

Analysis of Long Molecular Dynamics Simulations Using Interactive Focus+Context Visualization

J. Byška^{†1,2}, T. Trautner^{1,3}, S. M. Marques^{4,5}, J. Damborský^{4,5}, B. Kozlíková², and M. Waldner³

¹Department of Informatics, University of Bergen, Bergen, Norway; ²Faculty of Informatics, Masaryk University, Brno, Czech Republic;

³Faculty of Informatics, TU Wien, Vienna, Austria; ⁴Loschmidt Laboratories, Masaryk University, Brno, Czech Republic;

⁵International Clinical Research Center, St. Anne's University Hospital, Brno, Czech Republic;

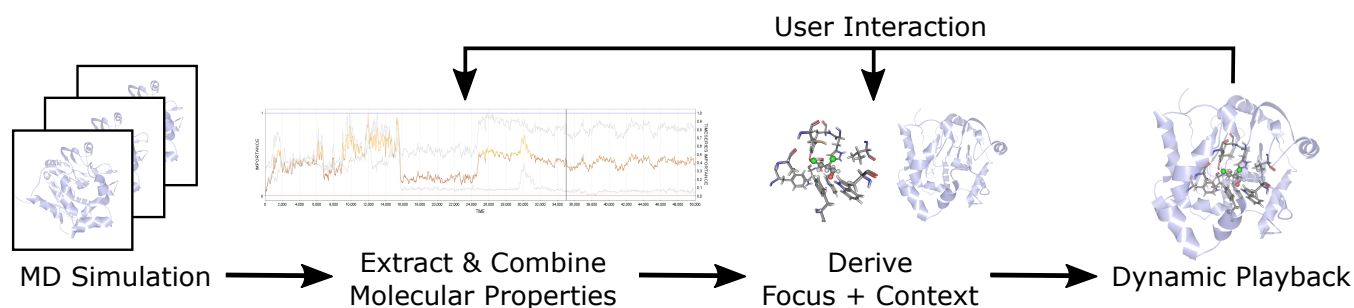


Figure 1: Workflow supported by our proposed approach. The input dataset, formed by long MD simulations, is analyzed by combining different properties in a 2D plot. From the graph, the user can extract the important events, which control the temporal and spatial focus and context 3D view by adjusting the visual representation and the animation speed.

Abstract

Analyzing molecular dynamics (MD) simulations is a key aspect to understand protein dynamics and function. With increasing computational power, it is now possible to generate very long and complex simulations, which are cumbersome to explore using traditional 3D animations of protein movements. Guided by requirements derived from multiple focus groups with protein engineering experts, we designed and developed a novel interactive visual analysis approach for long and crowded MD simulations. In this approach, we link a dynamic 3D focus+context visualization with a 2D chart of time series data to guide the detection and navigation towards important spatio-temporal events. The 3D visualization renders elements of interest in more detail and increases the temporal resolution dependent on the time series data or the spatial region of interest. In case studies with different MD simulation data sets and research questions, we found that the proposed visual analysis approach facilitates exploratory analysis to generate, confirm, or reject hypotheses about causalities. Finally, we derived design guidelines for interactive visual analysis of complex MD simulation data.

CCS Concepts

• **Human-centered computing** → **Scientific visualization**; **User centered design**;

1. Introduction

Protein dynamics and interaction with other molecules is of utmost importance for biochemists and molecular biologists. In protein engineering, the analysis of these phenomena helps the researchers to understand and design the desired function of a protein. In drug design, it can reveal crucial findings for the design of effective new

drugs with the suitable pharmacokinetic properties towards the target biomolecules [SBK*14, KCB*13]. Molecular dynamics (MD) is one of the most commonly used and powerful computational methods, enabling to simulate and capture the physical movements of molecules and their atoms. It mostly utilizes interatomic potentials or molecular mechanics force fields to calculate the forces between atoms and their potential energies. Such complex calculations require substantial computational power, creating the main limitation for capturing long MD trajectories. The time span be-

[†] jan.byska@gmail.com

tween two consecutive time steps, used in these simulations, is typically between 1 to 100 picoseconds. While small-scale interactions require such a high temporal resolution, many biological processes are happening rather in the order of milliseconds. To simulate both small-scale and large-scale processes, we therefore need hundreds of thousands or even millions of time steps. Nowadays' computing capabilities enable us to calculate such long simulations, and the next step is their thorough analysis by the domain experts, aiming to observe important behavior captured in the simulation. One of the most natural ways is to use visualization methods, accompanying the experts already for decades [Lev66].

3D animations of MD simulations have gained popularity for exploring molecular ensembles some decades ago [HDS96]. This was logical as the experts have a good understanding of the 3D representation, and nowadays the major molecular visualization tools are integrating them. However, with the increasing length and complexity of the simulations, this approach is becoming less and less feasible. From numerous conversations with the experts about their daily workflow, we are witnessing that, indeed, the relevance of 3D visualization for MDs is decreasing. Temporally-averaged or property-based clustering analysis of MDs based on different measures, such as the distance between atoms or molecules, angles, root-mean square deviation (RMSD) of protein residues or sections, electrostatic or van der Waals interaction energies, hydration, etc., are now a common practice. Since we today have the resources and tools, the use of multiple simulation replicates (sometimes hundreds) are a necessity for the sake of sampling and reproducibility of the observed events. Therefore, it is impractical to visualize all these simulations, and nowadays, the visualization is always accompanied (and sometimes even fully replaced) by statistical analysis [SS13]. A method that is gaining enormous popularity is the statistical analysis of the simulations to construct Markov states that can describe the dynamical process under study, and even allow the prediction of exchange kinetic rates between them [HP18].

However, history confirms that for proper understanding of the biological behavior, in-depth investigation of the 3D structure is crucial [WC53] and cannot be fully replaced by sole statistical analysis. In our focus groups, domain experts confirmed that 3D animations are still important as an exploratory tool, as well as for presenting their findings. However, we also found that with the increasing complexity of the MD simulations, the visual clutter – both spatially and temporally – makes it hard to focus on the phenomena under observation. To support an in-depth analysis of such cluttered structures and dynamics, we first need a deep understanding of the role of 3D animations in the domain experts' workflow to be able to design better visualizations for the visual analysis in a fluid interplay with classic statistical analysis – even for ensembles of long trajectories with a lot of elements.

In this work, we employed a design study methodology [SMM12] to improve 3D animations of long and crowded MD simulations. We first carefully analyzed the current workflow of protein engineers and derived requirements for 3D animations of MD simulations in this field. We then elicited the most suitable design choices based on these requirements and validated the prototype implementation for the analysis of different open research questions in the field of biochemical engineering. Finally, we de-

rived guidelines for the design of 3D animations of long and complex MD simulations. Following the guidelines of design studies [SMM12], our goal was not to come up with a radically novel visualization technique but rather to choose an effective visualization and interaction design in a problem-driven research approach. In summary, the contributions of this paper are:

1. Requirements for effective and efficient exploratory visual analysis of proteins and protein-ligand interactions with 3D animations based on focus groups characterizing the importance of dynamic 3D visualizations of MD, as well as the increasing challenges to use them in molecular modeling practise (Section 3).
2. The design of a spatio-temporal focus+context (F+C) technique augmenting existing 3D visualization approaches, where user-selected elements and events of interest are visualized in full spatial (elements) and temporal (events) detail, while the context remains abstracted, as well as its implementation in a visual analysis environment with smart visibility management and temporal smoothing for visual decluttering (Section 4).
3. Results from two informal evaluations, two case studies, and a small usability evaluation revealing which variants of the spatio-temporal focus+context visualization enables the domain experts to analyze MD simulations with hundred thousands time steps and ten thousands of elements on different levels of detail with varying investigation focus (Section 5).
4. Design guidelines for 3D animations of long and complex MD simulations based on the results of the evaluations (Section 6), as well as guidance for future work (Section 7).

The conceptual implications of our work are widely applicable to many other situations and it can be extended to more general problem-focused and user-defined contexts, as discussed in Section 6.

2. Related Work

Visualization and visual analysis of MD simulations have been in the scope of researchers for already many years, which is evidenced also in recent survey papers focusing on visualization of biomolecular structures [KKF*16] and multiscale molecular visualization [MKK*18]. Over the years, the focus of the domain experts has changed according to the capabilities to capture longer simulations and their ensembles. Early solutions were focusing mainly on the animation of the movements, making this a standard representation in widely used tools, such as VMD [HDS96], PyMOL [Sch15], Chimera [PGH*04], or YASARA [KV14]. This was feasible for small-scale simulations, consisting of few hundreds of time steps. Nowadays, the length of simulations spans up to hundreds of thousands, or even millions, of time steps, which is not feasible to observe in a step-by-step manner anymore. Therefore, the domain experts have to use abstract static representations, aggregating the movements to show the trends of movements or more importantly, navigate themselves only to the interesting parts of the simulations. Examples of static representations of MD simulation were proposed by Bryden et al. [BPG12], showing the main movements of protein domains by arrow glyphs, or by modulating the width of the tube of the protein backbone representation to express the flexibility of individual parts of protein chain, as introduced in the UnityMol tool [LTDS*13] and shown in [KKF*16]. Patro

et al. [PIVH10] are trying to extract events in MD simulations and create a summary overview. They are calculating the importance on each keyframe of the simulation, by measuring per-atom saliency and its change from the previous and consecutive keyframe. This computation is then performed over multiple scales of the simulation. The final result is a collection of significant keyframes that represent the MD simulation. Another approach proposed by Patro et al. [PIB*11] summarizes the MD simulations using a state-transition graph between the keyframes. Alharbi et al. [ALC16] present MolPathFinder, a tool for interactive filtering of paths of MD simulations. This tool also generates an overall representation of the paths and the user can select those path which fulfil given conditions. In these examples, the MD simulation is effectively summarized, and the 3D geometry of the most important events can be studied in detail. However, with the static representations, the ability to observe dynamic behavior is lost.

Additionally, the length of the MD simulations is not the only parameter influencing the selection of the appropriate exploration technique. The content of the simulation is another crucial factor. Some existing techniques are focusing on visual exploration of the void space inside proteins, without any interacting partner. These techniques include the work of Lindow et al. [LBBH13] where the authors study the migration of a selected inner cavity within protein over time and aggregate the occupied void space to form so called dynamic cavities. Byška et al. [BJG*15] present an abstracted representation of the evolution of the bottleneck of a selected protein tunnel over time. The concept was extended to study the whole tunnel along its centerline [BMG*16]. An overview of the existing methods for visualization and visual analysis of void space in molecules, for both static and dynamic cases, can be found in Krone et al. [KKL*16].

The input MD simulations can contain also the trajectory of ligands entering the protein inner part and travelling to the active site. In such cases, the exploration of the simulation is enhanced by studying parameters influencing the behavior of the ligand, i.e., its speed, direction, conformation changes, and many others. Furmanová et al. [FJB*17] presented a visual analysis tool for exploration of these parameters, revealing potentially interesting parts of ligand trajectories by using several interlinked views. This is useful for analyzing movements of a single ligand in the protein. However, we observed that the focus of investigation often changes to other aspects of the simulation, such as the evolution of the protein conformation, behaviors of multiple ligands, or even more complex combinations of multiple measurements, that cannot be covered by such a dedicated tool. Another visual representation of protein-ligand interactions was proposed by Vázquez et al. [VHG*18]. Their 2D compact representation of molecular simulation encodes several properties of the interaction. It also enables the comparison of two simulations side-by-side. Although the technique comprehensibly conveys several properties of the interaction, it is not suitable for the large trajectories and more ligands within one simulation. These limitations are addressed by Duran et al. [DHR*19] who designed a tool for visual exploration of more ligands at once. Each ligand trajectory is represented by one 2D plot, showing the information about the ligand speed, distance to the active site, and energies. However, their 3D animation does not contain focus+context or visibility management techniques so

that long and dense simulations easily get too cluttered for in-depth analysis.

None of these approaches can be utilized for observing the interaction between a protein and a large number of molecules, such as the solvent consisting of thousands of water molecules. In such cases, the experts are often more interested in the trends of flow of these small molecules rather than observing their behavior individually. For these purposes, Bidmon et al. [BGB*08] proposed a clustering of water trajectories and showing only the extracted principal paths of these clusters in 3D. Another visual analysis tool for studying the flow of water molecules inside a protein was proposed by Vad et al. [VBJ*17]. Their tool consists of a set of linked views, enabling the exploration and filtering of trajectories, based on user-defined criteria and thresholds. It integrates the clustering of trajectories, proposed by Góra et al. [MMG*17]. However, our case studies revealed that experts sometimes need to investigate small-scale interactions between a single water molecule and a ligand in more detail. By only abstracting the overall flow, these detailed interactions cannot be revealed. One of the most similar techniques to our approach is the work of Skånberg et al. [SLK*18]. Their VIA-MD tool serves for visual interactive analysis of MDs where they are aiming to find interesting patterns. Except for the possibility to plot a set of different descriptive functions, they visualize the density field depicting the spatial distribution of molecules over time. However, our goal is rather to create an integrated descriptive function which then controls the 3D animation of the simulation.

Although we evidenced that there are already several interesting approaches for exploration of MD simulations, which can contain the ligand or solvent trajectories, each of them is focusing on one specific investigation (e.g., only ligand trajectories) in one specific spatial and temporal scale (e.g., the flow of water in the protein). Our focus groups revealed that complex biological phenomena usually cannot be explained from a single investigation target and on a single spatial and temporal level. It is rather a complex interplay between elements and spatio-temporal scales that can explain certain behavior. We therefore seek to provide a flexible approach to support these varying investigation foci in a single interactive analysis environment.

3. Focus Groups

Over a period of three years, we conducted four focus groups with a group of protein engineers and researchers at a local protein engineering research group. The focus groups were attended by three to four domain experts and two to three visualization experts each. All focus groups were either audio recorded or transcribed on-the-fly. On average, each focus group lasted around 90 minutes.

The major topic of the focus groups was the characterization of the current workflow, the role of 3D visualizations for MD simulation data analysis, and the identification of important elements and events in the MD simulations. For each focus group, we prepared a set of questions for guiding the group. In addition to the discussions within the focus group, the domain experts also provided us with written descriptions of their current workflow steps, together with screenshots as follow-up information.

3.1. Summary of Results

The focus groups provided us with deep insights into multiple aspects of the domain experts' work. All our domain experts are focusing on protein engineering. Ultimately, their goals are to design mutations, which lead to changes of protein properties, such as its activity or stability. This is important, for instance, to accelerate transfer of toxic compounds into non-toxic compounds by increasing the interaction speed with a selected ligand.

Current workflows: MD simulation is one of the essential methods to study dynamic protein behavior. Through the MD simulations, the protein engineers can study, for instance, how tunnels form over time and how ligands move through these tunnels to and from the protein active site. Any MD simulation data is not merely visually inspected, but evaluated statistically as well. For statistical evaluation, different measurements are derived from an ensemble of simulations, such as distances or interaction energies over time. The space of potential measurements is enormous and it is not feasible to analyze all possible measurements. Therefore, the analysts need to have at least some idea where to start with the analysis. A typical first step is to plot a time-dependent distance measurement to identify first temporal regions of interest, as shown in Figure 2.




Figure 2: A 2D plot created by one of our domain experts as the starting point for his analysis. It shows the distance of the alcohol (green) and chloride Cl^- (blue) products to the buried active site of haloalkane dehalogenase DhaA (PDB entry 4E46) during an MD simulation. The protein boundary is approximated and depicted by the horizontal dashed lines. The running average of the distances over 0.5 ns for the alcohol and Cl^- are colored according to the location with respect to the protein boundary.

In general, our experts stated that for their tasks, elements, such as ligands, amino acids, or water molecules, are important when they influence the protein tunnel. Often, the experts therefore analyze the ligand movement as a first measurement to get more event indicators. Subsequently, they use different techniques to identify which protein regions the ligand is predominantly attending and where most interaction energy takes place.

The role of visualization: Apart from the 2D plots, shown in Figure 2, the prevalence of visualization in the experts' workflows is surprisingly low. One expert even stated that “3D visualizations are meaningless” and another one repeatedly stressed the point that 3D visualization is not an important tool for them any more; their research relies on statistics (“We trust in data, not visualization!”).

They explained that with nowadays' longer and more detailed simulations, “you lose detail, it is hard to know what you want to observe” because “it is too messy, it is not possible to focus on something”. The visual clutter results from frame skipping (showing only every 10th to 25th frame that is performed to keep the length of the dynamic visualization within a reasonable limit, to be able to observe the overall behavior of the protein. However, for observing specific smaller-scale events, a higher temporal resolution would be required. For instance, to observe some interactions between ions and the protein or the fast conformational changes of residues with small side chains, a 10 picoseconds resolution is not sufficient.

Nevertheless, the experts use tools like VMD [HDS96] or PyMol [Sch15] to occasionally consult 3D representations. In the course of the four focus groups, the users mentioned different purposes why they use 3D visualizations:

- to perform qualitative observations to formulate a hypothesis, which is then verified through classic statistical analysis,
- to estimate which elements' measurements to include in the statistical analysis to verify a hypothesis (e.g., whether a particular residue is likely to be responsible for tunnel opening or whether the analysis should focus on larger groups of residues),
- to better imagine the raw data (e.g., how much is 5 Å? “We have to see it sometimes”),
- to qualitatively explain the mechanisms leading to statistically observed differences, and
- for presentation, such as movie sequences for conferences or for teaching.

In general, the purpose is always to observe behavior of the protein or interactions between elements, i.e., dynamic phenomena that are hard to grasp from raw numbers or 2D plots. Even in a complex study with multiple MDs, the visualization of a representative trajectory can provide information that can be generalized over the entire set of simulations. The choice of the trajectory to be visualized depends on previous 2D data from a preliminary analysis performed by the domain experts or on the tracking by third-parties (e.g., by a detected transition between Markov states) of the events of interest.

However, the experts never watch a 3D visualization of the MD simulation from the beginning to the end. They rather use time series plots, as the one in Figure 2, to identify potentially interesting time steps and then use VMD's VCR (“Videocassette recorder”) controls and the timeline to navigate themselves to these time steps. Then they frequently switch between these two visualizations. In case of very crowded simulations, their current approaches cannot handle the dynamic visualization of the simulation data at all.

3.2. Requirements for 3D Animations of MD Simulations in Protein Engineering

From the focus groups, it became clear that traditional visualization techniques with uniform temporal resolution and standard rendering capabilities, as provided by tools like VMD [HDS96] or PyMol [Sch15], are not sufficient any more when dealing with increasingly long and complex simulations. Yet, our domain experts

still use them as the supportive instrument in various situations, despite their current shortcomings, when technically possible. Therefore, we believe that the role of 3D animations can be strengthened again by compensating for the increased visual clutter and better adjusting the visualization interface to the users' workflows. From our observations, we derived the following requirements:

R1: Link 3D animations with time-dependent MD simulation measurements to support efficient navigation to events of potential importance.

R2: Allow flexible exchange and combination of these measurements as the focus of the investigation progresses.

R3: Adapt the temporal resolution of the 3D animation to the mechanisms under observation.

R4: Maximize the visibility of the structures of interest even in very crowded and dynamic scenes while abstracting the context.

There are solutions that already address some of these requirements. For instance, Duran et al. [DHR*19] link 3D animations of protein-ligand interactions with 2D plots of associated MD simulation measurements (R1). However, the 3D animation is not adapted to the exploration focus (R3, R4) and therefore hard to comprehend for long and crowded simulations. In addition, our goal was to create a general solution for protein engineers that is not limited towards a specific application, such as protein-ligand interaction. By providing a tool addressing all these requirements, the protein engineers' workflow can be supported across a range of data sets and exploration tasks.

4. Spatio-Temporal Focus+Context for 3D Animations

Based on the above requirements, we designed and developed a novel interactive visual analysis system for MD simulations with the goal to better integrate 3D animations of MD simulation data into a typical biochemical analysis workflow. In an iterative design process, we elicited, discussed, and extended potentially useful techniques – not only from the molecular visualization field, but also from other visualization areas. The prototype was integrated into CAVER Analyst [JBB*18] as it already provides a lot of MD measurements that are crucial to fulfill requirements R1 and R2.

4.1. Spatial Focus+Context

Perception studies show that humans need focused attention to detect dynamic changes of an object [ROC97]. That means, in order to fully understand the dynamic behavior, the user's attention should be focused on a single element at a time. To satisfy requirement R4, we therefore need focus+context visualization [Hau06]. Classic focus+context techniques apply spatial distortions to the visualization, generating a fisheye or magnifier glass effect around the focus region [Fur86, WZMK05]. We are more interested in the simplifications of context elements, while showing the focus in full detail [CDF*06]. This principle has been applied to molecular visualization, where secondary structures of proteins can be selectively abstracted to generate a focus+context effect [VDZLBI11]. The authors also present how to smoothly interpolate between different representations to create a fluid transition between focus and context areas.

However, our domain experts disliked the intermediate abstrac-

tion stages and preferred to see well-known representations, such as cartoon or balls-and-sticks. These representations have distinct semantic meanings and are used to either observe secondary structures or atoms. We therefore render the focus elements in a known detailed representation and context elements in another, less detailed one. This is a standard feature in many MD visualization tools. However, protein engineers are especially interested in the interaction between the elements in focus and their surroundings. Therefore, we allow them to select two foci, each with a different representation. The choice of representation, of course, depends on the elements in focus and context. For example, for the analysis of ligand movement, our default representations are balls-and-sticks for the ligand, sticks for the surrounding amino acids (i.e., 5 Å from the ligand's center of gravity), and the cartoon representation for the rest of the protein, as shown in Figure 3. Nevertheless, we allow the users to change the representation for both foci and the context at any time. The biochemists can choose from *structure wireframe*, *solvent-excluded surface*, *cartoon*, or *alpha-trace* representation for the context and *sticks*, *balls-and-sticks*, *van der Waals surface*, or *solvent-excluded surface* representation for each focus. Our prototype also allows them to set the focus on water molecules traveling to or from the active site during the simulation, derived from [VBJ*17].




Figure 3: Three different levels of abstraction for focus and context: The whole protein (i.e., context) is rendered in cartoon representation, the main focus element (i.e., ligand) is rendered using balls-and-sticks, and its surroundings (i.e., nearby amino acids) are shown as sticks. The image shows the same closeup view with ghosting turned off (left) and on (right).

As our domain experts do not wish to have intermediate abstraction stages, there is no smooth spatial transition between the focus and context areas. As we have elements that gain and lose focus over time, such as amino acids getting in contact with the ligand, these elements may only appear or disappear for a few frames as the ligand is moving. To avoid flickering due to these fast movements, we were experimenting with gating signals with a signaling window of up to a second. In other words, we ignored the focus-context changes for elements losing or gaining focus for such a short time window. Nevertheless, we found out that such a solution can hide important short-time events or keep the elements in focus for too long, depending on the signaling window used. The second and more successful solution we implemented and tested with our domain experts was motion blur. The amount of the motion blur is based on the playback speed to suppress the extreme flickering but it is completely removed on the slow playback to allow easy tracking of the focus elements. The advantage of the motion blur

in comparison with the gating signals is that it does not completely remove the short-time events but only suppresses them.

To keep the user oriented in space, we additionally need visibility management to ensure that the focus remains visible across different viewpoints and changing protein conformations (see Figure 3). Classic visibility techniques in 3D are cutaways and ghostings [FS92]. In a comparative study, Baer et al. [BGCP11] found that ghosting techniques lead to more accurate depth perception and are preferred over simple alpha blending. In the field of molecular visualization, Le Muzic et al. [LMMS*16] presented a novel interface to selectively “cut away” macromolecules in a dense molecular visualization to reveal the embedded structures. However, according to our domain experts, it is “*better seeing more residues than too few*” in order not to miss relevant details. We therefore decided to maximize the visibility of the focus elements with ghosting rather than cutaways to minimize removal of potentially important context information.

4.2. Temporal Focus+Context

To fulfill requirement **R3**, we need not only to put a visual emphasis on the spatial structures, but also on events which may happen in different times scales. While focus+context in static visualizations is a well-known concept, changing the temporal resolution of a dynamic visualization is not very commonly used. As a notable exception, Wolter et al. [WAH*09] present a visualization of nasal airflow, where they first show an overview of the simulation in a coarse temporal resolution. The user can then select intervals of interest to generate visualizations with a higher temporal resolution so see more detailed dynamics. However, according to our domain experts, it is challenging to identify the potentially relevant events in the first place – especially when they are very short. Therefore, we need a mechanism to automatically adapt the temporal resolution of the animation to the events of interest. We designed and implemented two possible solutions for temporal focus+context, *event-driven* (Section 4.2.1) and *navigation-driven* (Section 4.2.2), and we allow users to freely switch between them.

4.2.1. Event-Driven Temporal Focus+Context

In the field of multimedia, *adaptive fast-forward* [CLCC09, HHWH11] has been introduced as means to express *temporal* focus and context information. It is defined as non-linear time mapping of a video to an animated visualization time. The mapping is adapted to the relevance of the content, such as the amount of motion, the visual complexity of the scene, or the commentator voice for sports events [DO07]. The more relevant the content, the slower the animation speed. We can apply the same concept to 3D animations: the more likely a relevant event is taking place, the slower the animation speed. A major challenge thereby is to detect important events.

As we learned in our focus groups, the users’ interest in different temporal regions changes over time as they investigate different MD simulation measurements (cf., requirement **R2**). The focus of investigation can be then described by a large variety of time-dependent measurements that are investigated to detect

events. Based on our domain experts’ research goals and their descriptions of potentially useful measurements, we included the following categories of time-dependent measurements: ligand-related (e.g., speed or distance from the active site), protein-related (e.g., RMSD), and tunnel-related (e.g., tunnel bottleneck). These measurements are computed for each simulation frame and linearly mapped to an importance range, spanning from 0 to 1. The importance values are then linearly mapped to the playback speed that is defined by a minimum and maximum number of simulation frames that are visualized per second (SFPS) and can be adjusted by the user at any time. By default, the minimum playback speed is 1 SFPS and the maximum playback speed is 20 SFPS. To ensure smooth playback, the users can optionally turn on Hermite interpolation between the simulation frames on small SFPS. By default, we are skipping simulation frames on high values (e.g., >60 SFPS) as they would not be all visualized anyway, due to the monitor refresh rates.

To communicate this information to the user, we utilize a time series visualization (see Figure 4). In this 2D chart, the time is shown on the horizontal axis and the importance values on the vertical axis. To preserve the information about the underlying measurements, we annotated the chart with two vertical axes. The left axis denotes the normalized importance that is used to steer the temporal resolution during the playback, while the real measurement values are shown on the right axis. The 2D chart is equipped with a vertical slider, denoting the currently visualized frame in the 3D view. The chart thereby allows the users to easily navigate to a part of the simulation expressing a specific behavior (e.g., a high concentration of water), by clicking on the corresponding time in the chart. The time series visualization therefore acts like a “scented widget” extension [WHA07] of the classic time slider, thereby fulfilling the requirement **R1**.

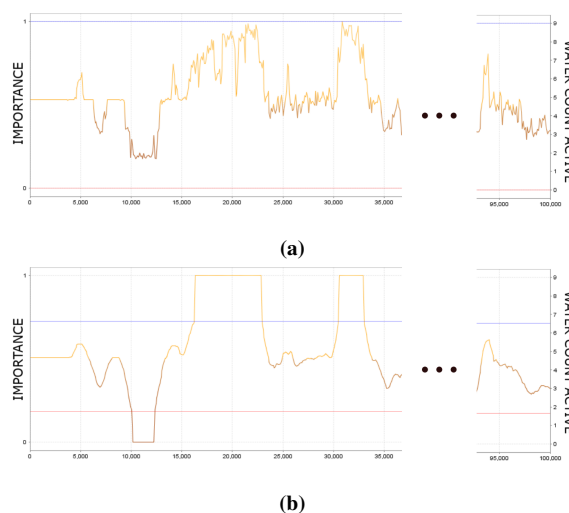


Figure 4: The 2D time series view showing the number of water molecules in the protein active site in the MD simulation consisting of 100,000 time steps without (a) and with (b) smoothing and active thresholds. The blue and red lines denote the upper and bottom threshold value, respectively.

As it is common to analyze simulations consisting of hundreds of thousands of time steps, we aggregate the neighbouring values in the time series until we reach the target resolution, derived from the width of the visualization window. We then depict only the average values of the aggregated points. Moreover, as the relevant physico-chemical measurements are usually noisy, the importance value and thus the playback speed can change frequently (see Figure 4a). Therefore, we also allow the users to *smooth* the importance function. We use the moving average technique with a window size derived from the length of the simulation. We plan to expose the window size parameter to the interface in the future.

During the feedback sessions, it became evident that the initially mentioned measurements are usually not sufficient to identify potentially interesting events without further modification. For instance, the classic ligand speed measurement can be interesting from various aspects. On the one hand, low speed can indicate that the ligand got “stuck” and was interacting with the surrounding amino acids, which is a potentially interesting state to locate amino acids that strongly influence the protein tunnel. On the other hand, high speed of the ligand is an indication that the tunnel is wide open, which is often a desired protein state to be studied. However, high ligand speed is irrelevant if the ligand is outside the protein, floating in the solvent. In other cases (e.g., the release of ligand from its stable configuration), an event is defined by a ligand changing its speed, i.e., the derivative of the speed measurement. We therefore added several operations to modify the MD simulation measurements on the fly: *invert* ($f'(t) = 1 - f(t)$), *derive* ($f'(t) = (f(t+1) - f(t-1))/2$), and *combine*.

We provide four options to combine multiple MD measurements: *max*, *min*, *average*, and *normalized sum*. The combination is always performed on the scaled importance range (i.e., 0 – 1) rather than on the original measurement values, to avoid skewing the results towards more prominent measurements. The *max* combination picks the maximum of any MD measurement for each frame, and the *min* combination picks the minimum value (Figure 5). The *average* balances the values of all measurements for each frame. Finally, the *normalized sum* sums up all the normalized measurements and scales the result back to the 0 – 1 range. To support easy modification and combination of a large number of MD simulation measurements, the time series visualization always highlights the currently modified measurement and shows its value range on the right axis (see Figure 4).

We additionally allow the users to set upper and lower *thresholds* by dragging the maximum and minimum lines in the time series window. All values above the upper and below the lower threshold lines are then mapped to the importance 1 and 0, respectively (see Figure 4b). To enable the users to set the thresholds more precisely, we also support zooming across both the time and importance axes.

These operations allow the user to express more complex investigation foci. For example, when observing a simulation with multiple ligands, the expert can combine the distance to the active site for each ligand using the *max* combination. In this way, the importance is high whenever *any* ligand gets to the vicinity of the active site. Similarly, the expert can find out if a ligand is moving fast

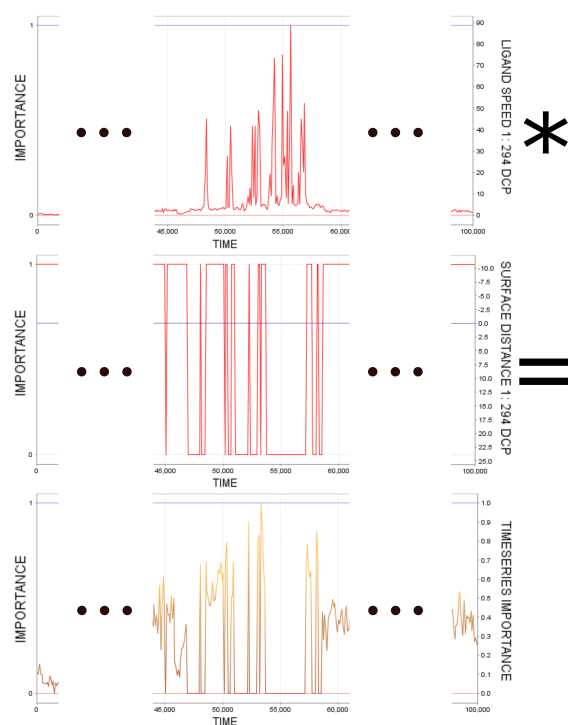


Figure 5: The min combination of the ligand speed (top) and its thresholded distance (middle) to the protein surface suppressing unimportant parts of the simulation where the ligand was not inside the protein. It is obvious that the high peaks from the original function representing the ligand speed (top) were removed, allowing the before insignificant changes to the ligand speed inside the protein to become more prominent.

within a protein by combining the ligand speed and ligand distance to the surface using the *min* combination (see Figure 5).

4.2.2. Navigation-Driven Temporal Focus+Context

A disadvantage of the event-driven F+C is that domain experts often initially cannot explicitly formulate interesting events based on the above mentioned measurements. They stated that they first need a quick overview of the protein dynamics to formulate a hypothesis. We therefore designed and implemented an alternative approach to the event-based F+C for untargeted exploratory analysis, based on the user’s spatial navigation.

The various phenomena observed by the experts take place on different temporal and spatial levels. It is well known that the timing of biological processes scales with size [SS13]. For example, the interaction between a protein and a ligand takes place on a different spatial and temporal scale than the interaction on an atomic level. Hence, we couple the temporal resolution of the 3D animation to the zoom factor in the visualization. To achieve this effect, we evaluate the camera distance from every element in the spatial focus, such as a ligand, and linearly map this distance to the minimum and maximum speed. We cap the minimum and maximum playback speed at the distances of 10 Å and 100 Å, respectively.

With this approach, the users can reduce the animation speed by zooming towards the element in focus. Therefore, small-scale events can be observed in high temporal resolution. The more the user zooms out, the more frames are skipped. This way, the users can obtain an overview of the behaviors happening on a larger spatial scale, such as the conformation changes of the entire protein or the behavior of ligands in the solvent.

5. Evaluations

In total, we performed two informal evaluations with a group of domain experts, as well as two case studies and one usability evaluation with a single expert to investigate which aspects of our approach are usable in practice. For the two informal evaluations, we used shorter simulations of one protein, the haloalkane dehalogenase DhaA (UniProt ID P0A3G2), with 2,162 frames and 1,000 frames, respectively. The first simulation contained a single ligand and was used by the experts to analyze amino acids influencing the main protein tunnel. The second simulation did not contain any ligand but was used to analyze the interaction between water molecules and the protein [MDP* 17]. For the two case studies, we used longer simulations of the same protein with 50,000 and 100,000 frames. In the first data set, there are three ligands concurrently interacting with the protein, as well as more than 11,000 water molecules. In the second data set, there are two different ligands and the same amount of water molecules.

For the informal evaluation and the case studies, we investigated only event-driven F+C. For all simulations, we set the minimum playback speed to 1 SFPS, which is the highest possible temporal resolution. For the two short simulations, the maximum speed-up was set to 5 SFPS. For the longer simulations, the maximum speed-up was initially set to 100 SFPS. We empirically determined these maximum values before the evaluations so that the context frames can be observed within a reasonable time, depending on the absolute number of simulation frames. For the usability evaluation, we systematically varied the temporal F+C approach.

5.1. Informal Evaluations

The informal evaluations were important to gather early feedback for the in-depth case studies to adjust the prototype accordingly. We demonstrated the initial prototype, asked three domain experts about what they would be interested in to observe, and gathered their feedback. The informal evaluations lasted around 30 minutes.

For the ligand-protein simulation, we set the spatial focus on the ligand and its surrounding contacts. As the temporal importance measurement, we used the ligand speed. The protein engineers particularly approved ghosting for maximizing the visibility of the focus elements. The experts also expressed interest to be able to load a large amount of additional measurements and to combine them arbitrarily. They urged that fast switching between measurements would be a strict requirement for the practicability of the tool. However, domain experts did not comment on the non-linear temporal resolution of the dynamic 3D visualization. Clearly, with a few thousand frames covering the entire simulation, the benefit was not really evident in this case.

The second simulation was also quite short, but it was much

more crowded with more than 11,000 water molecules (see Figure 6a). We set the spatial context to the water molecules within the protein (Figure 6b) and used the number of waters in the protein as the temporal importance measurement. This means the temporal resolution would be low when only few waters are within the protein. Compared to the protein-ligand use case, the experts had much less knowledge about the data so they found it hard to define how they would start the investigation. For this use case, the domain experts showed a lot of excitement, as they would now be able to observe if residues exhibit unexpected movements if a lot of water molecules flow into the protein, which is not possible without having control over the investigation focus due to visual clutter.




Figure 6: Visualization of all water molecules present in the simulation of DhaA (PDB entry 4E46) (a) and our proposed focus+context visualization (b) depicting only water molecules inside the protein.

5.2. Case Studies

In these case studies, a post-doc protein engineering expert was asked to investigate one of his current simulation data. Using this data, he had been investigating the binding and release cases of a ligand. However, he still was not able to fully explain what exactly triggers the binding and release of the ligands (if anything), the relationship between the ligand binding and the tunnel opening and closing, and which residues are most important for determining the changes in the tunnel. This data has not been investigated using 3D visualization so far as it was too long and crowded for their current tools, such as VMD [HDS96] or PyMol [Sch15], without focus+context support.

To support his analyses, we pre-computed several time-dependent MD measurements for the ligands (e.g. speed, distance to active site), the protein conformation (e.g. RMSD), the water molecules (e.g. the number of water molecules in the active site), and the tunnel (e.g., bottleneck size, throughput) to ensure fluid switching between the measurements. The focus could be set on all ligands with or without contacts or surroundings, as well as the water molecules. The expert was asked to investigate his open research questions with our prototype implementation. He was encouraged to think aloud and ask for assistance whenever necessary. On average, one case study lasted 90 minutes. We recorded the screen, transcribed the audio, and analyzed the transcripts. After the second case study, the expert wanted to continue his investigation over

a longer period of time. We therefore provided him with an executable of our prototype implementation. A full report of his findings from this individual analysis of the data sets can be found in the supplemental material.

The first investigated data set contained three 1,2,3-trichloropropane (TCP) substrate ligands moving to the active site of the haloalkane dehalogenase DhaA protein [MDP*17]. The expert started by observing the ligands and their speed, as well as their surroundings. The first unexpected observation was how far away the ligands are from the protein for a very long part of the simulation. However, this observation has no relevance for his work. After getting the initial overview, he wanted to investigate what keeps the ligand bound. He therefore used the inverted speed as the temporal importance function. Thereby, he could identify events where the ligand is close to the tunnel bottleneck. For instance, he observed that residue 149 changed its conformation (see Figure 7) and as the residue bent to the back the ligand was more mobile. He acknowledged that identifying this would be hard in a “plain MD simulation” and that this could be a starting point for a statistical verification. By stepping through multiple events, where the ligand got stuck, he identified two further residues that were always present around the stuck ligand.




Figure 7: Residue 149 of DhaA (PDB entry 4E46) (left) moves to the back (right) when the ligand (TCP substrate) enters the tunnel.

As he shifted his analysis focus towards the tunnel bottleneck, the 2D plots became more important. For instance, he checked the correlation between the ligand distances to the active site and the tunnel cost by superimposing the measurements (see Figure 8). As expected, the tunnel cost is mostly low with a ligand inside. However, he spotted one time period without any ligand and the tunnel still being open. He hypothesized that the tunnel opening could have been induced by waters in this case. Navigating to that time period confirmed that as the tunnel opened, there were many more water molecules inside the protein, and the number dropped again as the tunnel closed and “waters are flushing” out.

In his report from the individual analysis, the expert additionally reports on a particular and unexpected finding: By setting the temporal focus on the water concentration to track the water molecules inside the protein, he could observe that, if there is a ligand inside the protein, water molecules can still access the active site through a “backside tunnel”. So far, such tunnels were not considered relevant because they do not have a high throughput. Yet, this helps to understand how waters, which are highly relevant for the catalytic hydrolysis of TCP into 2,3-dichloropropan-1-ol (DCP), access the protein.

In the second case study [MDP*17], the expert was investigating

a data set simulating the release of DCP and chloride (Cl^-) products from the active site of the DhaA protein. As for the first data set, he does not yet fully understand what triggers the release of the ligands, especially in the presence of waters.

The first aspect the expert planned to observe was whether the hydration of chloride increased before its release. With the MD measurement set to the ligand’s distance to the active site, he found it interesting to observe that chloride is interacting a lot with the hydroxyl group of DCP, while they travel together to the tunnel mouth. Through careful observation, he spotted one water molecule that was interacting with the chloride over a longer period of time. He hypothesized that this water molecule “is going to drag the ligand out. But let’s see.” He continued watching the simulation until the chloride was released from the protein and the temporal resolution of the visualization was decreased again. He concluded that his hypothesis was not confirmed. “Which I mean that the water [molecule] was not responsible for the ligand to leave.” During his individual analysis, the expert could spot additional interesting interactions between ligands and waters, such as waters interrupting the interaction between DCP and Cl^- , which led to hypotheses for further investigations.

In general, it was often difficult for the domain expert to express his investigation focus by selecting the focus elements and a measurement or combination of multiple measurements even though he could express his interest on a high level. For instance, it was difficult to express that the temporal importance should be high whenever any ligand moves slowly within the protein. While this is possible to express for one ligand using the *min* operator (see Figure 5), it requires a combination of *min* and *max* for multiple ligands, which was not supported by our system. Especially when the temporal focus was not optimally defined, it was often unclear for the expert why the simulation was speeding up or slowing down. Overall though, he commented that he liked the fact that he could now watch longer simulation sequences without having to compromise details due to uniform frame skipping, which is the case in their current standard approach.

When interacting with the system, we observed that the user at no point was trying to observe very long parts of the simulation or even watch the simulation as a single movie from the beginning to the end. He rather used the 2D visualizations as navigation aid, explicitly looking for peaks and watching the simulation until the event of interest was over. This implies that linking the 3D animation with the 2D plots supports their current practice, where they use static plots to find interesting simulation frames (see Section 3). This also implies that temporal F+C does not fully eliminate the need to manually navigate in the 3D animation. One particularly useful feature was to show the derivative of the MD measurements to track major changes in the investigated property. One unexpected behavior was that the user tried to select elements from the 3D visualization to obtain MD measurements related to that particular element, such as an amino acid. Conversely, he expected that the visualization would explicitly highlight regions responsible for a slow-down, such as the ligand with the lowest speed, to quickly see not only that something important is happening but also where it is happening.

A highly appreciated aspect was the flexible handling of spatial




Figure 8: The superimposed 2D plots of the three ligands (gray – distance to active site) and tunnel cost (red): from time step 16,000 to $\sim 30,000$, there is at least one ligand close to the active site and the tunnel cost is low (left blue box). It increases as the ligands leave again (time steps 30,000 to 32,000) and drops again afterwards, when two ligands move back towards the active site (right blue box). At the beginning of the simulation (time step 5,000 to 12,000, yellow box), the tunnel cost is low even though there is no ligand close to the active site, which could be caused by the presence of water molecules.

focus selections by our system. The expert changed the spatial focus several times – the ligands only, ligands with their contacts or larger surroundings, as well as water molecules inside the protein alone or together with the ligands. He thereby sometimes adjusted the representation of the different types of focus elements, such as van der Waals radii for water molecules together with ligands represented as sticks, or ligands represented as balls-and-sticks together with the surrounding residues as sticks only, depending on his investigation focus.

5.3. Usability Evaluation

As the informal evaluations and case studies did not provide sufficient evidence initially to conclude whether temporal F+C is useful and which settings are the most appropriate ones, we finally performed a usability evaluation with the domain expert. We presented two different events to the expert using the simulation from the second case study: 1) the release of the DCP ligand and 2) water concentration changes. These events were presented to the user with three different settings: 1) manual speed control through a menu, 2) event-driven F+C defined by respective temporal importance functions (ligand distance to the active site and the derivative of the water count inside the protein, respectively), and 3) navigation-driven F+C. His task was to interactively explore and observe the 3D animation, describe the content, and give feedback on the appropriateness of the temporal resolution.

Ligand release: Using the **manual speed** control, the expert was able to observe the chloride moving from the tunnel to the outside, while always being surrounded by water molecules. He commented that the speed was a bit too fast. However, he chose not to change it through the menu as this would distract him from the animation. He suggested to control the speed through the arrow keys instead. Using **event-driven F+C**, he observed a lot of interaction of both, the ligand and the chloride, with the surrounding water inside of the tunnel. He commented that the slower playback before the ligand was released helped him to observe these more detailed interactions. However, he found it crucial but very challenging to define the correct function for the temporal resolution. Finally, using the **navigation-driven F+C**, he did not report any additional observations. On the one hand, he liked the concept as it is easily possible to change the speed without attending the menu. On the other hand,

he commented that it was hard to set an appropriate minimum and maximum speed and speculated that a logarithmic mapping of distance to speed could be more intuitive.

Water concentration: Using the **event-driven F+C**, the expert was able to observe movements of helices concurrent with water concentration changes and hypothesized that these movements could be the cause for the water concentration increase. Using the **navigation-driven F+C**, he could only observe the increase of water molecules, but was not able to detect any causes. He commented that it was not easy to spot the actual events using this method. Finally, we increased the maximum speed of the **event-driven F+C** from 20 to 1,000 SFPS so that the entire simulation could be watched within a few minutes, while slowing down when the water concentration changes. Using this setting, he was able to observe large movements of the helices, which “*feels like the water molecules are pushing [some of these helices] a bit when they change.*”

6. Discussion and Design Recommendations

By enhancing conventional 3D animations with a spatio-temporal focus+context technique, we were able to visualize crowded environments with complex interactions that our users had not been able to investigate in a 3D animation using their standard tools. The case studies revealed that 3D animations have the potential to identify possible targets for further statistical analysis, to perform qualitative observations that either confirm or contradict a hypothesis formed on the basis of the statistical data, and to generate hypotheses about causalities. Based on the informal feedback and the case studies, we revisit our design decisions to provide design recommendations and hints to interesting future research questions.

For our use cases, it is clear that they cannot be investigated without powerful spatial focus management. Visibility management through ghosting is thereby a highly appreciated aspect. Our domain expert of the case studies also liked the ability to easily select multiple categories of focus elements and control their representations. For animated 3D visualizations of MD simulations, we therefore recommend to **allow the users to easily select different focus elements, let them assign detailed standard representations, and maximize the visibility of these elements.**

From the focus groups and the case studies it is clear that the users need different temporal resolutions, depending on their current investigation focus. The usability study revealed that the navigation-driven F+C or manual control of the temporal resolution can be initially more intuitive. However, it does not lower the risk of missing important events that are relevant for the current investigation focus. Our observations show that the event-driven F+C has the potential to guide the user's attention to simulation time steps that are of high relevance. However, speed changes caused by the non-uniform temporal resolution were also often found confusing. Highlighting elements that are responsible for an event and a more apparent visualization of the current speed than just the motion blur, such as suggested by Höferlin et al. [HKH*12], could make the speed changes less disruptive. We therefore recommend to **consider to dynamically adjust the temporal resolution to guide the user's attention in complex 3D animations. However, also allow the user to easily adjust the speed and provide clear visual feedback.**

Linking 2D plots of MD measurements and the 3D visualization was a key concept for detecting events and finding entry points for the 3D animation. To our surprise, the domain experts very thoroughly specified events by manipulating single measurements and combining them into compound measurements to express their current focus of investigation. Unfortunately, the possibilities provided by our prototype were limiting them in their expressiveness, especially if multiple potential focus elements were present in the simulation. To satisfy all observed information needs, a far more sophisticated interface to design compound measurements will be necessary. It therefore seems that **dynamic 3D visualizations should be linked to (multiple) 2D time series visualizations of associated MD measurements to facilitate detection of and navigation to relevant events.**

It is important to stress that the usage of the event-driven F+C concepts can be generalized to many other research fields dealing with dynamics simulations, such as drug design, structural biology, or material sciences. With additional measures, such as angle, dihedral angle, contacts, hydration, or interaction energy, complex studies beyond the field of protein-ligand interaction, such as protein flexibility or stability under certain conditions or the transport of molecules through the channels in the membrane proteins, can be conducted. It will require more advanced interfaces so that users can compute their measures of interest and flexibly combine them. Once the user has found the optimal importance function that can properly describe their focus of study, such function can be exported to be used in batch analysis of a virtually unlimited number of simulations.

7. Future Work

We have demonstrated the usefulness of our spatio-temporal F+C concept for visual analysis of single MD simulations. We believe that this concept is a first step so that dynamic 3D visualizations will start to gain more importance in the scientific work again. To fully reach this goal, we foresee several important next steps based on our first experiences:

First, we will require more versatile and smart interfaces for the

user to express the high-level investigation focus. With a large number of additional MD measures, such as interaction energies or RMSD of various residues or groups of residues, it will become even more challenging to express the investigation focus.

Second, with a more mature and flexible interface, it will also be possible to let domain experts – even beyond the field of protein engineering – conduct longer-term visual analyses in their own environment. Using field logging and follow-up interviews, their analysis steps can be used to further improve the interface and to identify missing features.

Third, MD simulations are usually computed in parallel, and any kind of statistical analysis is conducted on an ensemble of simulations. While protein engineers consult 3D animations of single, selected MD simulations with the current workflow, integrated analysis of time series data and 3D animations has the potential to let them selectively access 3D animations of a simulation ensemble. This leads to a significantly increased computational demand when flexibly combining measurements (**R2**) and accessing the linked 3D animations on demand (**R1**). In this regard, our approach lends itself to be coupled with in-situ visualization, where the acquisition of the simulation frames can be coupled to the current investigation focus.

8. Conclusions

Using spatio-temporal focus+context visualization, we were able to visualize long and crowded MD simulation data that our collaborators were not able to visualize in 3D with their standard tools. We could show that a clear spatial focus and non-uniform temporal resolution has the potential to guide the user's attention to relevant events in different spatial regions and in different temporal granularities.

Our goal was to create an application-independent solution to analyze MD simulations. Using our prototype implementation, we could demonstrate the usefulness for two applications within the field of protein engineering. Clearly, the focus definitions are application-specific. Therefore, for other use cases than protein-ligand interaction and protein-water interaction, appropriate focus elements and MD measurements have to be identified. Conceptually, our approach can be applied to any simulation data, where time series data can be extracted to specify events, and spatial selections can be made to define a spatial focus. We do believe that our approach is generalizable not only to these other cases but even beyond the field of simulations, for instance to enable efficient storytelling in 3D animations.

Acknowledgments

The presented work has been supported by the Czech Science Foundation international project GC18-18647J and by the MŠMT projects no. LM2015055, LQ1605, LM2015047, LM2015042, LM2015085 as well as the Austrian Science Fund (FWF): T 752-N30. The Computational resources were provided by the CESNET (LM2015042) and CERIT Scientific Cloud (LM2015085). Parts of this work have been done in the context of the VIDJ project, which is supported by Bergens Forskningsstiftelse, the Mohn Medical Imaging and Visualization Center (MMIV), and the University of Bergen.

References

- [ALC16] ALHARBI N., LARAMEE R. S., CHAVENT M.: Mol-PathFinder: Interactive Multi-Dimensional Path Filtering of Molecular Dynamics Simulation Data. In *Computer Graphics and Visual Computing (CGVC)* (2016), Turkey C., Wan T. R., (Eds.), The Eurographics Association. 3
- [BGB*08] BIDMON K., GROTTTEL S., BÖS F., PLEISS J., ERTL T.: Visual abstractions of solvent pathlines near protein cavities. *Computer Graphics Forum* 27, 3 (2008), 935–942. 3
- [BGCP11] BAER A., GASTEIGER R., CUNNINGHAM D., PREIM B.: Perceptual evaluation of ghosted view techniques for the exploration of vascular structures and embedded flow. *Computer Graphics Forum* 30, 3 (2011), 811–820. 6
- [BJG*15] BYŠKA J., JURČÍK A., GRÖLLER M. E., VIOLA I., KOZLÍKOVÁ B.: MoleCollar and Tunnel Heat Map Visualizations for Conveying Spatio-Temporo-Chemical Properties Across and Along Protein Voids. *Computer Graphics Forum* 3, 34 (2015), 1–10. 3
- [BMG*16] BYŠKA J., MUZIC M. L., GRÖLLER M. E., VIOLA I., KOZLÍKOVÁ B.: AnimoAminoMiner: Exploration of protein tunnels and their properties in molecular dynamics. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 747–756. 3
- [BPG12] BRYDEN A., PHILLIPS G., GLEICHER M.: Automated illustration of molecular flexibility. *IEEE Transactions on Visualization and Computer Graphics* 18, 1 (2012), 132–145. 2
- [CDF*06] COLE F., DECARLO D., FINKELSTEIN A., KIN K., MORLEY K., SANTELLA A.: Directing Gaze in 3D Models with Stylized Focus. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques* (2006), EGSR'06, Eurographics Association, pp. 377–387. 5
- [CLCC09] CHENG K.-Y., LUO S.-J., CHEN B.-Y., CHU H.-H.: Smart-Player: user-centric video fast-forwarding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2009), ACM, pp. 789–798. 6
- [DHR*19] DURAN D., HERMOSILLA P., ROPINSKI T., KOZLÍKOVÁ B., VINACUA A., VÁZQUEZ P.-P.: Visualization of large molecular trajectories. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 987–996. 3, 5
- [DO07] DIVAKARAN A., OTSUKA I.: A video-browsing-enhanced personal video recorder. In *14th International Conference of Image Analysis and Processing-Workshops* (2007), IEEE, pp. 137–142. 6
- [FJB*17] FURMANOVÁ K., JAREŠOVÁ M., BYŠKA J., JURČÍK A., PARULEK J., HAUSER H., KOZLÍKOVÁ B.: Interactive exploration of ligand transportation through protein tunnels. *BMC Bioinformatics* 18 Suppl 2 (2017). 3
- [FS92] FEINER S. K., SELIGMANN D. D.: Cutaways and ghosting: satisfying visibility constraints in dynamic 3D illustrations. *The Visual Computer* 8, 5-6 (1992), 292–302. 6
- [Fur86] FURNAS G.: Generalized fisheye views. *SIGCHI Bull.* 17, 4 (1986), 16–23. 5
- [Hau06] HAUSER H.: Generalizing Focus+Context Visualization. In *Scientific Visualization: The Visual Extraction of Knowledge from Data*, Mathematics and Visualization. Springer Berlin Heidelberg, 2006, pp. 305–327. 5
- [HDS96] HUMPHREY W., DALKE A., SCHULTEN K.: VMD: visual molecular dynamics. *Journal of Molecular Graphics* 14, 1 (1996), 33–38. 2, 4, 8
- [HHWH11] HÖFERLIN B., HÖFERLIN M., WEISKOPF D., HEIDEMANN G.: Information-based adaptive fast-forward for visual surveillance. *Multimedia Tools and Applications* 55, 1 (2011), 127–150. 6
- [HKH*12] HÖFERLIN M., KÜRZHALS K., HÖFERLIN B., HEIDEMANN G., WEISKOPF D.: Evaluation of fast-forward video visualization. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2095–2103. 11
- [HP18] HUSIC B. E., PANDE V. S.: Markov state models: From an art to a science. *Journal of the American Chemical Society* 140, 7 (2018), 2386–2396. PMID: 29323881. 2
- [JBB*18] JURČÍK A., BEDNÁŘ D., BYŠKA J., MARQUES S. M., FURMANOVÁ K., DANIEL L., KOKKONEN P., BREZOVSKÝ J., STRNAD O., ŠTOURAC J., PAVELKA A., MAŇÁK M., DAMBORSKÝ J., KOZLÍKOVÁ B.: CAVER Analyst 2.0: analysis and visualization of channels and tunnels in protein structures and molecular dynamics trajectories. *Bioinformatics* 34, 20 (2018), 3586–3588. 5
- [KCB*13] KOUDELÁKOVÁ T., CHALOUPKOVÁ R., BREZOVSKÝ J., PROKOP Z., ŠEBESTOVÁ E., HESSELER M., KHABIRI M., PLEVAKA M., KULIK D., KUTÁ SMATANOVÁ I., ŘEZÁČOVÁ P., ETTRICH R., BORNSCHEUER U. T., DAMBORSKÝ J.: Engineering enzyme stability and resistance to an organic cosolvent by modification of residues in the access tunnel. *Angewandte Chemie International Edition* 52, 7 (2013), 1959–1963. 1
- [KKF*16] KOZLÍKOVÁ B., KRONE M., FALK M., LINDOW N., BAAEDEN M., BAUM D., VIOLA I., PARULEK J., HEGE H.-C.: Visualization of biomolecular structures: State of the art revisited. *Computer Graphics Forum* 36, 8 (2016), 178–204. 2
- [KKL*16] KRONE M., KOZLÍKOVÁ B., LINDOW N., BAAEDEN M., BAUM D., PARULEK J., HEGE H.-C., VIOLA I.: Visual Analysis of Biomolecular Cavities: State of the Art. *Computer Graphics Forum* 35, 3 (2016), 527–551. 3
- [KV14] KRIEGER E., VRIEND G.: YASARA View—molecular graphics for all devices—from smartphones to workstations. *Bioinformatics* 30, 20 (2014), 2981–2982. 2
- [LBBH13] LINDOW N., BAUM D., BONDAR A. N., HEGE H. C.: Exploring cavity dynamics in biomolecular systems. *BMC Bioinformatics* 14 Suppl 19 (2013), S5. 3
- [Lev66] LEVINTHAL C.: Molecular model-building by computer. *Scientific American* 214, 6 (1966), 42–52. 2
- [LMMS*16] LE MUZIC M., MINDEK P., SORGER J., AUTIN L., GOODSSELL D. S., VIOLA I.: Visibility equalizer cutaway visualization of mesoscopic biological models. *Computer Graphics Forum* 35, 3 (2016), 161–170. 6
- [LTDS*13] LV Z., TEK A., DA SILVA F., EMPEREUR-MOT C., CHAVENT M., BAAEDEN M.: Game on, science - how video game technology may help biologists tackle visualization challenges. *PLoS ONE* 8, 3 (2013), e57990. 2
- [MDP*17] MARQUES S. M., DUNAJOVÁ Z., PROKOP Z., CHALOUPKOVÁ R., BREZOVSKÝ J., DAMBORSKÝ J.: Catalytic Cycle of Haloalkane Dehalogenases Toward Unnatural Substrates Explored by Computational Modeling. *Journal of Chemical Information and Modeling* 57, 8 (2017), 1970–1989. 8, 9
- [MKK*18] MIAO H., KLEIN T., KOUŘIL D., MINDEK P., SCHATZ K., GRÖLLER M. E., KOZLÍKOVÁ B., ISENBERG T., VIOLA I.: Multiscale molecular visualization. *Journal of Molecular Biology* (2018). In press. 2
- [MMG*17] MAGDZIARZ T., MITUSIŃSKA K., GOŁDOWSKA S., PŁUCIENNIK A., STOLARCZYK M., ŁUGOWSKA M., GÓRA A.: AQUADUCT: a ligands tracking tool. *Bioinformatics* 33, 13 (2017), 2045–2046. 3
- [PGH*04] PETERSEN E. F., GODDARD T. D., HUANG C. C., COUCH G. S., GREENBLATT D. M., MENG E. C., FERRIN T. E.: UCSF Chimera—a visualization system for exploratory research and analysis. *Journal of Computational Chemistry* 25, 13 (2004), 1605–1612. 2
- [PIB*11] PATRO R., IP C. Y., BISTA S., THIRUMALAI D., CHO S. S., VARSHNEY A.: MDMaP: A system for data-driven layout and exploration of molecular dynamics simulations. In *2011 IEEE Symposium on Biological Data Visualization (BioVis)*. (2011), pp. 111–118. 3
- [PIVH10] PATRO R., IP C. Y., VARSHNEY A., HAGEN H.: Saliency guided summarization of molecular dynamics simulations. *Scientific Visualization: Advanced Concepts I* (2010), 321–335. 3

- [ROC97] RENSINK R. A., O'REGAN J. K., CLARK J. J.: To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science* 8, 5 (1997), 368–373. 5
- [SBK*14] SÝKORA J., BREZOVSKÝ J., KOUDELÁKOVÁ T., LAHODA M., FOŘTOVÁ A., CHERNOVETS T., CHALOUPKOVÁ R., ŠTEPÁNKOVÁ V., PROKOP Z., SMATANOVA I. K., HOF M., DAMBORSKÝ J.: Dynamics and hydration explain failed functional transformation in dehalogenase design. *Nature Chemical Biology* 10, 6 (2014), 428–430. 1
- [Sch15] SCHRÖDINGER, LLC: The PyMOL molecular graphics system, version 1.8. November 2015. 2, 4, 8
- [SLK*18] SKÅNBERG R., LINARES M., KÖNIG C., NORMAN P., JÖNSSON D., HOTZ I., YNNERMAN A.: VIA-MD: Visual Interactive Analysis of Molecular Dynamics. In *Workshop on Molecular Graphics and Visual Analysis of Molecular Data* (2018), The Eurographics Association. 3
- [SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 12 (2012), 2431–2440. 2
- [SS13] SECRIER M., SCHNEIDER R.: Visualizing time-related data in biology, a review. *Briefings in Bioinformatics* 15, 5 (2013), 771–782. 2, 7
- [VBJ*17] VAD V., BYŠKA J., JURČÍK A., VIOLA I., GRÖLLER E., HAUSER H., MARQUES S. M., DAMBORSKÝ J., KOZLÍKOVÁ B.: Watergate: Visual exploration of water trajectories in protein dynamics. In *Proceedings of Eurographics Workshop on Visual Computing for Biology and Medicine (EG VCBM)* (2017), pp. 33–42. 3, 5
- [VDZLBI11] VAN DER ZWAN M., LUEKS W., BEKKER H., ISENBERG T.: Illustrative molecular visualization with continuous abstraction. *Computer Graphics Forum* 30, 3 (2011), 683–690. 5
- [VHG*18] VÁZQUEZ P.-P., HERMOSILLA P., GUALLAR V., ESTRADA J., VINACUA A.: Visual analysis of protein-ligand interactions. *Computer Graphics Forum* 37, 3 (2018), 391–402. 3
- [WAH*09] WOLTER M., ASSENMACHER I., HENTSCHEL B., SCHIRSKI M., KUHLEN T.: A time model for time-varying visualization. *Computer Graphics Forum* 28, 6 (2009), 1561–1571. 6
- [WC53] WATSON J. D., CRICK F. H.: Molecular structure of nucleic acids: a structure for deoxyribose nucleic acid. *Nature* 171 (1953), 737–738. 2
- [WHA07] WILLETT W., HEER J., AGRAWALA M.: Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1129–1136. 6
- [WZMK05] WANG L., ZHAO Y., MÜLLER K., KAUFMAN A.: The magic volume lens: an interactive focus+context technique for volume rendering. In *Visualization, 2005. VIS 05. IEEE* (2005), pp. 367–374. 5