

Scale-Space Splatting: Reforming Spacetime for Cross-Scale Exploration of Integral Measures in Molecular Dynamics

Juraj Pálenik, Jan Byška, Stefan Bruckner, and Helwig Hauser

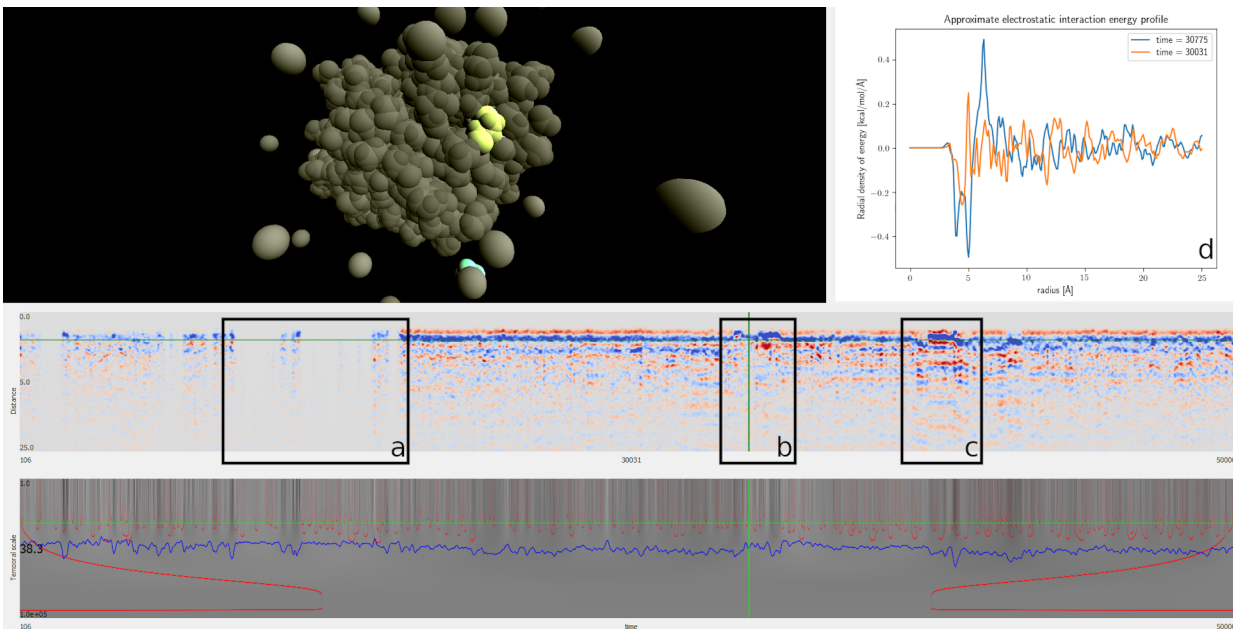


Fig. 1: Novel representation of electrostatic interaction energy (in the middle) helps navigating the simulation by spatially resolving contributions to the integral measure traditionally represented by a timeseries (blue line in the bottom). It captures different modes of the simulation better than the timeseries and helps localizing changes in the configuration of the simulation by providing spatial context. The multiscale behaviour of the data is explored by seamlessly navigating the spatiotemporal scale-space. The selected timestep shows the ligand escaping from the tunnel, before it is sucked back in by the protein.

Abstract—Understanding large amounts of spatiotemporal data from particle-based simulations, such as molecular dynamics, often relies on the computation and analysis of aggregate measures. These, however, by virtue of aggregation, hide structural information about the space/time localization of the studied phenomena. This leads to degenerate cases where the measures fail to capture distinct behaviour. In order to drill into these aggregate values, we propose a multi-scale visual exploration technique. Our novel representation, based on partial domain aggregation, enables the construction of a continuous scale-space for discrete datasets and the simultaneous exploration of scales in both space and time. We link these two scale-spaces in a scale-space space-time cube and model linked views as orthogonal slices through this cube, thus enabling the rapid identification of spatio-temporal patterns at multiple scales. To demonstrate the effectiveness of our approach, we showcase an advanced exploration of a protein–ligand simulation.

Index Terms—Scale space, time-series, scientific simulation, multi-scale analysis, space-time cube, molecular dynamics

1 INTRODUCTION

The study of molecular dynamics (MD) simulations involves computation of descriptive integral measures such as energy, temperature, pressure, density, etc. These measures are, traditionally, computed over the length of the simulation and analysed as timeseries. The purpose of

such aggregations is to simplify the large amounts of data, providing an overview of the simulation, yet it comes at the cost of hiding possibly relevant structure. This often results in degenerate timeseries that are not able to discriminate different modes in the simulation.

In order to enhance the analysis of particle simulations, we propose a novel representation that captures spatial information of the aggregate measures. Moreover, as the object of study is usually a large-scale dataset, the spatiotemporal dependency of the considered phenomena should be explored and analysed at multiple temporal and spatial scales.

For example, the electrostatic energy captures information about favourable and unfavourable states by summing contributions from spatially distributed electric charge. However, all the information about the molecule's configuration in the spatial domain is lost by performing the spatial aggregation. Furthermore, the interaction energy time-series captures various scales such as Brownian oscillation, amino-acid motion, and high-level changes, which all are relevant for different types of analysis tasks.

- Juraj Pálenik is with University of Bergen. E-mail: juraj.palenik@uib.no.
- Jan Byška is with University of Bergen and Masaryk University. E-mail: jan.byska@gmail.com.
- Stefan Bruckner is with University of Bergen. E-mail: stefan.bruckner@uib.no.
- Helwig Hauser is with University of Bergen. E-mail: helwig.hauser@uib.no.

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

Traditionally, this complexity in the data is poorly supported by the current MD analysis programs. Advanced computational tools such as NAMD [37] support evaluation of the integral measures at predefined spatial and temporal resolutions, in practice however, the visual analysis [13, 22] defaults to time series representations, where the values are reconstructed at a single temporal scale, without regard for the spatial dimensions. Conversely, scale-space approaches are well-established for investigating the behaviour of continuous functions across space and time [11, 21, 31, 32]. They provide valuable means for exploring global trends as well as local changes, and for treating noise in the data.

In this paper, we present a novel approach to the simultaneous exploration of temporal and spatial scales of integral measures in MD data. In analogy to the projection of reconstruction kernels for splatting-based volume rendering [53], we propose to project the Gaussian kernel used for scale-space construction into a reformed coordinate space that captures the space-time evolution of an integral measure. This enables us to provide a set of simple yet effective visual analysis mechanisms to explore the correspondences between the measure at different scales and its footprint in space and time.

The main contributions of this work can be summarized as follows. We introduce a new approach to the visual analysis of trajectory data that specifically takes into account its multi-scale nature (in both time and space) and present scale-space splatting as a novel computational method for reconstructing a reformed space of integral measures. Furthermore, we show how this reformed space can provide insight into complex multi-scale phenomena, such as ligand–protein interaction, and demonstrate that it enables the visual identification of relevant features in real-world data.

2 RELATED WORK

This paper contributes to the analysis of data that usually would be treated as a time series of spatially aggregated values. A good overview of visualization techniques for time-dependent data can be found in a book by Aigner et al. [1]. They focus primarily on time-dependent data in general, without specifically addressing spatial localization or scale-space approaches. Lee et al. [28] proposed an analysis based on multivariate trend identification in time dependent data without spatial dependence. Wegenkittl et al. [51] explore the use of parallel coordinates for the visualization of trajectories in multi-dimensional dynamical systems.

Our approach is based on the concept of the space-time cube, coined by Hägerstrand [20] in 1970. Since then it was repeatedly used, either explicitly or as an underlying concept. Recently, Bach et al. [3] presented a useful overview of related techniques. They describe the theoretical concept of a generalized space-time cube together with a taxonomy of all elementary space-time operations and their combinations that can be performed on such a space-time cube. A more abstract approach to time-dependent volume data analysis has been proposed by Woodring et al. [55], based on hyperplane slicing of a four-dimensional space-time hypercube. Other examples of slicing higher-dimensional data (not necessarily time dependent) include Sliceplorer [47], HyperSlice [48], Hypersliceplorer [46] and HyperMoVal [38].

2.1 Scale-Space Construction

The idea of a scale-space representation was first suggested by Witkin [54]. He proposed an approach to constructing a scale space by continuous blurring of the original time series with a Gaussian kernel of increasing size. In his work, the scale-space is segmented by finding the zero-crossings of the second derivative, and an interval tree is constructed and visualized as a tessellation of the scale-space.

In the image processing field, a well-known generalization of the scale-space construction has been proposed by Perona and Malik [36], who slackened the uniformity constraint and built the scale-space using anisotropic diffusion. Although their designated use-case was edge detection in image data, the method itself has been successfully applied to time-series scale-space construction as well [30]. A spatiotemporal extension of the scale-space approach has been presented by Laptev [27]. He extended the space for 2D time-dependent data and

his work has been successfully applied to feature detection in videos. For more advanced scale-space solutions, including an n -dimensional time-dependent scalar field scale-space, we refer the reader to Tony Lindeberg’s book “Scale-Space Theory in Computer Vision” and his more recent works [31, 32]. The respective works on a spatiotemporal scale-space construction by Laptev and Lindeberg are unfortunately not directly applicable to particle data.

Meijers and Oosterom [34] described a scale-space construction for polygonal maps, which they named *space-scale cube*. They describe a hierarchical level-of-detail method for geographical chart construction, but are not concerned with temporal data.

More scale-space techniques can be found in an overview by Holmström [21]. Vuollo et al. [49] present the construction of a scale-space for spherical data. In their work the data populates the surface of a sphere which is topologically different from our reformed coordinates.

2.2 Scale-Space Applications in Visualization

Scale-space approaches have been successfully applied in the field of visualization. For example, a scale-space surface extraction for 3D density fields was proposed by Kindlmann et al. [24]. They utilize a precise scale interpolation technique for the detection of creases by an energy-function driven particle sampling. They construct the scale-space for a continuous, static 3D field and sample the scale-space with particles as discrete agents, but they do not construct a scale-space for particle data.

Klein and Ertl [25] describe the scale-space tracking of critical points in 3D vector fields. Their technique is based on 3D continuous static data. Miao et al. [35] address the multi-scale characteristics of large molecular simulations (an HIV virus modelled at the atomic level). The authors are, however, concerned with the multi-scale nature of the spatial domain only and do not address the temporal aspect of the simulation. Bremer et al. [5] proposed a method for spatial scale selection in grid-based simulations, founded on Morse theory. In their work, a tree decomposition of the spatial domain is constructed, and a visual exploration of threshold values is enabled as opposed to a priori value guessing. The method is based on a grid representation and thus not directly applicable to particle simulations.

Pinus, by Sips et al. [42], is a multiscale visual analytics technique for finding patterns in time series data. In this work, a tree representation is constructed by computing hierarchical aggregates of the time series. While this approach is applicable to time series data, their paper is not concerned with the spatial aspect.

2.3 Space Reformations

Space reformation techniques are valuable tools for gaining insight into the spatial structure of the data. Recent publications present methods for a decomposition using a space-filling curve in MotionRugs [7] and Dynamic Volume Lines [52]. These works transform continuous volume data into one-dimensional representations by following a space-filling curve through the volume. As this kind of space reformation suffers from “delocalization”, where points close to each other can end up far apart in the representation, Dafner et al. [14] tried to overcome this limitation by context-based space-filling curves. A space-filling curve is not applicable to our case as it does not provide the desired aggregation required for the statistical treatment of particle data.

Several works address space reformation techniques for volume data, of which we highlight two: A general approach to volume transformation by Chen et al. [12], via the use of spatial transfer functions for 3D volume warping, and a curved planar reformation by Kanitsar et al. [23], that was used in medical visualization. For more examples, Kreiser et al. [26] provide a survey of flattening based techniques in medical visualization.

2.4 Particle Simulation Data

Substantial work has been done on the visualization of trajectory-based simulation data, including projects such as OVITO [45], Trillion Particles [8], Multiscale HIV [35]. These are primarily concerned with the sheer volume of the data and its direct visualization, rather than the computation of derived properties and their temporal analysis.

Focusing on the trajectories, hierarchical particle grouping for large datasets has been done by Fraedrich et al. [18]. Schirski et al. [39] extract prominent trajectories from large particle data. Kottravel et al. present a specialized tool for exploring Monte Carlo simulations of photo-voltaic cells.

Lichtenberg et al. [29] present a visual analytics tool for exploring molecular structures based on the Solvent Accessible Surface (SAS). This is a geometric approach to the extraction of the spatial configuration of a protein–ligand interaction. MoleCollar and Tunnel Heat Map [9] are works by Byška et al. where they reform the properties of a protein tunnel into the tunnel's centre line and the tunnel's cross section. In AnimoAminoMiner [10], the authors focus on amino acids lining the tunnel and their temporal development.

A visual analytics tool for long molecular dynamics simulation data, described by Duran et al. [15], provides an importance-driven time series analysis of integral measures of particle simulation data. They rely on time series aggregation and interval clustering based on ligand's position. This approach is not concerned with resolving the spatial dependency of the integral measures. VIA-MD [43] uses histograms to capture the temporal structure of complex spatio-temporal data from molecular dynamics and volume rendering to show the spatial probability distribution.

Our approach is complementary to those above and focuses on resolving the spatial dependency of the integral measures for the cases when the timeseries representation fails to capture the relevant changes.

3 ON SCALE-SPACE REFORMATION

We see that large-scale and long-time particle simulations are common in a variety of fields, especially in molecular dynamics (MD). The exploration of both the spatial configuration and aggregate measures is of importance to domain experts. The traditional representation of an aggregate measure as a time series, used for its simplifying character, cannot convey any spatial configuration. Furthermore, it exhibits multiscale behaviour that is usually not explicitly addressed. Scale-space approaches are used in continuous field analysis for the treatment of multiple scales in the data, however, they are expensive to compute for large datasets and they are not easily combined with spatial reformations that could bring out the most important spatial dependencies. Enabling an efficient computation of scale-space in reformed coordinates can significantly speed up the algorithms and allow for simplified visualizations exploiting dimensionality reduction. In this section, we address these issues and describe the construction of a scale-spaces in reformed coordinates.

The canonical scale-space is defined as the evolution of a diffusion equation over a continuous field [32]. The standard way to recover a continuous field from particle data (for example in Smoothed Particle Hydrodynamics) is by kernel density estimation [19]. In this approach, each particle is replaced by a kernel function, centred at the particle position, and the contributions from all the particles are summed up into a single scalar field. The size of the reconstruction kernel affects the resulting function. The bigger the kernel used, the smoother the resulting function becomes, with less detailed features.

One can obtain a scale-space of the data by evolving the reconstructed field with the diffusion equation. The solution to the diffusion equation can be obtained by convolving the field with a Gaussian kernel. If the reconstruction kernel is also a Gaussian, the same scale space can be obtained by a series of reconstructions with increasing kernel size. This gives a straightforward and easy to grasp concept of a scale-space for particle data. By gradually increasing the size of the reconstruction kernel we achieve the same results as by using the convolution on the reconstructed field.

Adapting this to time-dependent data means stepping up to four-dimensional space-time and evolving the diffusion equation in time, as well. As Lindeberg et al. [32] explain, the size of the kernel in the spatial dimensions does not behave the same as the one along the temporal dimension. In their work on video data, they managed to bind the two scale dimensions which enabled a joint extrema analysis and an automatic identification of space-time points of interest. Finding such a formula is application dependent and restricts the flexibility of the

analysis. Therefore, we treat the two scales separately, ending up with two independent scale dimensions, which we address by the means of visual exploration.

Our scale-space construction leaves us with a six-dimensional continuous scalar field with three spatial dimensions, a single time dimension and two scale dimensions. High dimensional slicing techniques could be utilized to explore such data, for example Sliceplorer [47] or Hyperslice [48]. However, constructing this six-dimensional space means engaging with an enormous analytical challenge, while it is not always necessary. We find that often the explored phenomena are distance-based and/or come with a spatial symmetry that can be exploited to reduce the spatial complexity of the system.

The analysis with respect to reference structures leads us to three standard transformations of coordinates, known from calculus. In the first approximation, these reference structures are a point, a line and a plane. For example, we might want to study a molecular dynamics simulation with respect to the position of a ligand interacting with the main molecular structure (e.g; a protein). Alternatively, a probe in a plasma physics scenario is usually formed by a straight wire, immersed in ionized matter, which can be considered as a line of interest/reference. Further, the distance from the piston in an engine combustion simulation can define a reference plane of interest.

In order to analyse these examples, one would simplify their spatial dependency by projecting the particle positions onto reformed coordinates. Since computing the scale-space in three dimensions in high resolution is a challenge, especially for large datasets, our approach aims at recognizing useful symmetries and realizing an according lower-dimensional scale-space in the reformed coordinates. This amounts to projecting the Gaussian kernel onto the reformed axes and using the projected version to convolve the data in the lower dimensional space.

3.1 Scale-space basics

A scale-space is usually a higher dimensional space constructed from an original field by promoting a reconstruction-filter parameter to a new continuous dimension. In the original idea of a scale-space described by Witkin [54], a one-dimensional input function $f(t)$ is filtered with a Gaussian blurring kernel with increasing kernel size, creating a two-dimensional scale-space “image” $G(t, \tau)$:

$$G(t, \tau) = \int_{-\infty}^{\infty} f(u) \frac{1}{\sqrt{2\pi\tau^2}} \exp\left[-\frac{(t-u)^2}{2\tau^2}\right] du \quad (1)$$

with τ being the kernel size. Note that, throughout the paper, we use σ for the spatial kernel size and τ for the temporal kernel size.

In order to detect and trace important features across scales, level sets of scale-space derivatives can be used. Originally, the zero crossings of the second derivative were proposed to identify peaks and valleys:

$$\frac{\partial^2 G(t, \tau)}{\partial t^2} = 0 \quad (2)$$

In the case of a one-dimensional scale-space this leads to one dimension for time and one dimension for the scale. The zero crossings of the second derivative then yield one-dimensional contours, tracing the time location of the inflexion point over the scales.

3.2 Definitions

In this paper, we treat MD simulation as trajectory-based particle data together with a descriptive integral measure. By an integral measure, we understand any function that can be computed from the trajectory data for each time-frame as a sum over all particles (atoms). Examples are: energy, potential, temperature, pressure, density, and more.

The particles “live” in the spatial domain, denoted $\Omega \subseteq \mathbb{R}^3$. By considering each particle in its own copy of Ω , we get a Cartesian product of N copies representing a spatial configuration of N particles: Ω^N . We denote the integral measure as a function $F : \Omega^N \rightarrow \mathbb{R}$, defined as:

$$F(\vec{x}_1(t), \dots, \vec{x}_N(t)) = \sum_{i=1}^N \mathcal{F}_i(\vec{x}_i(t)) \quad (3)$$

where Ω^N is the configuration space, \mathcal{F}_i is a function evaluated per particle with position $\vec{x}_i(t)$ at the time t and i is indexing the particles. The first requirement is, therefore, that the integral measure is computed as a sum of per particle contributions.

The traditional representation of the integral measure, by a timeseries $f(t)$, is obtained by summing all the contributions for each timestep.

$$f(t) = F(\omega(t)), \quad (4)$$

where $\omega(t) \in \Omega^N$ is a particular spatial configuration of the particle system at time t . For example, evaluating the energy for each timestep yields the an energy timeseries.

A continuous representation of the integral measure capturing the spatial information can be obtained from particle data by kernel density estimation (KDE) [40, 41, 50]. The KDE with kernel size σ yields a smooth scalar valued function $\Phi_\sigma(t, \vec{x}) : \mathbb{R} \times \Omega \rightarrow \mathbb{R}$

$$\Phi_\sigma(t, \vec{x}) = \sum_{i=1}^N \mathcal{F}_i(\vec{p}_i(t)) g_\sigma(\vec{x} - \vec{p}_i(t)) \quad (5)$$

where g_σ is the kernel function with kernel width σ and $\mathcal{F}_i(\vec{p}_i)$ is the contribution from the particle at the position \vec{p}_i . In case of the energy, this construction yields the energy density function. The use of Gaussian kernel provides us with two important properties described below.

First, the convolution of a Gaussian kernel with the KDE yields a reconstruction with a bigger kernel:

$$\Phi_{\sigma+\tau}(t, \vec{x}) = g_\tau(\vec{x}) * \Phi_\sigma(t, \vec{x}) \quad (6)$$

This allows us to avoid the a priori specification of a single reconstruction scale and instead call the family of KDE reconstructions a scale-space with a positive, continuous scale parameter $\sigma \in \mathbb{R}^+$.

Second, integrating the KDE $\Phi(t, \vec{x})$ over the domain Ω yields the time series f :

$$f(t) = \int_\Omega \Phi_\sigma(t, \vec{x}) d\vec{x} \quad (7)$$

Since all the $\mathcal{F}_i(\vec{p}_i)$ contributions are constant at a fixed time t , we just need to integrate over the normalized kernel. The kernel always integrates to one, resulting in the sum over all particle contributions, which is the timeseries f . This expedient behaviour with respect to integration motivates the following section.

3.3 Space Reformation

Our main goal is to understand the relevant dependencies of the integral measure F on the spatial configuration of the particles. Since the function F may describe different kinds of interaction between the particles (long-range interactions, short-range interactions, collective behaviour, etc.), different spatial structures will emerge.

An initial attempt at exploring the spatial dependency would be through the kernel density estimate $\Phi_\sigma(t, \vec{x})$. This function could be reconstructed at multiple scales and explored by direct visualization techniques, using, for example, efficient SPH rendering [17] or implicit Gaussian surfaces [6]. The full spatial reconstruction of this field, however, imposes big challenges on the analysis, where this is not necessarily required to capture the relevant spatial dependency of the studied integral measure. Many measures can, for example, be meaningfully studied as distance-based, where only the relative distance to a reference structure is important.

If we can capture this dependency with a model of the domain using a transformation of coordinates denoted $\psi : \Gamma \rightarrow \Omega$, we can reform the coordinates, such that the relevant spatial information is contained in the first dimension of the reformed space Γ . We can then project the KDE by a partial integration of the domain:

$$\Phi_\sigma^*(t, u) = \int_{\psi^{-1}(\Omega)} \Phi_\sigma(t, \psi(u, v, w)) |\det(D\psi)(u, v, w)| dv dw \quad (8)$$

A coordinate transformation of an integral is performed by substituting the new variables and including the determinant of the Jacobian matrix

$|\det(D\psi)|$. The projection is then obtained by integrating in the new coordinates u, v, w over everything except the new, independent variable u . This results in a simplified representation of the kernel density estimate, where the relevant spatial structure is captured in a single, new spatial dimension. For example, in case of a spherical transformation the new representation would be called a radial density.

Looking at the definition of Φ in Eq. 5, it consists of a weighted sum of kernels. Hence, just as in Eq. 7, the projection onto the relevant axis can be computed using the projections of the kernels. The reconstructed KDE in the projected coordinates then boils down to

$$\Phi_\sigma^*(t, u) = \sum_{i=1}^N \mathcal{F}_i(\vec{p}_i(t)) g_\sigma^*(u(\vec{x}), u(\vec{p}_i(t))) \quad (9)$$

where g_σ^* is the projected kernel and $u(\vec{x})$ is the coordinate u of the transformation $\psi^{-1} : \Omega \rightarrow \Gamma$.

We take the traditional integral transformations from calculus [44] based on the dimensionality of the region of interest: 1) point of interest \rightarrow spherical transformation; 2) line of interest \rightarrow cylindrical transformation; 3) area of interest \rightarrow orthogonal transformation.

The resulting projections of the Gaussian kernels can be seen in Fig. 2 and the derivations in the supplementary material. Looking at Fig. 2 we can observe that the projection onto both the spherical and cylindrical radial distance mostly takes care of the boundary condition when $r \rightarrow 0$ and behaves like a traditional Gaussian for $r \rightarrow \infty$.

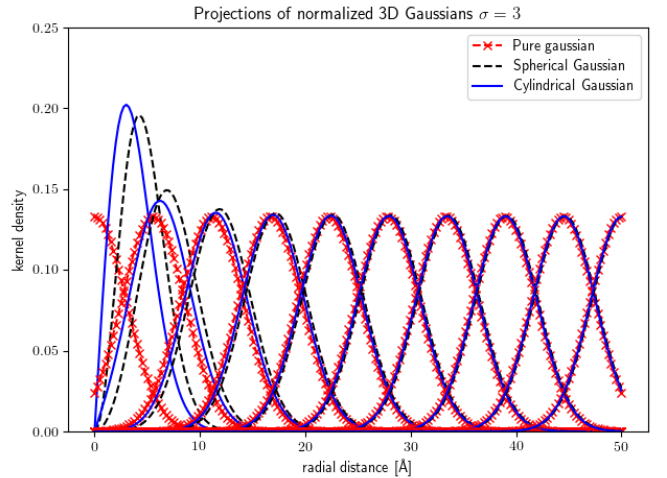


Fig. 2: Spherically (black) and cylindrically (blue) projected Gaussians compared with pure Gaussian kernels (red crosses) centred at positions uniformly sampling the radial axis. We can observe that the projection takes care of the boundary condition when $r \rightarrow 0$. Each kernel that would extend into the negative domain ($r < 0$) gets “pushed out” such that it zeros out at $r = 0$.

3.4 Time-Scale Representation

Treatment of the temporal evolution is illustrated in Fig. 3. First the traditional representation of the integral measure by a timeseries $f(t)$ is computed. The timeseries describes the temporal evolution of the integral measure, disregarding the spatial structure. To capture the temporal scales in the timeseries the scale-space representation is generated as described in Sec. 3.1, obtaining a scale-space image and a level set of zero crossing curves. The scale-space image is used for the temporal scale exploration as detailed in Sec. 4.2. This corresponds to the first row of Fig. 3.

Constructing the temporal scale-space for the KDE representation, that captures the spatial structure, is done by repeating the very same process. Considering that the function $\Phi(t, \vec{x})$ (representing the KDE) consists of a timeseries in each point \vec{x} , we apply Eq. 1 at each of these points. This amounts to convolving the data along the temporal axis

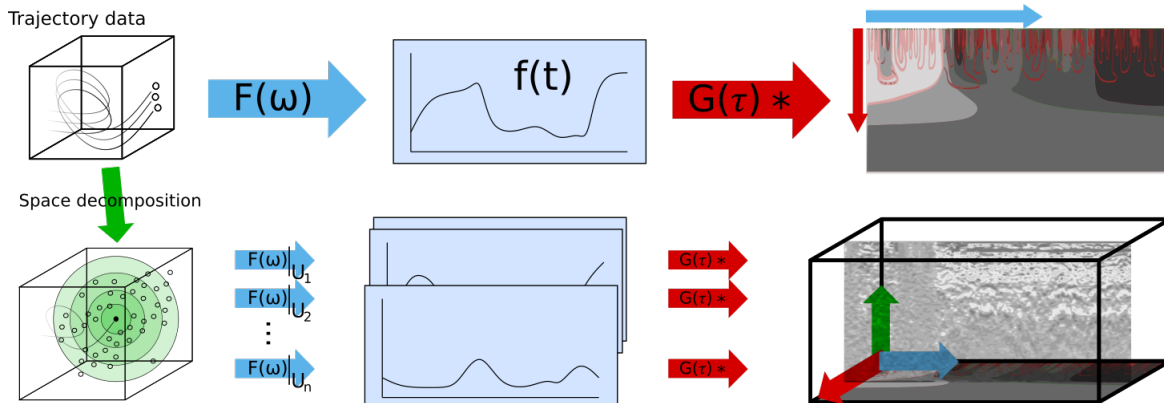


Fig. 3: Overview of the temporal scale-space construction. First row: the integral measure $F(\cdot)$ is evaluated for each frame of the particle data, yielding a time series $f(t)$, which is then convolved with a set of Gaussian kernels, resulting in a traditional scale-space; Second row: A point location of interest is identified; the spatial domain is decomposed into shells; the integral measure is evaluated for each region giving a set of time series; the scale-space construction is repeated for each and the results are stacked into the SSST cube.

with Gaussian kernels of increasing sizes, as expected from treating the time and space independently. This construction, however, creates a large amount of data that would necessarily be projected and for analytical purposes. To avoid this large amounts of intermediate data, we construct the temporal scale-space of the projected KDE representation directly. Looking at Eq. 8, we see that the projection is done by integrating over the spatial domain only. Plugging the definition of $\Phi^*(t, u)$ into Eq. 1 and using Fubini's theorem [44] to exchange the order of the integrals, we prove that the scale-space construction and the space reformation are order independent. The projected temporal scale-space is therefore obtained by a direct convolution of $\Phi^*(t, u)$ with Gaussian kernels along the temporal axis.

Even though the order of the operations is not important, the transformation itself will introduce temporal dependency if it changes over time. It is important to keep this mind when interpreting the results, however, it is still a satisfying result as it enables us to first project the data and construct the scale-space of much smaller representation. The temporal dependency introduced by the projection will be the same as the one introduced by projecting the temporal scale-space of the full KDE. The correct choice of a space reformation will capture the spatial dependency in the data. Whereas the scale-space representation helps identifying features at different scales. The scale-space in the reformed coordinates is the basis for the slice view detailed in Sec. 4.3.

The construction of the temporal scale-space in the reformed coordinates is illustrated in the second row of Fig. 3.

3.5 Numerical Treatment

This section so far provides a rigorous approach to constructing a scale-space for particle datasets in reformed coordinates. Following the definitions for the numerical implementation would, however, require the recomputation of the function Φ^* using Eq. 9 for a range of scales. As summing over all the particles is the most demanding part of the computation, we can compute the values of the projected KDE only once, for a small value of σ and obtain the higher scales by integral transformation (generalized convolution) with the projected kernels.

$$\Phi_{\sigma+s}^*(t, u) = \int g_s^*(u, u_0) \cdot \Phi_{\sigma}^*(t, u_0) du_0 \quad (10)$$

Looking at the shapes of the projected kernels in Fig. 2, we conclude that they copy the Gaussian kernels, except for the boundary. With care, we can take the optimization even further and approximate the scale-space by convolving the function $\Phi^*(t, u)$ with a simple Gaussian kernel:

$$\Phi_{\sigma+s}^*(t, u) \approx g_s(u) * \Phi_{\sigma}^*(t, u) \quad (11)$$

This is an approximation only in the spatial case. For the temporal scale the Gaussian convolution is the exact result.

4 VISUALIZATION & INTERACTION DESIGN

In this section, we present the use of scale-space splatting as a method to bring spatial information into aggregate measures in the analysis of molecular dynamics (MD) data. We do not provide a replacement of well-proven tools, such as VIA-MD [43], which are clearly useful for many purposes. Instead, we demonstrate that considering the spatial context of the aggregate measures provides additional information, which can be exploited to build improved tools for visual analytics of MD simulations.

The goal of our visual analytics solution is to enable the exploration of multi-scale features using the aggregate measures of trajectory-based spatio-temporal data in molecular dynamics. The aim is not to promote the exploration of scale-space over the original simulation, but to enhance the analysis with a spatio-temporal scale-space approach. The standard way of exploring such data is through 3D animated rendering of the spatial configuration, steered in time using an interactive time axis. An integral measure, such as interaction energy, represented through a timeseries $f(t)$ using a line plot, is often used for navigating the time axis. However, a simple line plot is not sufficient to capture the multiscale spatio-temporal behaviour.

We propose a multiple coordinated view solution for exploring multiscale features of the MD simulation. This is done by navigating and brushing in a scale-space representation linked to a 3D view showing the selected spatiotemporal intervals. Users can select temporal ranges to loop over in an animated 3D view where the spatial extent is expressed in focus and context highlighting manner. Efficient computation of the spatiotemporal scale-space of the integral measure using the scale-space splatting algorithm enables interactive exploration of scales even for large datasets.

The views include a 2D representation of the temporal scale and a 2D representation of the spatiotemporal slices of the scale-space space-time cube. In order to navigate the temporal scale, we propose a representation combining a line plot enhanced with a temporal scale-space image, which provides the context of different scales (see Fig. 4). We provide interaction with the scale dimension and visual cues to support the exploration, as further detailed in Sec. 4.2.

In order to explore the spatial dependency of the integral measure at different scales, we render slices of the scale-space space-time cube (Sec. 4.1). We do this by encoding the density field as a pixelmap, where the colour of each pixel encodes the values of the integral measure at corresponding spatial and temporal coordinates. This is described in Sec. 4.3. Both views share the temporal scale axis, which enables the exploration of the temporal scale in the scale view, where it is easier to grasp the effect of the reconstruction kernels. As it is difficult to read off exact values from the density plots using colormaps, plotting of the slices of the slice view is also supported.

The description of the linking and investigating the spatial configu-

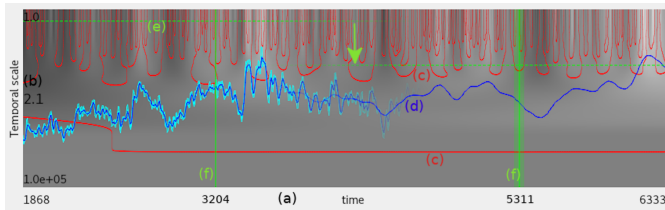


Fig. 4: The scale view provides the investigation tool for the temporal scale of the data. The horizontal axis is time (a), the vertical axis is temporal scale (b). The background is the scale-space image with its second derivative level-set highlighted in red (c). The value of the timeseries at a selected scale is in blue (d) with the min-max interval in cyan. The change of the scale control (e) from value $\tau = 2.1$ to $\tau = 16.1$ results in the disappearing of the min-max interval and a thicker window around the time pointer (f). The picture is a horizontal blending (left to right) of two screen-shots with the time indicator at different positions.

ration of the particles in a 3D rendering is described in the Sec. 4.4.

An overview of our application can be seen in Fig. 1. A demonstration of supported interactions can be found in the video in the supplementary materials.

4.1 The Scale-Space Space-Time Cube

Bach et al. [3] introduced a unified framework of space-time cubes. In their work they summarize the nomenclature for data manipulation and visualization of spatio-temporal data. A conceptual image of a cube is introduced to unify the spatial and temporal axes. An example would be stacking the frames of a movie into a box. The x and y coordinates correspond to the screen space coordinates, whilst the z axis corresponds to time. Playing the movie then corresponds to slicing the box (cube) perpendicular to the z axis. Applying the space time cube to the movie data allows for exploring novel visualization techniques such as using non-orthogonal slicing, drilling, flattening (aggregating) and other data transformation techniques.

In order to manage our construction in a visualization, we introduce the concept of a scale-space space-time (SSST) cube. An SSST cube builds on the concept of the space-time cube as described by Bach et al. but treats one of the axes as temporal scale and one as spatial scale. In our case, the cube has 4 dimensions: time, reformed space, temporal scale, spatial scale. The temporal scale is constructed by convolving the data with Gaussian kernels along the temporal dimension. The spatial scale is constructed by convolving with Gaussian kernels along the new spatial axis. We slice the cube in time-space planes in order to visualize spatiotemporal features. We flatten the cube along the spatial axis, to obtain the standard time series representation of the integral measure as computed over the whole spatial domain, enhanced with the time-scale scale-space information. We drill along the spatial axis for direct visualization of values at given time and scales.

4.2 Scale View

We find that there are many visualization papers utilizing scale-space approaches, however, there is not much research done on direct visualization and interaction techniques for scale-space datasets. There are two examples, both using a colourmap to show a scale-space image. An example from statistics [11], uses a scale-space image coloured based on statistical properties of the reconstructed timeseries to find prominent features. The second example is Pinus [42], which computes a hierarchical aggregation of the investigated timeseries. They employ a triangular-shape image representation of their scale-space, where pixels are coloured by the aggregated values.

For showing the temporal scale-space of the timeseries $f(t)$, we adapt the direct colourmapping of the scale-space image, similar to both SiZer [11] and Pinus [42]. Since we are dealing with long timeseries data and the representation in Pinus scales quadratically with the number of timesteps, we adopt the log-scaled scale axis from SiZer.

This way we end up with showing a long strip $n \times \log(n)$ pixels, rather than a huge $n \times n$ matrix, (the background image in Fig. 4).

The log-scaling is justified by linearizing the effect of the blurring of the time-series. A small change $\Delta\tau$ in the size of a large kernel $\tau \gg \Delta\tau$ has very little effect on the resulting time series. On the other hand, the same change has much bigger effect for kernel size comparable to the change itself ($\Delta\tau \approx \tau$). Using the exponential function to explore the scale dimension keeps the relative change of the kernel size $\frac{\Delta\tau}{\tau}$ constant. Exponential sampling of the dimension leads to the log-scaled axis, ensuring that a fixed change anywhere along the re-scaled axis has a proportional effect on the appearance of the timeseries. In practice, using a linearly scaled axis would result in the uninteresting, grey, bottom part of the image being much wider, up to the point of not having anything else in the picture.

Since any colourmap itself is insufficient for reading off exact values of the timeseries, we provide slicing of the image, similar to the approach in Pinus [42]. The only difference is that we plot the timeseries on top of the rectangular scale-space image (Fig. 4d). The selected scale is indicated by a horizontal dashed line (Fig. 4e) and the exact value of scale is displayed on the left (Fig. 4b). Due to the finite resolution of the display, it is impossible to show the reconstructed function for some combinations of zoom and scale. This effect is well known in the field and can be addressed for example by using a band graph [1] as in BinX [4] and Duran [15].

Sliding the green dashed line up and down provides means for exploring the temporal scale space. We found that the scale-space image in the background of the line plot successfully highlights prominent regions in the data, but it provides only a vague indication of interesting scales. We therefore include the red curves (Fig. 4c), which are the contours of zero crossings of the second derivative computed according to Eq. 2. They provide a visual cue for regions of peaks and valleys in the data, together with corresponding scale significance. The “deeper” the contour reaches, the more prominent the given peak or valley is. The feature disappears from the line plot completely when the selected scale does not cross the contour.

In order to interact with the time axis, the selected time indicator is drawn as a vertical green line (Fig. 4e), which can be dragged along the time axis and changes the currently selected time step. Since the selected time step at the current time scale is the result of averaging over a neighbourhood the indicator is enhanced with a semi-transparent green box showing the size of this neighbourhood. It is drawn around the vertical green line (Fig. 4f) showing the $(t - \tau, t + \tau)$ interval corresponding to the width of the Gaussian kernel of the currently selected scale.

4.3 Slice View

The slice view provides the spatiotemporal overview of the simulation with respect to the chosen integral measure. It does its job by conveying the information contained in the scale-space space-time cube. The cube is a four dimensional density field which is traditionally visualized

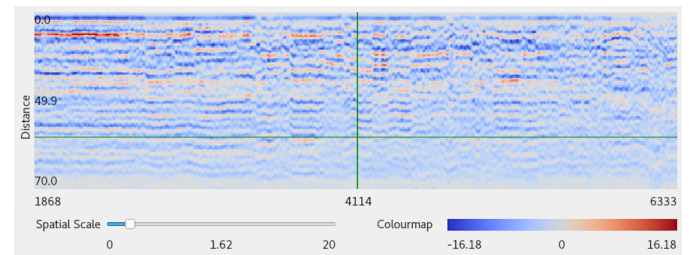


Fig. 5: The slice view: a pixelmap showing a spatio-temporal slice of the SSST cube with horizontal time axis and vertical spatial axis. The green cross consists of a horizontal line for navigation in the spatial axis and a vertical line for navigation along the time axis. A divergent colourmap is used to encode the values. The spatial scale is set to $\sigma = 1.62 \text{ \AA}$ and temporal scale is shared with the scale view $\tau = 2.1$.

with 1D [47] and 2D [48] slices. The idea is to incorporate the spatial information into the lineplot representation of the timeseries. To preserve coherence between the two representations, we keep the temporal axis of the line plot and use the second axis for the reformed spatial dimension. Since we can no longer use the second axis for encoding the values, as in the line plot, we utilize a pixelmap equipped with an appropriate colourmap. The colour of each pixel is determined by the value in the scale-space space-time at given spatial and temporal scales, with the x coordinate corresponding to the time axis and the y coordinate corresponding to the reformed spatial axis. To support the exploration of exact values in the pixelmap we enable plotting of the spatial dependency at the given timestep and the particular scales in a separate view (Fig. 1 d).

When the frequencies in the data are higher than what is possible to capture on the screen, the pixelmap suffers from the same overdraw as the line plot. Since there is not an easy solution for including the overplotted ranges in the pixelmap, we rely on the joint properties of the scale view and the slice view. The slice shares its values with the timeseries representation – the values in the slice aggregated over the spatial domain yield the timeseries – they will also roughly share the same frequencies in the data. We exploit this correspondence by sharing the temporal scale axis between the two views. It is thus possible to perform the exploration of the temporal scale in the scale view and propagate the interaction to the slice view. The scale view thus serves as a proxy for exploration of the temporal scale of the slice view with explicit indication of the effect of the overdraw.

A similar scale exploration tool should, in principle, be available also for the spatial scale, where the aggregate of the slice along the temporal axis would guide the exploration. In our case, where the temporal axis is much larger than the spatial axis, we found it sufficient to enable the spatial scale exploration by a slider and investigating the effect of the scale selection in a simple line plots of values of the cube along the spatial dimension for a selected time (see Fig 1 d). An advanced tool, such as the scale view, is necessary for exploration of a scale axis when interacting with various level of zoom.

The second link between the slice view and the scale view is realized through the shared temporal scale axis, which supports panning and zooming. The green cross in Fig. 5 indicates the temporal and spatial selection used for the 3D view.

4.4 3D view

There is a body of work on 3D visualizations of particle data in MD, SPH, and also other fields [2, 17]. Since 3D representations are not central in our research, we default to the simplest, useful, particle representation by shaded spheres at particle positions with user-adjustable sizes. We render the configuration of the particles at the time step selected in the shared temporal axis in slice and scale view. The way of focusing on a given spatial scale is by hiding inner or outer regions with respect to the decomposition. The spheres in the focused region are coloured with the same colourmap as in the slice view, with the values based on the contribution of each particle to the integral measure. By this, the user is able to pinpoint a particular interaction both spatially and temporally. An example of peeling the concentric spheres from around the point of interest in a protein–ligand simulation can be found in Fig. 6.

4.5 Implementation

The implementation consists of two parts, a pre-processing module and an interactive visual analytics tool. Both the data pre-processing and the visual analytics tool are implemented in Python. For a given measure and a reformation, the pre-processing consists of filling the data into one spatiotemporal slice of the scale-space space-time cube at the highest resolution. For this the whole simulation is processed frame by frame, where each particle is transformed into the new coordinate system and its contribution is projected onto the chosen axis. The timeseries is obtained by aggregating the slice along the spatial axis. The scale-space image of the timeseries is constructed by repeated convolutions with Gaussian kernels and the zero crossings are traced with sub-pixel precision.

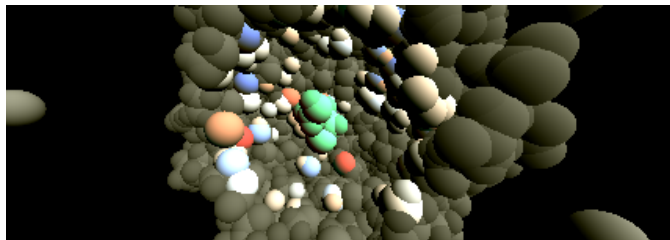


Fig. 6: An example of the 3D view from the molecular dynamics simulation showing the peeling of spatial scales around the point of interest. The particles in the selected region are coloured based on their contribution to the integral measure. The atoms are coloured based on their contribution to the integral measure and the colourmap is shared with the slice view. The ligand representing the point of interest is in green.

All the other slices of the scale-space space-time cube are computed on demand by convolution with Gaussian kernels. As demonstrated in Sec. 3.3, the projected kernels are also Gaussians, except for when they overlap with the boundary $r \rightarrow 0$. Since in our data all the values near the $r = 0$ are zero (see Fig. 1, top right), we do not explicitly implement the boundary condition as described in Eq. 10.

The user interface has been designed using bindings for the Qt framework. It runs interactively with the only demanding parts being convolving the matrix in the slice view with 1D Gaussian filters, recomputed on the scale change. This could be easily optimized by convolving the matrix on the GPU.

5 DEMONSTRATION

In the following, we describe the use of the new method in our proof-of-concept application. We worked on two datasets, each in collaboration with a domain expert: one protein engineer and one expert in computational biology with background in quantum chemistry. The data consists of the trajectories of the atoms of a protein and of a smaller molecule called a ligand. One of the integral measures commonly used for the analysis of MD simulations is electrostatic interaction energy.

The electrostatic interaction energy is computed from the pairwise interaction of the atoms of the ligand and the protein by the formula:

$$F_{el} = \sum_i \sum_k \underbrace{\frac{1}{4\pi\epsilon_0} \frac{q_i q_k}{\|\vec{p}_k - \vec{p}_i\|}}_{\mathcal{F}_i(\vec{p}_i)}$$

where i is indexing the atoms of the protein, k is indexing the atoms of the ligand, \vec{p} is the atom position, and q is the atom partial charge. The letter F does not stand for force, but for the integral measure, in accordance with the definition in Eq. 3.

Since the position of the ligand is of importance in the analysis, we identify the position of the ligand as the reference point and present an analysis of the data using the spherical transformation. The details of the transformation are described in Sec. 3 of the supplementary material. Further results obtained using cylindrical and orthogonal transformations are demonstrated in Sec. 4-5 in the supplementary material.

5.1 Three Ligands

The first dataset we studied contains 50k timesteps of an MD simulation of a protein with 4650 atoms and three ligands, each consisting of 12 atoms. The research question addressed with the simulation is the analysis of the ligands' propagation to the active site. The data come from the research published in the paper by Marques et al. [33].

Exploring the spatio-temporal scale-space of the integral measure improves navigation through the simulation. Comparing the timeseries representation of the electrostatic interaction energy with the slice view, the slice view provides an immediate indication of the intervals when the ligand is interacting with the protein and when not (see Fig. 1 a). The visualization also shows, that there is no interaction past the 25 Å

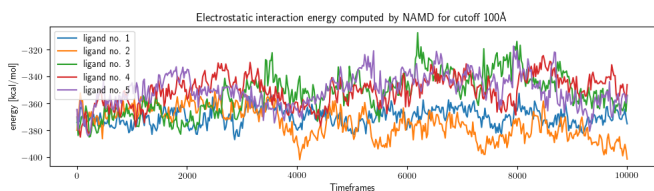


Fig. 7: Comparison of the electrostatic interaction energies for the ligands no. 1 - 5. There is no clear deviation of the energy values for the two cases that undergo the change.

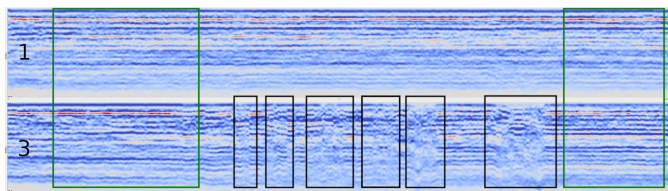


Fig. 8: The comparison of the slice views for ligands no.1 and no.3 providing an overview of the simulation. The ligand no.1 exhibits stable behaviour during the whole simulation as opposed to the ligand no.3 which undergoes several changes, framed in black. Both ligands show similar behaviours towards the beginning and the end of the simulation.

distance, which is important information for correctly setting up the simulation and is an object of study on its own [16]. Zooming on the 0 – 25 Å interval is performed at the data level, but interactive spatial navigation would be necessary for a final-product application. When investigating ligand no. 3 at appropriate spatial and temporal scales, there are three distinct features that draw instant attention. The first one are the empty intervals towards the beginning of the simulation, where the ligand is too far away from the protein (Fig. 1 a). The second one starts with the ligand trying to escape the tunnel and as it approaches the surface it is pushed back into the tunnel. A second ligand appears at the entrance, locking the ligand no. 3 in the active site (Fig. 1 b). The third interesting region occurs at times when all three ligands are present near the same tunnel opening (see Fig. 1 c). Comparing these features to the very weak indications in the timeseries representation, our proposed visualization makes the identification of relevant features in the data more apparent.

5.2 Five Ligands

We have obtained an MD simulation dataset from our collaborators consisting of 5 ligands (26 atoms each) and a protein of circa 16 thousand atoms, with the duration of 10 000 timesteps. The research question that led to the computation of this simulation was to study the ligand’s dynamics in the binding pockets over time. We received the dataset with a conclusion: “There is nothing [to be seen] in the electrostatic energy. It has to be coupled with other measures.”.

The plot of the energies can be found in Figure 7. Indeed, the overview plot does not separate the ligands into distinct categories, nor does it segment the time axis into different modes.

Loading the pre-computed data into the tool we find that 10 000 timesteps are causing overdraw on the time series representation and there is little to be seen in the initial scale slice. Interacting with the temporal scale reveals that in fullscreen on our 1920px wide monitor the τ of 9 timesteps results in a smooth plotting of the energy function. Due to the joint treatment of the temporal scale axis we also automatically obtain a good setting for the slice view. Hence, a preliminary vertical pattern appears also in the slice view. Interacting with the spatial scale, after some attempts we identify that above 4 Å a vertical pattern is dominant in the data and the horizontal stripes get blurred out. The scale size of 2 Å preserves the structure and enhances the contrast. Comparing the overview visualizations for all the ligands we find that ligands no. 1, 2 and 4 exhibit stable behaviours throughout the

simulation, whereas ligands no. 3 and 5 experience various changes.

The cross-scale exploration enables the identification of prominent features in the data. For example, applying a considerable amount of blurring in the spatial domain ($\sigma = 5.86$ Å), and a relatively low amount in the temporal domain ($\tau = 10.7$), we see spatially significant changes that possibly happen over short time periods. The effect of these settings on the slice view for the ligand no. 5 can be seen in Fig. 9 (the mostly blue pixelmap). Two significant events are apparent in the scale view. The first change in the pattern is a shear occurring at timestep 2600 (Fig. 9a,b). Investigating the 3D configuration shows a big movement of the ligand inside of the binding pocket. This event is not apparent in the energy timeseries, however, it allowed the ligand to reach a better position where the energy reaches its minimum after this change (9 c). The biggest contribution to this energy minimum comes from atoms located at distances between 16.6 and 22.2 Å from the ligand’s centre.

The second significant event happens towards the end of the simulation, where the original pattern is interrupted and it is paired with a spike in the energy timeseries (Fig. 9 1,2,3). The snapshots of the ligand before, during and after the change show the ligand changing its orientation.

The following can be observed from the example described above. The descriptive power of the fully aggregated integral measure was not sufficient to distinguish the interesting behaviour. Computing the spatial dependency in full resolution and simply showing the data does not suffice for finding patterns in the data. Exploring the scale-space representation enabled rapid identification of prominent features, which lead to spatial and temporal localization of interesting behaviour. This example demonstrates the analytical power of accessing spatial information of an integral measure. The domain expert confirmed that the timeseries representation does not provide sufficient guidance in this particular case. On the other hand, being able to spatially resolve the respective contributions to the timeseries at correct scales, we are able to gain more insight from exactly the same data. We have thus demonstrated that there is more information accessible through the integral measure when the spatial dependency is also included and relevant scales are considered.

5.3 Domain Expert Feedback

The presented use case scenarios were worked out in cooperation with two domain experts in order to validate our results. To evaluate the usability of the prototype tool, we have conducted a 45 minutes demonstration followed by an interview with one of the domain experts. The domain expert answered questions to ensure his understanding of the application and was tasked to guide a short exploration of an MD simulation, while the authors were interacting with the application.

For the “three ligands” dataset two spherical projections were available, one centred at the ligand and one centred at the protein. The expert chose to work in the spherical projection centred at the protein and pointed out that the slice view provides a good overview of the simulation. He navigated to a time interval corresponding to an extremum in the energy and observed that for a detailed analysis more context in the 3D view, such as amino-acid/atoms labels and different 3D representations would be helpful. The strong point of the slice view was that the spatial orientation of the ligand with respect to the protein could be deduced. He also mentioned that “It would be interesting to compare the Van der Waals and electrostatic using this view.” While the Van der Waals energy is not at the moment supported, it is an example of an integral measure that can be easily incorporated in the future. He found the treatment of the scales appropriate as there are many different time scales in the simulation and one could easily examine the persistence of features across scales. With respect to the usability he mentioned, that “It would, of course, require some training, but it does not seem more difficult than other analytical software available.”.

5.4 Performance

The pre-computation of the spatial decomposition along the ligand trajectory is the most demanding part. For the decomposition of 5 ligand trajectories against a protein with 16 thousand atoms and 10 000

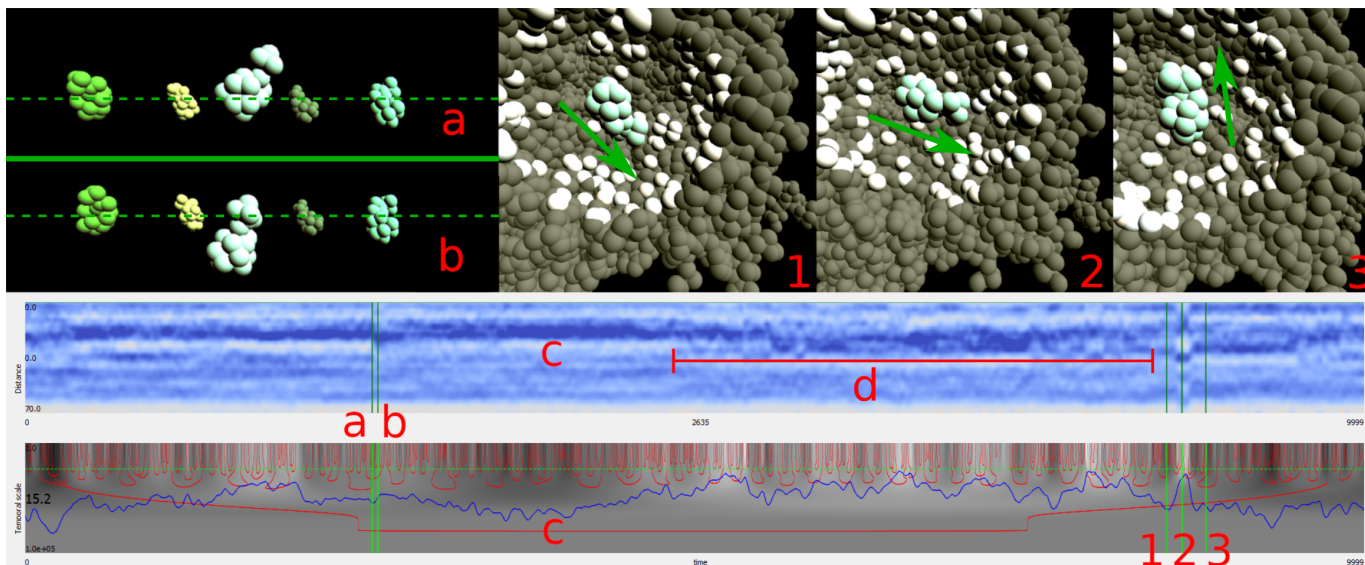


Fig. 9: The analysis of behaviour of ligand no. 5 at large spatial scale. The two most prominent features in the slice view (middle) are the shear towards the beginning of the simulation (a,b) and the abrupt modality change towards the end (1,2,3). The shear is caused by movement of the ligand in the binding pocket. Protein is not shown to avoid occlusion, instead the movement is referenced by the plane that ligands lie in. The shear follows by energy minima which is largely contributed by the protein atoms located at a distance between 16.6 - 22.2 Å (c). The change of modality (1,2,3) results after a period of unstable behaviour (d) and is resolved as ligand's rotation (the arrows indicate the ligand's orientation).

time-frames the pre-computation took 20 minutes per ligand in a single core application on a system with following parameters: CPU: Intel(R) Core(TM) i7-7700K @ 4.20GHz; RAM: 2 x 16GiB DIMM DDR4 2400 MHz; GPU: GeForce GTX 1080, 8 GB GDDR5X; Storage: 500GB SSD Samsung 860 Evo.

A full ligand-protein electrostatic interaction for a ligand consisting of 26 atoms took 6.5 hours. We acknowledge that the implementation is not optimized and better tools for computing these values are available [37]. The domain experts do not find the computational time to be an issue, since our precomputation time is negligible with comparison to the computational time of the simulation. The result of the precomputation is scalar 2D array of $n \times m$ values, which is easily dealt with both on CPU and GPU.

6 DISCUSSION AND LIMITATIONS

We have demonstrated that the simultaneous treatment of spatial and temporal scales in the analysis of MD simulation can lead to an improved analysis. Of course, this information does not have to be accessed in every analysis, namely when the timeseries representation sufficiently captures the studied phenomena. Being able to access the additional information is vital when other approaches fail.

The reason that our method works better in these cases lies in enabling exploration of the spatial structure by efficient computation of scale-space. Where previous approaches would be able to reconstruct the spatial representation of the integral measure, they would do so at a fixed resolution [37], without an interactive visual support. We, on the other hand, enable a multiscale exploration of both spatial and temporal dependencies of the integral measures, which allows more flexibility in the analysis.

The biggest obstacle for a widespread use of the method is a steep learning curve regarding the scale-space. In particular, we had hard time explaining the scale-space basics and the mechanics of the scale view to the domain experts. Yet, they found our demonstration of the analysis powerful and the results relevant. The steep learning curve of scale-spaces needs to be balanced with an attractive, well-designed, user-friendly interface that can bridge the gap with ease of use over years of knowledge in calculus and differential equations. This would ideally be combined in one of existing MD analysis tools, since the domain experts rely on more than one approach in a successful analysis.

7 CONCLUSION AND FUTURE WORK

In this work, we presented a method for investigating the spatial dependence of integral measures of particle data across temporal and spatial scales. Based on the theoretical framework we have implemented a sample demonstration application showcasing the potential benefits of simultaneous spatial and temporal scale exploration of trajectory-based particle data in molecular dynamics. We have shown that this kind of analysis is suitable for the exploration of one-dimensional spatial dependencies in aggregate measures. Consequently, such exploration can be a valuable tool for real world applications, helping in understanding long time simulations. As the core of the method depends only on trajectory based simulation, equipped with an integral measure, exhibiting a simple symmetry, we expect that it could be applied to other fields outside of molecular dynamics. This will be subjected to further research.

There are several research directions open for future work. A (semi-) automatic scale identification could be investigated for pre-selection of relevant scales to enhance the user experience. Furthermore, the unification of spatial and temporal scales might allow for automatic feature detection in the spatiotemporal scale-space similar to the works of Laptev [27] and Lindeberg [31, 32], however, this might be strictly domain or even data dependent. There is also a possibility for adapting the framework to non-uniform scale-space constructions, similar to the edge enhancing diffusion of Perona and Malik [36].

Also more visualization techniques could be adapted for the scale-space space-time cube, namely non-planar slicing – adaptively changing the reconstruction precision based on an importance function. Furthermore, promoting the scale-space as the prime object of interest and using direct volume rendering, or iso-surface extraction to trace its structure can be of interest. This might combine greatly with the Gaussian derivatives of the data.

ACKNOWLEDGMENTS

We would like to thank our collaborators: Dr. Sérgio Marques from Loschmidt Laboratories and Hanif Muhammad Khan from Computational Biology Unit at University of Bergen, for providing us with the data and their invaluable feedback.

REFERENCES

- [1] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of Time-Oriented Data*. Springer Publishing Company, Incorporated, 1st ed., 2011. doi: 10.1007/978-0-85729-079-3
- [2] N. Alharbi, M. Alharbi, X. Martinez, M. Krone, A. S. Rose, M. Baaden, R. S. Laramee, and M. Chavent. Molecular Visualization of Computational Biology Data: A Survey of Surveys. In B. Kozlikova, T. Schreck, and T. Wischgoll, eds., *EuroVis 2017 - Short Papers*, pp. 133–137. The Eurographics Association, 2017. doi: 10.2312/eurovisshort.20171146
- [3] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale. A descriptive framework for temporal data visualizations based on generalized space-time cubes. In *Computer Graphics Forum*, vol. 36, pp. 36–61. Wiley Online Library, 2017. doi: 10.1111/cgf.12804
- [4] L. Berry and T. Munzner. BinX: Dynamic exploration of time series datasets across aggregation levels. In *IEEE Symposium on Information Visualization*, Posters Compendium, pp. 5–6. IEEE, 2004.
- [5] P.-T. Bremer, G. Weber, J. Tierny, V. Pascucci, M. Day, and J. Bell. Interactive Exploration and Analysis of Large Scale Simulations Using Topology-based Data Segmentation. *IEEE Transactions on Visualization and Computer Graphics*, 17(9):1307–1324, 12 2010. doi: 10.1109/TVCG.2010.253
- [6] S. Bruckner. Dynamic Visibility-Driven Molecular Surfaces. *Computer Graphics Forum*, 38(2):317–329, 2019. doi: 10.1111/cgf.13640
- [7] J. Buchmüller, D. Jäckle, E. Cakmak, U. Brandes, and D. A. Keim. Motion-rugs: Visualizing collective trends in space and time. *IEEE Transactions on Visualization and Computer Graphics*, pp. 76–86, 2018. doi: 10.1109/TVCG.2018.2865049
- [8] S. Byna, J. Chou, O. Rubel, Prabhat, H. Karimabadi, W. S. Daughter, V. Roytershteyn, E. W. Bethel, M. Howison, K. Hsu, K. Lin, A. Shoshani, A. Uselton, and K. Wu. Parallel I/O, analysis, and visualization of a trillion particle simulation. In *SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pp. 1–12, Nov 2012. doi: 10.1109/SC.2012.92
- [9] J. Byška, A. Jurčík, M. E. Gröller, I. Viola, and B. Kozlíková. MoleCollar and Tunnel Heat Map Visualizations for Conveying Spatio-Temporo-Chemical Properties Across and Along Protein Voids. *Computer Graphics Forum*, 34(3):1–10, 2015. doi: 10.1111/cgf.12612
- [10] J. Byška, M. L. Muzic, M. E. Gröller, I. Viola, and B. Kozlíková. AnimoAminoMiner: Exploration of Protein Tunnels and their Properties in Molecular Dynamics. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):747–756, Jan 2016. doi: 10.1109/TVCG.2015.2467434
- [11] P. Chaudhuri and J. S. Marron. SiZer for Exploration of Structures in Curves. *Journal of the American Statistical Association*, 94(447):807–823, 1999. doi: 10.1080/01621459.1999.10474186
- [12] M. Chen, D. Silver, A. S. Winter, V. Singh, and N. Cornea. Spatial transfer functions: A unified approach to specifying deformation in volume modeling and animation. In *Proceedings of the 2003 Eurographics/IEEE TVCG Workshop on Volume Graphics*, VG '03, pp. 35–44. ACM, New York, NY, USA, 2003. doi: 10.1145/827051.827056
- [13] E. Chovancová, A. Pavelka, P. Beneš, O. Strnad, J. Brezovský, B. Kozlíková, A. Gora, V. Šustr, M. Klvaňa, P. Medek, L. Biedermannová, J. Sochor, and J. Damborský. CAVER 3.0: A tool for the analysis of transport pathways in dynamic protein structures. *PLoS Computational Biology*, 8(10):1–12, 2012. doi: 10.1371/journal.pcbi.1002708
- [14] R. Dafner, D. Cohen-Or, and Y. Matias. Context-based Space Filling Curves. *Computer Graphics Forum*, 19(3):209–218, 2000. doi: 10.1111/1467-8659.00413
- [15] D. Duran, P. Hermosilla, T. Ropinski, B. Kozlikova, A. Vinacua, and P. P. Vazquez. Visualization of Large Molecular Trajectories. *IEEE Transactions on Visualization and Computer Graphics*, pp. 987–996, 2018. doi: 10.1109/TVCG.2018.2864851
- [16] E. Fadrná, K. Hladečková, and J. Koča. Long-range electrostatic interactions in molecular dynamics: an endothelin-1 case study. *Journal of Biomolecular Structure and Dynamics*, 23(2):151–162, 2005. doi: 10.1080/07391102.2005.10531229
- [17] R. Fraedrich, S. Auer, and R. Westermann. Efficient High-Quality Volume Rendering of SPH Data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1533–1540, Nov 2010. doi: 10.1109/TVCG.2010.148
- [18] R. Fraedrich and R. Westermann. Motion visualization in large particle simulations. In *Proceedings of IS&T/SPIE Electronic Imaging 2012*, vol. 8294, pp. 1–12, 2012. doi: 10.1117/12.904668
- [19] R. Gingold and J. Monaghan. Kernel estimates as a basis for general particle methods in hydrodynamics. *Journal of Computational Physics*, 46(3):429–453, 1982. doi: 10.1016/0021-9991(82)90025-0
- [20] T. Hägerstrand. What about people in regional science? In *Papers of the Regional Science Association*, vol. 24, pp. 6–21. Springer, 1970. doi: 10.1111/j.1435-5597.1970.tb01464.x
- [21] L. Holmström. Scale space methods. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(2):150–159, 2010. doi: 10.1002/wics.79
- [22] W. Humphrey, A. Dalke, and K. Schulten. VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics*, 14:33–38, 1996. doi: 10.1016/0263-7855(96)00018-5
- [23] A. Kanitsar, R. Wegenkittl, D. Fleischmann, and M. E. Gröller. Advanced curved planar reformation: flattening of vascular structures. *IEEE Transactions on Visualization and Computer Graphics*, pp. 43–50, Oct 2003. doi: 10.1109/VISUAL.2003.1250353
- [24] G. L. Kindlmann, R. S. J. Estépar, S. M. Smith, and C. F. Westin. Sampling and visualizing creases with scale-space particles. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1415–1424, 2009. doi: 10.1109/TVCG.2009.177
- [25] T. Klein and T. Ertl. Scale-Space Tracking of Critical Points in 3D Vector Fields. In H. Hauser, H. Hagen, and H. Theisel, eds., *Topology-based Methods in Visualization*, pp. 35–49. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. doi: 10.1007/978-3-540-70823-0_3
- [26] J. Kreiser, M. Meuschke, G. Mistelbauer, B. Preim, and T. Ropinski. A survey of flattening-based medical visualization techniques. *Computer Graphics Forum*, 37(3):597–624, 2018. doi: 10.1111/cgf.13445
- [27] I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2-3):107–123, 2005. doi: 10.1007/s11263-005-1838-7
- [28] T. Lee and H. Shen. Visualization and exploration of temporal trend relationships in multivariate time-varying data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1359–1366, Nov 2009. doi: 10.1109/TVCG.2009.200
- [29] N. Lichtenberg, R. Menges, V. Ageev, A. A. P. George, P. Heimer, D. Imhof, and K. Lawonn. Analyzing residue surface proximity to interpret molecular dynamics. *Computer Graphics Forum*, 37(3):379–390, 2018. doi: 10.1111/cgf.13427
- [30] Y. Lin. *Subseries Join and Compression of Time Series Data Based on Non-uniform Segmentation*. PhD thesis, University of Waterloo, Ontario, Canada, 2008.
- [31] T. Lindeberg. *Scale-Space Theory in Computer Vision*. Springer US, 1 ed., 1994. doi: 10.1007/978-1-4757-6465-9
- [32] T. Lindeberg. Generalized Gaussian scale-space axiomatics comprising linear scale-space, affine scale-space and spatio-temporal scale-space. *Journal of Mathematical Imaging and Vision*, 40(1):36–81, 2011. doi: 10.1007/s10851-010-0242-2
- [33] S. M. Marques, Z. Dunajova, Z. Prokop, R. Chaloupkova, 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. doi: 10.1021/acs.jcim.7b00070
- [34] M. Meijers and P. van Oosterom. The space-scale cube: An integrated model for 2D polygonal areas and scale. In *28th urban data management symposium*, vol. 38, pp. 95–102, 2011. doi: 10.5194/isprsarchives-XXXVIII-4-C21-95-2011
- [35] H. Miao, T. Klein, D. Kouřil, P. Mindek, K. Schatz, M. E. Gröller, B. Kozlíková, T. Isenberg, and I. Viola. Multiscale molecular visualization. *Journal of Molecular Biology*, 431(6):1049–1070, Mar. 2019. doi: 10.1016/j.jmb.2018.09.004
- [36] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990. doi: 10.1109/34.56205
- [37] J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kalé, and K. Schulten. Scalable molecular dynamics with NAMD. *Journal of Computational Chemistry*, 26(16):1781–1802, 2005. doi: 10.1002/jcc.20289
- [38] H. Piringer, W. Berger, and J. Krasser. HyperMoVal: Interactive Visual Validation of Regression Models for Real-Time Simulation. *Computer Graphics Forum*, 29(3):983–992, 2010. doi: 10.1111/j.1467-8659.2009.01684.x
- [39] M. Schirski, T. Kuhlen, M. Hopp, P. Adomeit, S. Pischinger, and C. Bischof. Efficient visualization of large amounts of particle trajectories in virtual environments using virtual tubelets. In *Proceedings of the 2004 ACM SIGGRAPH*, pp. 141–147. ACM, New York, NY, USA, 2004.

doi: 10.1145/1044588.1044615

- [40] D. W. Scott. *Multivariate density estimation*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons, Inc., New York, 1992. Theory, practice, and visualization, A Wiley-Interscience Publication. doi: 10.1002/9780470316849
- [41] J. S. Simonoff. *Smoothing methods in statistics*. Springer Series in Statistics. Springer-Verlag, New York, 1996. doi: 10.1007/978-1-4612-4026-6
- [42] M. Sips, P. Köthur, A. Unger, H. C. Hege, and D. Dransch. A visual analytics approach to multiscale exploration of environmental time series. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2899–2907, Dec 2012. doi: 10.1109/TVCG.2012.191
- [43] 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
- [44] M. Spivak. *Calculus on manifolds: a modern approach to classical theorems of advanced calculus*. Mathematics monograph series. Perseus Books, Cambridge, Mass, 1998.
- [45] A. Stukowski. Visualization and analysis of atomistic simulation data with OVITO—the Open Visualization Tool. *Modelling and Simulation in Materials Science and Engineering*, 18(1):015012, 12 2010. doi: 10.1088/0965-0393/18/1/015012
- [46] T. Torsney-Weir, T. Möller, M. Sedlmair, and R. M. Kirby. Hypersliceplorer: Interactive visualization of shapes in multiple dimensions. *Computer Graphics Forum*, 37(3):229–240, 2018. doi: 10.1111/cgf.13415
- [47] T. Torsney-Weir, M. Sedlmair, and T. Möller. Sliceplorer: 1D slices for multi-dimensional continuous functions. *Computer Graphics Forum*, 36(3):167–177, 2017. doi: 10.1111/cgf.13177
- [48] J. J. van Wijk and R. van Liere. Hyperslice: Visualization of scalar functions of many variables. In *Proceedings of the 4th Conference on Visualization '93*, VIS '93, pp. 119–125. IEEE Computer Society, Washington, DC, USA, 1993.
- [49] V. Vuollo and L. Holmström. A scale space approach for exploring structure in spherical data. *Computational Statistics & Data Analysis*, 125:57–69, 2018. doi: 10.1016/j.csda.2018.03.014
- [50] M. P. Wand and M. C. Jones. *Kernel smoothing*, vol. 60 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, Ltd., London, 1994. doi: 10.1201/b14876
- [51] R. Wegenkittl, H. Löffelmann, and E. Gröller. Visualizing the Behavior of Higher Dimensional Dynamical Systems. In *Proceedings of the 8th Conference on Visualization '97*, VIS '97, pp. 119–125. IEEE Computer Society Press, Los Alamitos, CA, USA, 1997.
- [52] J. Weissenböck, B. Fröhler, E. Gröller, J. Kastner, and C. Heinzl. Dynamic Volume Lines: Visual Comparison of 3D Volumes through Space-filling Curves. *IEEE Transactions on Visualization and Computer Graphics*, Li:1040–1049, 2018. doi: 10.1109/TVCG.2018.2864510
- [53] L. Westover. Footprint evaluation for volume rendering. In *Proceedings of the 1990 ACM SIGGRAPH*, pp. 367–376. ACM, New York, NY, USA, 1990. doi: 10.1145/97879.97919
- [54] A. P. Witkin. Scale-space filtering. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'83, pp. 1019–1022. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1983.
- [55] J. Woodring, C. Wang, and H.-W. Shen. High dimensional direct rendering of time-varying volumetric data. *IEEE Transactions on Visualization and Computer Graphics*, pp. 417–424, Oct 2003. doi: 10.1109/VISUAL.2003.1250402