Elmqvist, J. Stasko, and P. Tsigas. Datameadow: a visual canvas for analysis of large-scale multivariate data. *Information Visualization*, 7(1):18–33, 2008.
- [9] T. Ertl, M. Krone, S. Kesselheim, K. Scharnowski, G. Reina, and C. Holm. Visual analysis for space-time aggregation of biomolecular simulations. *Faraday Discussions*, 01 2014.
- [10] K. Furmanová, M. Jarešová, J. Byška, A. Jurčík, J. Parulek, H. Hauser, and B. Kozlíková. Interactive exploration of ligand transportation through protein tunnels. *BMC Bioinformatics*, 18 Suppl 2, 2017.
- [11] U. Hensen, T. Meyer, J. Haas, R. Rex, G. Vriend, and H. Grubmüller. Exploring protein dynamics space: the dynasome as the missing link between protein structure and function. *PLoS ONE*, 7(5):e33931, 2012.
- [12] S. A. Hollingsworth and R. O. Dror. Molecular dynamics simulation for all. *Neuron*, 99(6):1129–1143, 2018.
- [13] W. Javed and N. Elmqvist. Explates: spatializing interactive analysis to scaffold visual exploration. In *Computer Graphics Forum*, vol. 32, pp. 441–450. Wiley Online Library, 2013.
- [14] D. Jönsson, P. Steneteg, E. Sundén, R. Englund, S. Kottravel, M. Falk, A. Ynnerman, I. Hotz, and T. Ropinski. Inviwo—a visualization system with usage abstraction levels. *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [15] A. Jurčík, D. Bednář, J. Byška, S. M. Marques, K. Furmanová, L. Daniel, P. Kokkonen, J. Brezovský, O. Strnad, J. Štourač, et al. Caver Analyst 2.0: analysis and visualization of channels and tunnels in protein structures and molecular dynamics trajectories. *Bioinformatics*, 34(20):3586–3588, 2018.
- [16] S. Kaushik, S. M. Marques, P. Khirsariya, K. Paruch, L. Libichová, J. Brezovský, Z. Prokop, R. Chaloupková, and J. Damborský. Impact of the access tunnel engineering on catalysis is strictly ligand-specific. *The FEBS journal*, 285(8):1456–1476, 2018.
- [17] J. Kodosky. Labview. *Proceedings of the ACM on Programming Languages*, 4(HOPL):1–54, 2020.
- [18] P. Kokkonen, J. Sýkora, Z. Prokop, A. Ghose, D. Bednář, M. Amaro, K. Beerens, v. Bidmanová, M. Slánská, J. Brezovský, J. Damborský, and M. Hof. Molecular gating of an engineered enzyme captured in real time. *Journal of the American Chemical Society*, 140(51):17999–18008, 2018.
- [19] T. Koudeláková, R. Chaloupková, J. Brezovský, Z. Prokop, E. Šebestová, M. Hesseler, M. Khabiri, M. Plevaka, D. Kulik, I. Kutá Smatanova, P. Řezáčová, R. Ettrich, U. T. Bornscheuer, and J. Damborský. Engineering enzyme stability and resistance to an organic cosolvent by modification of residues in the access tunnel. *Angewandte Chemie International Edition*, 52(7):1959–1963, 2013.
- [20] S. M. Marques, L. Daniel, T. Buryška, Z. Prokop, J. Brezovský, and J. Damborský. Enzyme tunnels and gates as relevant targets in drug design. *Medicinal Research Reviews*, 37(5):1095–1139, 2017.
- [21] S. M. Marques, Z. Dunajová, Z. Prokop, R. Chaloupková, J. Brezovský, and J. Damborský. Catalytic cycle of haloalkane dehalogenases toward unnatural substrates explored by computational modeling. *Journal of Chemical Information and Modeling*, 57(8):1970–1989, 2017.
- [22] X. Martinez, M. Chavent, and M. Baaden. Visualizing protein structures—tools and trends. *Biochemical Society Transactions*, 48(2):499–506, 2020.
- [23] H. Mei, Y. Ma, Y. Wei, and W. Chen. The design space of construction tools for information visualization: A survey. *Journal of Visual Languages & Computing*, 44:120–132, 2018.
- [24] J. Meyer-Spradow, T. Ropinski, J. Mensmann, and K. Hinrichs. Voreen: A rapid-prototyping environment for ray-casting-based volume visualizations. *IEEE Computer Graphics and Applications*, 29(6):6–13, 2009.
- [25] R. Patro, C. Y. Ip, S. Bista, D. Thirumalai, S. S. Cho, and A. Varshney. MDMap: A system for data-driven layout and exploration of molecular dynamics simulations. In *2011 IEEE Symposium on Biological Data Visualization (BioVis)*, pp. 111–118, 2011.
- [26] R. Patro, C. Y. Ip, A. Varshney, and H. Hagen. Saliency guided summarization of molecular dynamics simulations. *Scientific Visualization: Advanced Concepts*, 1:321–335, 2010.
- [27] M. Pavlová, M. Klvaňa, R. Chaloupková, P. Banáš, M. Otyepka, R. Wade, Y. Nagata, and J. Damborský. Redesigning dehalogenase access tunnels as a strategy for degrading an anthropogenic substrate. *Nature Chemical Biology*, (5):727–733, 2009.
- [28] F. Ritter, T. Boskamp, A. Homeyer, H. Laue, M. Schwier, F. Link, and H.-O. Peitgen. Medical image analysis. *IEEE Pulse*, 2(6):60–70, 2011.
- [29] K. Schatz, J. J. Franco-Moreno, M. Schäfer, A. S. Rose, V. Ferrario, J. Pleiss, P.-P. Vázquez, T. Ertl, and M. Krone. Visual analysis of large-scale protein-ligand interaction data. In *Computer Graphics Forum*, vol. 40, pp. 394–408. Wiley Online Library, 2021.
- [30] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE transactions on visualization and computer graphics*, 18(12):2431–2440, 2012.
- [31] C. T. Silva, J. Freire, and S. P. Callahan. Provenance for visualizations: Reproducibility and beyond. *Computing in Science Engineering*, 9(5):82–89, 2007.
- [32] R. Skånberg, M. Linares, C. König, P. Norman, D. Jönsson, I. Hotz, and A. Ynnerman. VIA-MD: Visual Interactive Analysis of Molecular Dynamics. In J. Byska, M. Krone, and B. Sommer, eds., *Workshop on Molecular Graphics and Visual Analysis of Molecular Data*. The Eurographics Association, 2018. doi: 10.2312/molva.20181102
- [33] J. Someone and T. Something. VVVV a multipurpose toolkit. <https://vvvv.org/>, 2019. Accessed: 2019-03-09.
- [34] C. Upson, T. A. Faulhaber, D. Kamins, D. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz, and A. van Dam. The application visualization system: a computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4):30–42, 1989.
- [35] V. Vad, J. Byška, A. Jurčík, I. Viola, E. Gröller, H. Hauser, S. M. Marques, J. Damborský, and B. Kozlíková. Watergate: Visual exploration of water trajectories in protein dynamics. In *Proceedings of Eurographics Workshop on Visual Computing for Biology and Medicine (EG VCBM)*, pp. 33–42, 2017.
- [36] P.-P. Vázquez, P. Hermosilla, V. Guallar, J. Estrada, and A. Vinacua. Visual analysis of protein-ligand interactions. *Computer Graphics Forum*, 37(3):391–402, 2018.
- [37] J. Waser, H. Ribicic, R. Fuchs, C. Hirsch, B. Schindler, G. Bloschl, and E. Gröller. Nodes on ropes: A comprehensive data and control flow for steering ensemble simulations. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):1872–1881, 2011.
- [38] B. Yu and C. T. Silva. Visflow - web-based visualization framework for tabular data with a subset flow model. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):251–260, 2017.
- [39] B. Yu and C. T. Silva. Flowsense: A natural language interface for visual data exploration within a dataflow system. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1–11, 2019.