

# Physics-based Visual Characterization of Molecular Interaction Forces

Pedro Hermosilla, Jorge Estrada, Victor Guallar, Timo Ropinski *Member, IEEE*, Àlvar Vinacua and Pere-Pau Vázquez

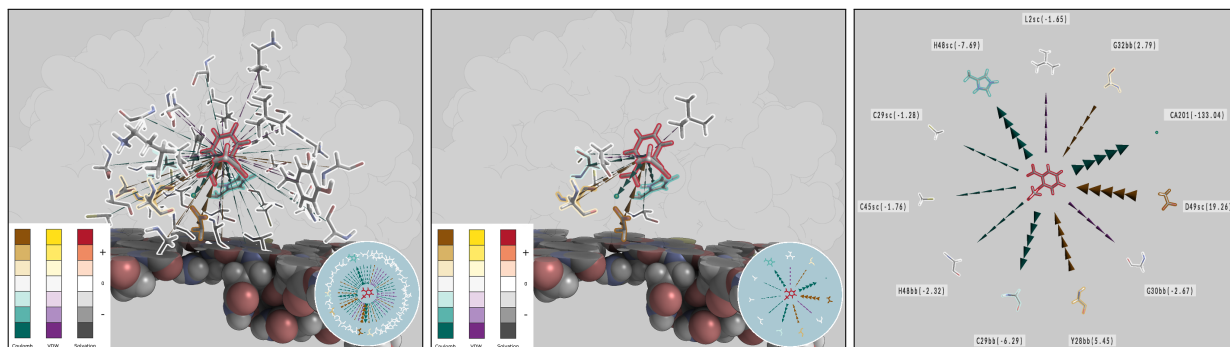


Fig. 1. Successive steps during the visual analysis of the binding of aspirin and the phospholipase A2 protein. We compute and visualize all essential interaction energies represented by 2D and 3D arrows. The orientation of the depicted arrows encodes the sign of the energy, i.e., attracting vs. repelling force. The width of the arrows as well as the color of the residue's silhouettes support energy quantification. During the visual analysis, energies are computed and depicted on-the-fly to support interactive hypothesis testing (*left*), and residues can be filtered based on energy and distance to obtain a more focused view (*middle*). Additionally, a 2D visualization helps to obtain total energy values in an uncluttered manner (*right*).

**Abstract**—Molecular simulations are used in many areas of biotechnology, such as drug design and enzyme engineering. Despite the development of automatic computational protocols, analysis of molecular interactions is still a major aspect where human comprehension and intuition are key to accelerate, analyze, and propose modifications to the molecule of interest. Most visualization algorithms help the users by providing an accurate depiction of the spatial arrangement: the atoms involved in inter-molecular contacts. There are few tools that provide visual information on the forces governing molecular docking. However, these tools, commonly restricted to close interaction between atoms, do not consider whole simulation paths, long-range distances and, importantly, do not provide visual cues for a quick and intuitive comprehension of the energy functions (modeling intermolecular interactions) involved. In this paper, we propose visualizations designed to enable the characterization of interaction forces by taking into account several relevant variables such as molecule-ligand distance and the energy function, which is essential to understand binding affinities. We put emphasis on mapping molecular docking paths obtained from Molecular Dynamics or Monte Carlo simulations, and provide time-dependent visualizations for different energy components and particle resolutions: atoms, groups or residues. The presented visualizations have the potential to support domain experts in a more efficient drug or enzyme design process.

**Index Terms**—Molecular visualization, binding analysis.

## 1 INTRODUCTION

Molecular design procedures, such as drug design and protein engineering, are complex processes, largely benefiting from computational resources but also from human analysis. In drug design, for example, a costly iterative loop involves simulations requiring long computation times, followed by a data analysis phase, which is conducted by domain

scientists using numerical analysis and visualization tools. Once some clues favoring or hindering binding have been understood, the ligand is modified by taking into account these clues, and another iteration is performed. In typical cases, computer simulation times range from hours to weeks, depending on the complexity of the molecules and the methods used. Recent advances have led to a reduction in computation time by exploiting modern supercomputers [8] in combination with modern simulation algorithms [34]. These advances in parallel computation thus allow for the computation of many simulation paths at the same time. This increase in available data also implies that the requirements for analysis tools become more demanding.

As a result of the computational resources becoming more affordable, human resources are becoming more and more the limiting factor in the computer-assisted molecular design process. In fact, most of the data analysis is performed in meeting rooms where different specialists discuss the outcomes of the simulation and the next design step. To enable these experts during comprehension and decision making, it is of great importance to provide effective data examination and visualization tools. In this paper we focus on one of the key aspects required to make informed decisions in the molecular design process: understanding which parts of the molecule influence the binding of the ligand. This information is key, as it enables the domain expert

- Pedro Hermosilla is with the ViRVIG Group and the Barcelona Supercomputing Center. E-mail: [pedro.hermosilla@cgstarad.com](mailto:pedro.hermosilla@cgstarad.com)
- Jorge Estrada is with the Barcelona Supercomputing Center. E-mail: [jorge.estrada@bsc.es](mailto:jorge.estrada@bsc.es).
- Victor Guallar is with the Barcelona Supercomputing Center. E-mail: [victor.guallar@bsc.es](mailto:victor.guallar@bsc.es).
- Timo Ropinski is with the Visual Computing Group, Ulm University. E-mail: [timo.ropinski@uni-ulm.de](mailto:timo.ropinski@uni-ulm.de).
- Àlvar Vinacua and Pere-Pau Vázquez are with the ViRVIG Group, UPC Barcelona. E-mail: [{pere.pau|alvar}@cs.upc.edu](mailto:{pere.pau|alvar}@cs.upc.edu).

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](mailto:reprints@ieee.org). Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

to hypothesize which residues can be altered in the subsequent design process in order to improve the ligand's affinity. Unfortunately, communication of this information results in several challenges. First, the binding information must be available instantly; for instance, if the domain expert moves the ligand, this information must be updated. Second, visual clutter, resulting from the multitude of displayed forces, which especially arises when considering long range forces, needs to be reduced. Third, domain experts must be able to identify the involved residues and to quantify the related energies. When considering the usually dense representation of complex molecules, it becomes clear that a single 3D visualization will not be able to meet all these challenges. Therefore, we combine 2D and 3D visualizations together with brushing-and-linking in order to communicate which residues influence the ligand. Furthermore, we propose how to perform real-time computations of the three main energy components, i.e., electrostatic, van der Waals (vdW) and solvation, which enables the domain expert to interactively explore entire trajectories consisting of multiple snapshots. Furthermore, with the proposed brushing-and-linking setup, it becomes possible for the first time to also analyze long-range interactions, which play an important role when a ligand is initially approaching a molecule. We not only hope that this long range analysis sheds new light on the entire docking process, but also expect that it helps to reduce required computation resources, as it allows for early intervention with the running simulation. Thus, we support a more effective, computer-based molecular design process by making the following main contributions in this paper:

- We propose visual analysis techniques for the real-time computation and inspection of interaction energies arising between a molecule's residues and the ligand.
- We propose a linked visualization setup communicating the computed interaction energies, by reducing visual clutter and enabling direct identification of the individual residues.
- We enable domain scientists through the means of brushing-and-linking to explore the underlying interactions, which in particular allows them for the first time to also inspect long range energies.

As illustrated in Figure 1, we combine these contributions, such that the user can interact with the input data to gather new knowledge, to formulate and assess hypotheses, and provide visual explanations of the discoveries.

## 2 RELATED WORK

In this section we will discuss work related to our approach. We first focus on visualization techniques used to communicate molecular structures, before we address more relevant techniques visualizing molecular interactions in 2D and 3D.

**Molecular structure visualization.** A lot of research has been conducted to improve visualization of complex molecular structures [23, 12, 53, 5, 18, 39]. As discussing all these techniques would be beyond the scope of this paper, we refer the interested reader to the state-of-the-art report by Kozlíková et al. [26], and only focus on those techniques directly related to the visualizations proposed in this paper.

Given the complexity of the molecular structures and the consequent increase in polygon count when atoms are represented as spheres or more complex shapes, effort was initially focused on their efficient rendering. The proposed algorithms span a wide range of underlying concepts, such as level-of-detail [29, 32, 38], impostors and glyphs [40, 51, 3, 15, 11] or geometric instancing [10, 33]. Often, geometry or tessellation shaders, provided by modern GPUs, are also exploited to create spherical or ribbon-based representations of molecules [27, 52, 21]. Furthermore, deferred shading is widely applied to accelerate shading of molecular models [16, 33, 11]. To realize a more intuitive communication, also modified shading algorithms are applied when visualizing complex molecular models. One of the most popular techniques used in this context is ambient occlusion [51, 28, 15, 25, 14, 20], but also contour lines [47, 29] or hatching [7, 31] have been applied.

**2D molecular interaction visualization.** Many of the modern visualizations proposed to communicate interaction energies exploit 2D views. For example LigPlot+ generates 2D views of ligand-protein interactions for a static frame [30], whereby some idioms are shown in the 2D view to illustrate interaction forces. Similar results can be obtained by using LeView [4], Maestro [46] or PoseView [50]. In all cases, the outcome is based on a close contact analysis with little or no interactivity. They do not allow filtering molecular interactions by strength or distance. This forbids, for example, the exploration of interactions outside the binding site (long range) and a comprehensive study of the binding process through residue interactions. Furthermore, the projection techniques provided make mental linking to the 3D structures difficult.

**3D molecular interaction visualization.** While nowadays, most molecular visualization techniques are able to communicate shapes and spatial arrangements at interactive frame rates, molecular interactions have received far less consideration in 3D despite their essential role in molecular design. Such interactions, for example, help to understand why a docking simulation succeeds or not, so their visualization and biophysical basis are essential for finding effective solutions. Also, no previous 3D interactive framework provides filtering by strength or distance. Visualizing the physics is also paramount in related fields, such as materials science, where Grottel et al. [13] represent scalar fields through color overlaid on the molecular surfaces, and represent electrostatic dipoles with arrows. Similar representations using iso-surfaces are common to other crystallography applications such as Vesta [35, 36], where the overlaid color represents electrostatic potential or hydrophobicity. Hyde [45] uses a similar representation to code the total affinity energy contribution, whereby the process takes several seconds. Cipriano and Gleicher [6] illustrate charges over the molecular surface by stylizing both the surface shape and the charge values. Günter et al. also focus on the atom level, where the signed electron density and reduced gradient fields are computed and then simplified to illustrate van der Waals and steric repulsion forces between atoms [17]. By showing colored dots at the contact surface between atoms, representing the van der Waals and hydrogen-bonding interactions, Word et al. allow evaluation of atom packing in biomolecular structures [55]. Another tool, LigandScout, exploits several views to support drug designers when screening chemical databases [54]. It allows for the interactive creation of so called pharmacophores (which are based on known ligands) that act as templates for finding new ligands. In the creation process, LigandScout highlights the ligands key features that interact with the protein, and supports surface coloring based on lipophilicity, hydrogen bonding or charge, using predefined scoring functions. The PLIP system is a web service that generates 3D views focused at the atom level and showing several interaction types [41]. However, as discussed later, no interaction nor filtering are supported and it is not designed to deal with a sequence of frames. More recently, Skanberg et al. have proposed to visualize energy interactions between atoms through diffuse interreflections computed for the surfaces of the atoms [49]. Falk et al. visualize molecule reactions by means of arrows augmenting paths representing molecule trajectories [10]. Khazanov and Carlson exploit tables to communicate molecule interactions [24]. They also communicate the interaction between ligand and binding site through modification of color and van der Waals radii on an atom scale, and indirectly address the residue scale by performing this depiction individually for each amino acid. Sarikaya et al. also take into account the residue scale, by visualizing classifier performance with respect to protein chains on which a classifier has operated [42]. Finally, to communicate the differences of surface projected parameters, Scharnowski et al. propose to use deformable models [44].

Unfortunately, these tools are restricted to the few close strong interaction contacts between molecules, not allowing the detection of long range interactions. These might, for example, guide the entrance of the ligand to the binding site (electrostatic steering) or, by adding up several weak interaction energies, contribute to the overall binding affinity. Although they deal with the main biophysical types of interactions, usually in a geometric and quantitative way (hydrogen bonds, van der Waals interactions), these tools do not allow a detailed

quantitative understanding of all the energy terms (including solvation) on which those and other interaction types are based, and which are the fundamental parts of the energy functions that guide the molecular simulations. Moreover, without a dynamic view of the binding process, these tools hinder a deeper understanding of the key actors (residues) and their roles in the process. Our aim is thus to enable this fine grain data analysis at the meaningful residue level, both for single frame studies and for the dynamic setting represented by a sequence of frames, and with filtering tools to allow the exploration of the different energy terms and interaction distances.

### 3 APPLICATION-DRIVEN VISUALIZATION DESIGN

In this section, we will provide the background for our application-driven visualization design. Before we analyze the visualization design requirements in Section 3.2, we discuss the application background.

#### 3.1 Application Background

In computational drug design, as well as in other molecular modeling areas, a key aspect is to estimate the interaction energy between the protein and the ligand (or substrate in, e.g. enzymatic catalysis).

The free energy  $G$  of the protein-ligand system is a powerful tool to understand the binding process and, by computing  $\Delta G_{\text{bind}}$ , the difference of free energy between the bound state and the unbound state where the protein and ligand stay free in the solvent, we can know the binding strength or affinity (more negative values of the energy mean a stronger binding). Typical energy models are additive, allowing to understand the main contributions to binding or the key interactions that would favor (or disfavor) ligand binding.

The energy model we follow is based on three terms:

- Van der Waals interaction energy,  $E_{\text{inter,VDW}}$ , which shows how well the protein and ligand molecules pack together.
- Electrostatic interaction energy (the interaction in vacuum  $E_{\text{inter,ele}}$  plus the solvent screening) which shows the strength of the interaction between the protein and ligand charges, screened by the effect of the solvent.
- Change in solvation energy, which includes a polar ( $\Delta G_{\text{pol}}$ ) and a non-polar term ( $\Delta G_{\text{np}}$ ). It shows how much the protein and ligand prefer to be bound together, instead of being free in the solvent, regarding exclusively the interaction with the solvent molecules, and including entropy.

So the energies computation results in

$$\Delta G_{\text{bind}} \approx E_{\text{inter,VDW}} + E_{\text{inter,ele}} + \Delta G_{\text{pol}} + \Delta G_{\text{np}} \quad (1)$$

To compute the energy of a given protein-ligand conformation the application needs the atomic 3D coordinates and a table of parameters per atom type (such as the van der Waals radius or electrostatic charge) as input. Besides the ligand atoms, which are typically user-provided, these parameters are standard for the models used. Refer to the supplementary materials for details on the computation of the energy terms. Here it suffices to know that the computation has a cost of  $O(N^2)$ , where  $N$  is the number of atoms.

When representing the electrostatic interaction energy, we calculate the sum of the actual  $E_{\text{inter,ele}}$  plus the screening due to the  $\Delta G_{\text{pol}}$  energy term assignable to the specific studied interaction. Interaction energies (the electrostatic ones including the screening) are thus represented for given interacting pairs of atoms or residues (one atom or residue from the protein, and the other from the ligand, considered as a single residue). However, we assign solvation energies to individual atoms or residues. Since some terms used in the calculation of the solvation energies involve a pair of atoms, we have assigned half that energy to one of the atoms and the other half to the other atom.

#### 3.2 Design Requirements

In this paper, we propose visualizations which have been developed with the goal to help domain experts understand the forces acting in molecular design, through analyzing the main components of the binding energy. In drug design, for example, domain experts need to analyze whether the ligand will or will not dock at the intended position. To answer this question, domain experts have to inspect the numerical results of the main interaction energy components which are usually provided in result tables. While such a table could be analyzed for a single ligand, comparing them for several ligands is not practical. However, since it is often necessary to study more than one bound structure, as an ensemble of structures will aid in a better characterization of the bound complex, domain experts need to be able to effectively analyze this type of data. In addition, studying the drug migration pathway, from the solvent to the bound complex, might better help in addressing the binding (or its absence) mechanism and locate key interactions that could facilitate (or hinder) binding [9]. Similar conclusions can be observed in enzymatic catalysis [1]. The detailed mechanistic knowledge provided by the binding energy analysis should locate those parts of the molecule/receptor that enhance or prevent docking, facilitating the following design steps.

The interaction of molecules is atom-based, but atoms group naturally into residues or chemical groups accountable of collective responses —such as ionic groups (carboxylic groups, etc.), aromatic groups (phenyl, etc.) or an entire residue— for which one is often interested in the whole group interaction. This is the reason why we analyze interaction energies in chemical groups and/or residues, not single atoms.

In a cooperation between domain researchers and visualization experts, we have identified the following questions as being essential to be answered in the visual analysis process:

- Q1: Which are the most active groups in the interactions between molecules?
- Q2: Which are the most powerful binding energy components at a certain simulation configuration?
- Q3: Is the proximity of the drug causing instability in any residue of the protein?
- Q4: Is the ligand solvation force favouring or rejecting binding?
- Q5: Which residues (if any) prevent drug delivery?

By creating a visualization method that illustrates the different components of the molecular interaction energies, we enable domain experts to answer these and other questions, thus gaining a detailed knowledge on the binding mechanism. Importantly, our system not only computes the energy components on the fly, but also provides a series of filters that let the user select distances or energy ranges to inspect them in fine detail. To support such an interactive visual analysis, besides effective visualizations, an efficient implementation is essential. Therefore, we exploit a data structure computed by the GPU that facilitates queries such that residues and groups can be filtered in real-time for all elements of interest. The thus filtered information is then communicated using the proposed visualization techniques. This way we can provide an interactive visual analysis system, that lets the user inspect the set of energies that are interacting at any time in a simulation. In the following sections, we will describe this system by focusing on the two main parts:

**Interaction energy calculation:** To compute the interaction energies on the fly, a GPU-based set of programs is used that computes an array of energies for each residue (and for each atom in the case of solvation energy), each time a new step of the simulation is selected (or a structure is loaded). Furthermore, a specialized energy data structure is used to accelerate the queries issued by the domain expert through the visualization front end.

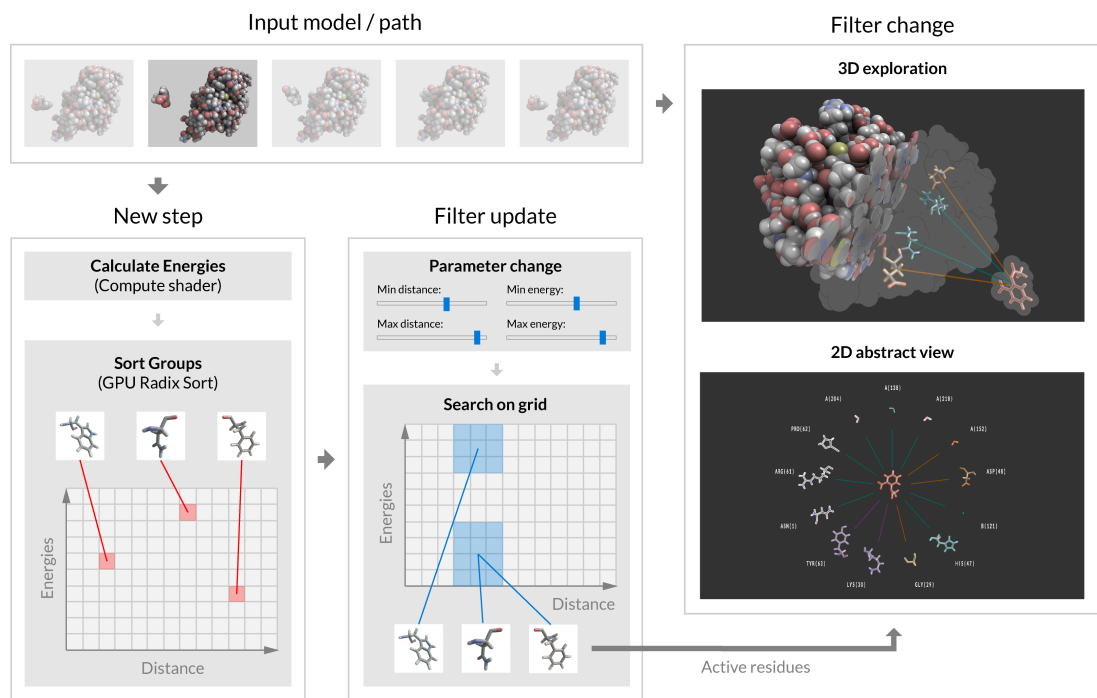


Fig. 2. Overview of the data flow underlying our application. When the current step or the configuration change, the system automatically computes the forces being exerted by each residue in a compute shader (*left*). When the user modifies the filters, the selected residues are quickly gathered from the indexing data structure computed by the compute shaders (*center*) and then used to render the 3D or the 2D abstract views (*right*).

**Interaction energy visualization:** During the visualization phase, the domain expert can interact with the data by means of widgets. These provide several visualization motifs and filters that facilitate the data analysis and inspection, as well as data presentation.

#### 4 INTERACTION ENERGY CALCULATION

Given a configuration of the molecules, the computation of the energies involved is not inexpensive. Some of the quantities involved—the Born radii and the per-atom solvation energy change due to interacting charges ( $\Delta G_{pol}$  in Equation (1))—have a cost which is quadratic in the number of atoms. Since we want our application to be able to adapt to changes in the configuration (because of user interactions or a change of frame in a simulation trajectory) interactively, we employ GPU computation at each change. Therefore, when the user selects a new step of the simulation, first we dispatch the computation of the Born radii to a compute shader. In order to compute the Born radius of each atom, we must visit all other atoms in the molecule, yielding a quadratic cost. Each thread of our compute shader calculates 16 of these interactions and adds the result to a shader storage buffer using atomic operations. After computing the Born radii, our system computes the solvation energy term for each atom. Again, this requires  $N^2$  separate operations, so we execute another compute shader to perform this task—assigning 16 of these computations to each thread as in the Born radii calculation. When all of these parameters have been computed for all atoms, we launch another compute shader that computes the VDW and Electrostatic terms and the final value of the energy between the ligand and the atom groups. With this strategy we are able to compute the energy of the system more than 20 times per second for molecules up to 18k atoms—using a computer with a processor Quad Core i7 at 3.7 GHz, 16Gb of RAM and a GeForce GTX 980.

Once the energies have been computed for each group, and to enable the interactive modification of filters by the user and their reflection on the visualization, we build an auxiliary data structure that allows us to find the selected groups in constant time. Each time the configuration changes, or the user shifts his attention to a different aspect—turning on or off some energy component—the array holding the energies of all the groups is sorted according to a key made up of the sum

of the chosen energies and the distance to the ligand. Both energy levels and distances are discretized into 128 possible segments with an energy range of  $[-3.0, 3.0]$  and a distance range of  $[0, 50]$ . With those parameters we assure a fixed number of groups per cell for all the examples we have tried. Then the sorting key of each group is constructed concatenating the energy and the distance bucket index (this is represented by the bottom portion of the “New step” box in Figure 2). The sort is carried out using a radix sort implemented in four compute shaders corresponding to the three steps in [19] plus an additional step to build an index table that holds in each cell the start and end position in the buffer of the entries corresponding to that cell. The table is indexed by the same 128 intervals of energy or distance.

Once the sort is complete, the results obtained are downloaded to the CPU. When the user modifies some of the filter ranges (“Filter update” box in Figure 2), the cells containing the chosen ranges are determined, and using the index table, all the groups in those cells are checked for the current filter ranges (as the cells will contain some neighbor values as well, because of the discretization). The groups thus filtered are made visible, and are given an identifier according to their energy level, to determine the appropriate colors of the silhouette. The start and end points of the connection arrows are also uploaded to the GPU with the same identifier that will determine their color.

#### 5 INTERACTION ENERGY VISUALIZATION

The usual workflow of the users is this (see Figure 2): They open a file that contains a structure or a simulation path. At this point and every time the user switches between path frames (“New step” box), the GPU updates the energies-distances data structure and downloads it to the CPU. Then, the user may freely update the filters. When those are changed, the data structure is queried (“Filter update” box) to determine active groups, energies, and so on. The user can then freely inspect the 3D and 2D views or update the filters again. All of this happens in realtime. In this section we describe how the visualization tool has been designed and implemented.

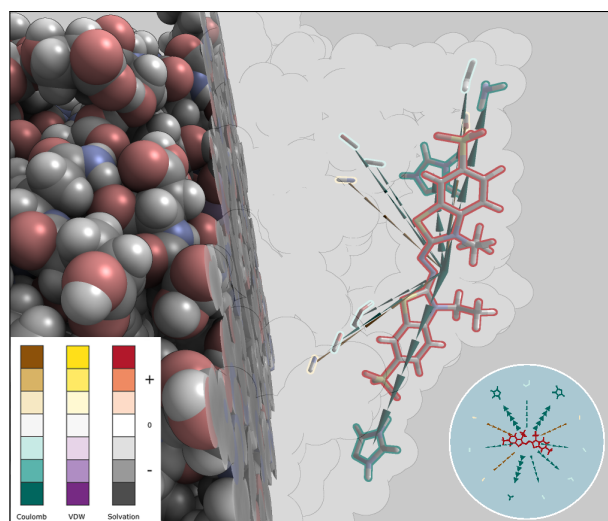


Fig. 3. The snapshot corresponds to the docking position of an artificial substrate, ABTS, to the manganese peroxidase 4. In Figure 4 the docking path is illustrated. This zoomed view shows the main idioms used to communicate energy: cones for direction of forces, their thickness and color to encode energy and intensity, and colored highlights of residues as described in Section 5.

## 5.1 Idioms and Filters

In order to support domain experts in answering the questions listed in Section 3.2, we have developed interactive visualizations communicating the real-time computed binding energy factors. By employing filtering, we can ensure that only residues currently of interest appear in the view. The visualization of these elements is then enhanced with visual idioms that provide information so that domain experts can easily understand the important forces in the current step of the simulation path. We have realized this interactive visual analysis by including the following visual idioms: filtering, focus and context, feature enhancement, and interaction. Most of them are illustrated in Figure 3

**Filtering.** In several cases the interaction energies are omnipresent, despite the fact that for several groups the absolute energy is rather low, e.g., the ligand will have electrostatic interactions with most of the groups of the protein. If all these energies were communicated, our visualization would suffer from cluttering. Therefore, we support filtering to restrict the energies and the groups to be visualized to those which fulfill certain criteria. Currently, we support three types of filters: *i*) distance filtering, *ii*) energy level filtering, and *iii*) energy type filtering. Thus, by selecting any (or all) of the energy types, changing the range of distances at which interactions are considered and the amount of energy, the user can finely analyze individual or group interactions. In Figure 1 we have applied distance filtering to confine the visualized residues to those having a short distance interaction with the ligand only.

**Focus and context.** The central entity that guides any exploration within our application is the energy enhancement. As a result, we always enhance the groups that are active, i.e., whose energy is between the limits determined by the used energy level filter. However, to embed the currently selected groups, it is important to add context to the focused elements. Therefore, we provide the visual context in two flavors: *i*) Eliding information using a user-defined clipping plane, and *ii*) Superimposing a layer that renders, using semi-transparency and silhouette enhancement, the information concerning the non-active groups. We allow the user to define the plane direction by selecting the current view orientation, and, once the plane direction has been fixed, the user is able move the plane along it using a slider. Moreover, the parts of the molecule that are hidden by the plane, are then rendered using the semi-transparent layer.

**Feature enhancement.** In order to facilitate comprehension of the affecting energies, we color code the dominant type of energy on the

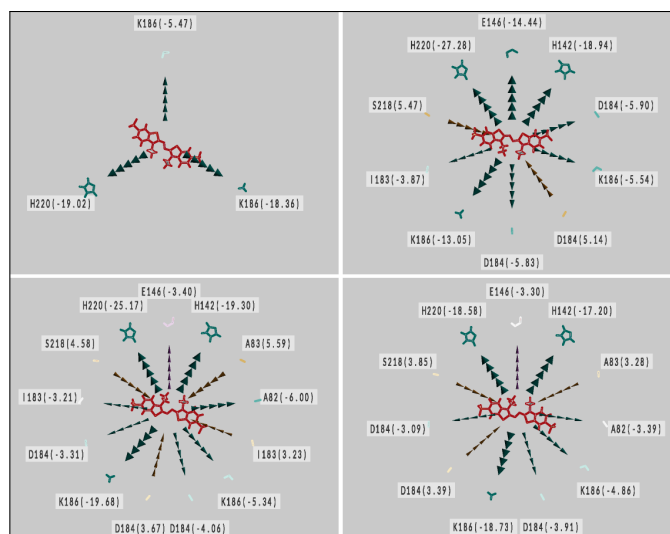


Fig. 4. The interaction through the 2D views (top left, top right, bottom left and bottom right indicate four steps of the interaction) visually explains the docking procedure of the ABTS, an artificial substrate to the manganese peroxidase 4. Once the histidine 220 (H220) has established an attraction connection (top left), the substrate does not leave the surface of the protein and finally docks (bottom right) also attracted by another histidine (H142). This can be seen interactively by hopping between path steps.

silhouettes of the respective groups (see Figure 3), where each energy type is coded using a color scale with two different hues—which represent the positive and negative ranges. Further communication of the energy and sign is provided by means of arrows, drawn along the axis that links the ligand with the active group. These cones are color coded with the hue of the energy type, and indicate the direction in which the ligand would move due to a single interaction. The size of their base also encodes the amount of total energy, so harder interactions are more likely to stand out in cluttered scenes where many groups are active. For an individual energy analysis or only repulsion/attraction analysis, the user could also modify the application to only highlight the total energy using a color scale with two hues, where one hue represents the positive range and the other the negative.

**Interaction.** During the whole visual analysis process, the user may freely inspect the 3D view by modifying the viewpoint, zoom, pan, and so on. Filters can also be changed interactively. And the path step can be manually selected or an animation can be triggered to see the full path when analyzing a docking simulation path. To support a more exact quantification, a 2D view, with details on the amino acids is also provided. This view is created by projecting all active residues in a circular layout around the ligand. By exploiting linking, we ensure that when filters are updated or the step of the path changes, the view is updated accordingly. Figure 4 shows the last steps of the docking path of an artificial substrate.

## 5.2 3D Visualization

To provide visual cues for the users to understand the energies involving each residue at any point in the simulation path, we need to visualize both, the binding energy and the elements affecting it. Therefore, we choose a default representation that lets the user identify the groups and facilitates the incorporation of energy information around the molecules.

**Atom representation.** We chose to render the active elements using a licorice visualization, with thin cylinders and atoms to reduce the footprint on the screen. This representation is a good balance between space occupied and information. Since the cylinders are coded with the traditional colors of the atoms, they are easy to identify, and additional information is communicated via a thick silhouette around them. To



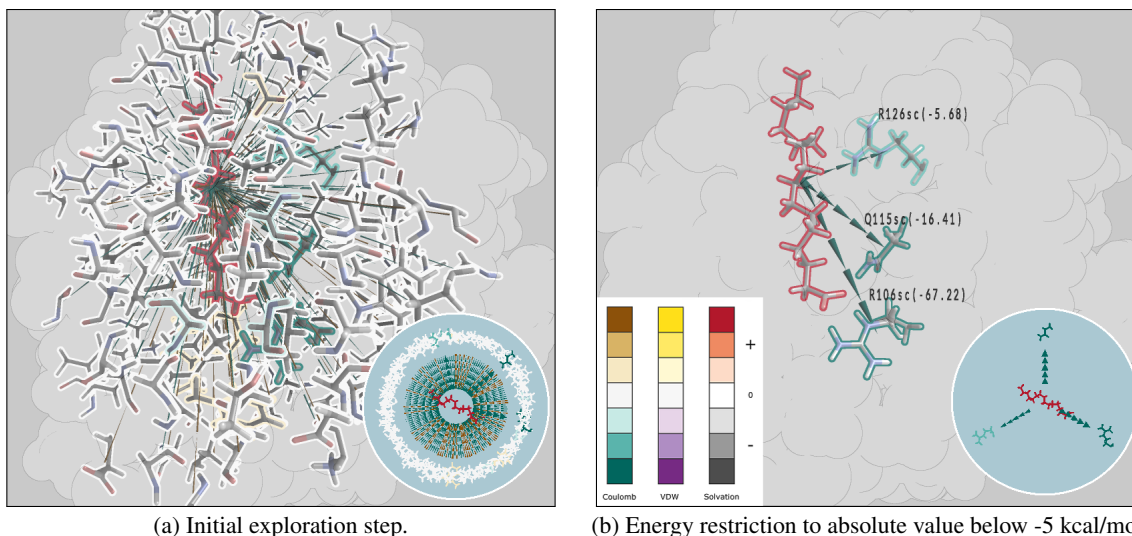


Fig. 5. Interaction between the palmitate ligand and the intestinal fatty acid protein. Initial exploration does not let us see the important interactions happening close to the palmitate acid. By carefully filtering energy terms larger than -5 kcal/mol, we can see how palmitate interacts strongly at electrostatic level (dark green interactions) with several residues, notably with arginines (R126 and very strongly with R106) as shown in the image. The values between parentheses indicate the total energy.

reduce clutter, only the active residues or groups are shown, and to provide 3D context, the rest of the molecule is visualized using semi-transparent van der Waals surfaces.

**Energy representation.** Energies are typically signed, so we will use diverging hue representations to show them. The hues for the range selections fulfill the following requirements: *i*) Avoid blue-red hues, since these are commonly used for polarization, *ii*) Avoid the colors commonly used to represent proteins (e.g. some grey, red and blue hues), *iii*) enhance color distinction by reducing the amount of tones and using perception-based selections, *iv*) avoid the repetition of hues to make them unambiguous. Under these conditions, we decided to represent the dominant energy of each interaction with a 7-point diverging hue scale where white represents neutral or close to neutral values.

We selected green-brown for electrostatic energies (typically spread in both sign directions), violet-yellow hues for van der Waals energies keeping the violet hues for negative, far more common in vdW than positive values. Finally, a grey-desaturated red scale was used for solvation energy. The solvation energy is encoded in the color of the silhouette of the ligand, since it is an important information that communicates whether the ligand is comfortable in the solvent or uncomfortable, which might favor binding. Although having the same neutral color for all the energy types might seem confusing, the energies represented by these colors are not relevant for the understanding of the simulation as their magnitudes are small (Fig 5). This guides the attention of the domain experts to the groups with high magnitudes, as they are represented with highly saturated colors. These magnitudes are also communicated through other more precise means (see Visualization Configuration below, and the 2D view described in Section 5.3).

In all cases, the energy magnitudes are rendered as a thick silhouette around the licorice representations of the active bonds. These color combinations and molecular representations result in a quite understandable way to encode interactions, facilitating the comprehension of simulation results. The main idea behind this is to avoid the requirement of checking other regions of the screen (tabular representations of values are also commonly used) and thus keeping the attention of the user onto the task.

**Visualization configuration.** To further guide the attention of the users to active groups, we also highlight the interactions with geometric elements that go from the center of the ligand to the center of the group of interest. The user can freely toggle these elements on or off. These cones indicate the direction the ligand would move as consequence of

the influence of the corresponding group of interest. The cones have base areas proportional to the total energy level.

### 5.3 2D Visualization

When generating a 2D projection of the groups of interest, it is important to facilitate an easy mental linking with the 3D visualization. As a consequence, our proposed algorithm for generating the 2D visualizations takes into account the 3D arrangement, and exploits the same connections (cones or cylinders) as used in the 3D view, as well as the same color coding for the silhouettes of the groups. The individual steps of the algorithm can be summarized as follows:

1. Calculate the vector that goes from the center of the ligand to the center of the scene. This vector is then used as the direction of the virtual viewing plane of step 4.
2. Determine the number of active groups.
3. Subdivide the virtual space around the ligand in as many equal sectors as there are active groups.
4. For each group, calculate the projection to a virtual viewing plane centered in the ligand, and exploit clockwise sorting to assign the respective partition.
5. Project the residues at a fixed distance from the center, and centered in their sector, maximizing the projected area.

To facilitate the interpretation of the projected groups, we project each residue to the 2D view with a (different) projection direction that maximizes its area. The optimal direction is achieved by performing a Principal Component Analysis of the group's atoms' positions, whereby the smaller eigenvector of the matrix determines the optimal projection direction. With this strategy, we achieve an ordering of the residues that is directly related to their 3D position in space, thus facilitating the inspection in both 2D and 3D views at the same time.

### 6 APPLICATION CASES

The proposed visual analysis techniques have been integrated into an application which is flexible and offers a large range of features, and thus allows to analyze different aspects of biomolecular interactions, with applications, for example, in drug design and protein engineering processes. We can analyze data from the point of view of the agnostic




Fig. 6. Interaction of the palmitate ligand with several leucines and phenylalanines residues at vdW level. Note the strong yellow silhouettes that indicate strong interactions.

scientist, just trying to gather new knowledge, or we can use it to assess some hypothesis. Usually, hypothesis testing will lead to simpler scenarios because we already have an initial guess about which parameters to analyze. In the following subsections, we will discuss application cases describing how the presented visual analysis techniques enable new insights. We will start by discussing a single conformation analysis process, before discussing the insights achieved when applying our approach to a more complex trajectory analysis.

## 6.1 Single Conformation Analysis

An initial scenario where we can use our visualization is to understand which molecular forces are predominant at a given structure, such as the bound conformation obtained from measurements or simulations. The proposed visualizations enable to spot key residues and chemical groups in the interaction between protein and ligand, both enhancing or opposing binding. This information is paramount for the scientist in order to understand the binding mechanism, as well as to choose ligand groups in drug design or protein amino acids in enzyme engineering for mutation, whereby the mutations are performed with the goal to improve or disfavor molecular interaction. Suggested changes can later be confirmed or discarded by analyzing a new simulation with the modified protein-ligand system.

With the following example, investigating the binding of palmitate, a fatty acid ligand, to the intestinal fatty acid-binding protein, we further illustrate how we can quickly assess molecular interaction hypotheses from a crystal structure inspection. The bound structure was obtained from a Nuclear Magnetic Resonance Spectroscopy (NMR) experiment (PDB id 1ure, [22]), after minimization with the OPLS-AA force field using the PELE software [34]. This system is interesting since palmitate's binding could show important contributions from the electrostatic, vdW and solvation energy terms, and our analysis is focused on question Q1, asking for the most active groups. Thus, the first hypothesis is that several arginine (Arg or R) residues, which are positively charged amino acids, should have an important role at the electrostatic level because they are in the vicinity of the negative polar extreme of the fatty acid, a carboxylic group. We can verify this assumption in Figure 5, where we checked all Coulomb contributions lower than -5 kcal/mol. Through the visual analysis, three arginines having large electrostatic stabilizing contributions can be clearly identified, whereby the strength




Fig. 7. The solvation term in the palmitate ligand, when inside the protein, indicates that, contrary to the hypothesis, it does not favor binding. Note the high energy value color coded in its silhouette. Being a fatty molecule, the expected value would be low (grey).

directly correlates with the cone radii. This is also seen in Figure 5, where the numerical value for each contribution indicates that Arg106, the one closer to the polar ligand group, is the main stabilizing residue.

A second hypothesis that can be easily tested is the nature of some of the vdW interactions. Since palmitate has a long aliphatic tail, some hydrophobic residues should have important vdW interactions, such as leucines (Leu or L) and phenylalanines (Phe or F). We can assess this by inspecting the vdW energies as shown in Figure 6. Here, several contributions can be identified, and we can see the interacting residues, i.e., leucine side chains (sc) and backbones (bb) (L36sc, L38sc, L72bb, and L72sc), and phenylalanine side chains (F62sc and F55sc). Interestingly, we can observe a strong destabilizing vdW component from arginine 106 (R106), induced by the large ionic attraction seen above. This example constitutes a nice (didactic) illustration of force field terms balance, and is directly related to our question Q2 asking which ones are the most powerful binding energy components.

We can see another example of how the visualization can help to confirm or reject hypothesis. In this case, related to question Q4, questioning the solvation force, the fatty acid has a long aliphatic chain. Consequently, it is expected that removing it from a water environment, which is polar, and placing it in its bound protein conformation would be associated to a (de)solvation energy gain. However, when using the visualization to assess the results, the domain experts were surprised by a desolvation loss instead. As a consequence, further detailed study was triggered and then it was discovered that the charged Carboxylate group in the ligand actually opposes this, and the total effect is a desolvation energy loss. This finding could be made by referring to Figure 7, where all the energies are removed except the solvation force that is encoded in the silhouette of the ligand. We can observe that the energy is strongly positive, and can thus reject the hypothesis that solvation energy terms would facilitate palmitate's binding.

## 6.2 Trajectory Analysis

A more complex scenario involves the analysis of multiple structures, obtained, for example, from molecular dynamics or Monte Carlo Simulations. We can run the whole ligand migration path, asking the application to highlight residues with dominating interaction energies at each frame of the simulation by simply stepping through simulation time. Note, that the term frame usually refers to a step in the simulation performed for the study of the protein-ligand binding. Similarly, multiple experimental structures can be analyzed simultaneously.

Figure 8 shows three different snapshots along the aspirine migration simulation in the phospholipase A2 protein. The top left image shows the ligand in the bulk solvent, far from the protein surface, where no energies acting on the ligand are detected. The top right image illustrates the ligand approaching the surface and how our visualization shows the initial protein-ligand recognition forces, two long-range electrostatic contributions that guide the ligand towards the protein. Coulomb

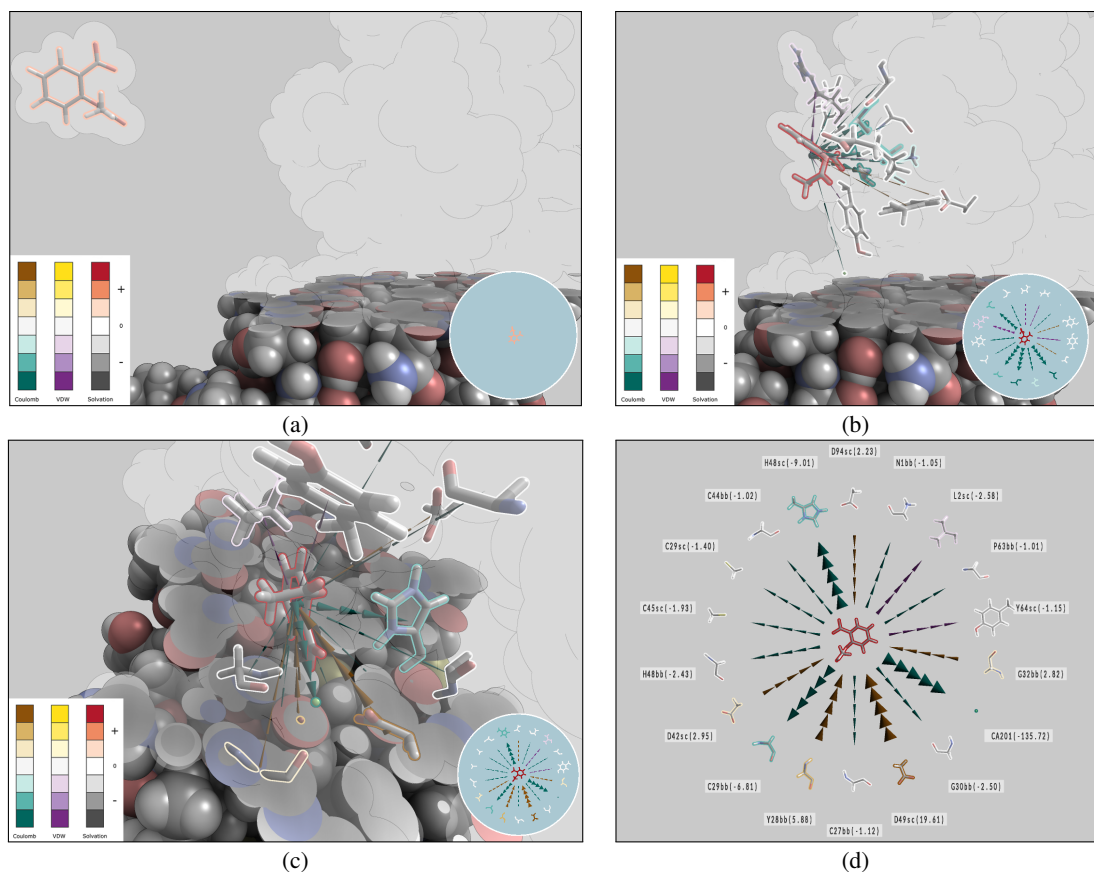


Fig. 8. Three different stages of a path of aspirin docking to the phospholipase A2 protein. The top left image shows the aspirin in the bulk solvent, so no energies are exerted. The top right image shows the ligand closer to the protein, thus electrostatic energies appear, and finally, the bottom left image shows a frame near the docking position. At this point, the calcium is strongly attracting the aspirin (green thick cones), where other groups exhibit repulsive electrostatic energies (brown cones). Bottom right shows the abstract view with all the active residues uncluttered.

forces are dominant usually when the ligand is relatively far from the protein. This is one example of question Q1 mentioned earlier. Here we also observe again the key stabilizing role of the calcium ion (green cones) together with some minor destabilizing electrostatic contributions (brown cones) from other calcium coordinated groups (having the same sign as aspirine). This partially deals with our questions Q2 and Q3, asking for the most powerful binding energy component and about drug-induced instability. In addition we can clearly observe some vdW smaller interactions from hydrophobic residues with the aromatic group of the aspirine ligand. This addition of van der Waals forces only appears at short range interactions. Finally, the bottom left image shows the ligand close to the docking position. We observe a strong interaction between the aspirin ligand and a calcium ion associated to the protein.

The crystal structure of the complex (PDB id 1oxr, [48]) shows, from a structural perspective, that this interaction exists. Our application enriches this information from an energetic point of view, as well as it allows to relate, in a qualitatively manner, the strength of this interaction with that of the aspirin to other important interacting residues, such as histidine and aspartic acid residues. The calcium ion is clearly visible as the green sphere in the middle of the protein. An inspection of the abstract view (bottom right) reveals all the interacting residues, and we can see clearly how histidine (H48sc and H48bb) and aspartic acid (D42sc, D49sc, and D94sc) residues are also interacting, as predicted. It can be further conclude, that some of these strong electrostatic attractive contributions, importantly, from the active site calcium (Ca) ion (with and overall +2 charge), are responsible for driving the ligand to its final bound position. Thus, this study constitutes a nice example of how this tool allows studying the ligand binding mechanism at atomic detail.

As we could show in the discussed application cases, the proposed

visual analysis techniques can be used to answer the stated questions Q1 to Q4. Unfortunately, we could not discuss any application case, where we could answer question Q5, asking for residues preventing drug delivery. This is due to the fact, that the simulation data analyzed in this paper has already been analyzed before with conventional methods. Accordingly, only those simulations resulting in a successful binding were at our disposal.

### 6.3 Evaluation

The development of the system has been done in close collaboration with domain experts. However, we also asked other experts external to our team in order to gather opinions on the design and implementation of our visualization tool. In particular, we asked the opinion of two more experts: a chemist that works in protein engineering, and a computational chemist working in computer aided drug design. After a small demo of the tool, we let them play for thirty to forty minutes with it. During and after the session, we gathered their comments.

Both of them found the tool very useful, with large applicability in protein engineering tasks (e.g. to inspect the result of protein mutations and analyzing protein-protein interactions), molecular dynamics (to better understand the molecules behavior), electrostatic steering (in order to analyze the effect of electrostatic fields in guiding the ligand's path), and also enzyme engineering. Moreover, they also found the application useful for results presentation, to help other collaborators "understand what is really happening". They also commented on the interface, where they referred to it as "very intuitive", since the interacting groups are easily identified. Finally, the semi-transparent visualization of the molecule was deemed very interesting, since it helps grasping how the ligand is getting close to the molecule without rendering it completely.






Fig. 9. Visual results obtained from Maestro (a) and PLIP (b) when using the systems to visualize the binding of aspirin and the phospholipase A2 protein. In comparison to our system, no linked views, no interactive exploration and no long range interaction analysis is supported.

From the design point of view, the use of cones of different sizes was found appealing because the users can identify the interactions and infer their strength quite simply. Furthermore, the domain experts valued very positively the possibility of filtering the energy magnitudes, energy types, and distances.

#### 6.4 Comparison with other methods

The benefits of our system stand out when we compare our results to the ones obtained by two popular systems that visualize molecular interaction forces, Maestro [46] (see Figure 9 (a)) and PLIP [41] (Figure 9 (b)) with the same simulations. Besides the different visualization design communicating interaction intensity and signage, three conceptual differences immediately become clear. First, while our system employs linked views in order to communicate the interaction forces, Maestro and PLIP use a single view only. Interestingly, Maestro uses a 2D visualization, while PLIP uses a 3D visualization. We see this as an indicator, that both 2D and 3D have their benefits, and that our design considering the linking of a 2D and a 3D view might be helpful. Furthermore, we believe that having two linked views makes it more intuitive to map the binding information to the 3D structure of a molecule. The second major difference is the consideration of time-varying data. While Maestro and PLIP show a single frame only, our system enables domain experts to interactively explore arbitrary frames in a time-varying data set—an essential prerequisite for understanding protein dynamics. Finally, in contrast to Maestro and PLIP, we support a full interactive analysis. Thus, the outcome is not only based on a close contact analysis, but molecular interactions can also be filtered by strength or distance, and visualizations are updated in real-time. This enables for instance for the first time the exploration of long range interactions.

The supplementary material shows larger views of those two software packages that include all the text that usually appears in the working window.

#### 7 CONCLUSIONS AND FUTURE WORK

In this paper we have presented visualization concepts developed for the analysis of binding forces in drug design and protein engineering. The proposed visual analysis workflow provides domain experts with several tools that let them perform a detailed analysis of the most relevant energies that intervene in a docking simulation: electrostatic energies, van der Waals energies, and solvent energies. This way, it becomes possible to gain an understanding on how simulations perform, why a ligand is getting or not getting to the docking position, and which are the residues crucial to such reactions. We have discussed the interaction computations used, outlined how to visualize the obtained results, and showed how filtering can help in the analysis process. We have further shown examples where domain experts can see at a glance which are the dominant energies throughout the simulation, or how easy it is to determine the strongest interactions at a certain point of the path. Through the filtering, the system instantly highlights the important residues and the energies can be seen as color-coded to indicate their relevance. Thanks to our GPU-based energy calculation,

which evaluates the binding equations in real-time, we may extend the application to any simulation path generated by other systems. The only necessary changes would be in the input files. In contrast to current approaches, our approach is able to deal with full simulation paths, instead of only single frames. Moreover, we may deal with information by residue, as well as per atom. This facilitates the analysis, since most forces are exerted at residue level. Moreover, we have a set of widgets that further facilitate the inspection, such as user-defined colors and transparency, configurable clipping plane, high quality illumination, and different rendering modes to add contextual information to the selected view.

In the future, we see several opportunities for future research. One direct extension could be to also consider regions of the ligand, and show the respective energies per region. The ligand could also be broken into groups, to separately show the interactions of its different parts. While this would increase the complexity of the visualization, it would also allow for a more detailed inspection. Furthermore, we plan to extend our visualization approaches to other molecular simulation problems.

#### ACKNOWLEDGMENTS

This project has been supported by TIN2014-52211-C2-1-R by the Spanish Ministerio de Economía y Competitividad with EU FEDER funds. We thank EAPM group members (BSC) and especially Sandra Acebes, for help in data preparation.

#### REFERENCES

- [1] S. Acebes, E. Fernandez-Fueyo, E. Monza, M. F. Lucas, D. Almendral, F. J. Ruiz-Dueñas, H. Lund, A. T. Martinez, and V. Guallar. Rational enzyme engineering through biophysical and biochemical modeling. *ACS Catalysis*, 6(3):1624–1629, 2016.
- [2] A. Bryden, G. Phillips Jr., and M. Gleicher. Automated illustration of molecular flexibility. *IEEE Transactions on Visualization and Computer Graphics*, 18(1):132–145, Jan. 2012.
- [3] S. Caboche. Leview: automatic and interactive generation of 2d diagrams for biomacromolecule/ligand interactions. *Journal of cheminformatics*, 5:40, 2013.
- [4] M. Chavent, B. Lévy, M. Krone, K. Bidmon, J.-P. Nominé, T. Ertl, and M. Baaden. GPU-powered tools boost molecular visualization. *Briefings in Bioinformatics*, 12(6):689–701, 2011.
- [5] G. Cipriano and M. Gleicher. Molecular surface abstraction. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1608–1615, Nov. 2007.
- [6] G. Cipriano, G. N. Phillips Jr., and M. Gleicher. Multi-scale surface descriptors. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1201–1208, Nov. 2009.
- [7] R. O. Dror, R. M. Dirks, J. Grossman, H. Xu, and D. E. Shaw. Biomolecular simulation: a computational microscope for molecular biology. *Annual review of biophysics*, 41:429–452, 2012.
- [8] K. Edman, A. Hosseini, M. K. Bjursell, A. Aagaard, L. Wissler, A. Gunnarsson, T. Kaminski, C. Köhler, S. Bäckström, T. J. Jensen, et al. Ligand binding mechanism in steroid receptors: From conserved plasticity to differential evolutionary constraints. *Structure*, 23(12):2280–2290, 2015.
- [9] M. Falk, M. Klann, M. Reuss, and T. Ertl. Visualization of signal transduction processes in the crowded environment of the cell. In *Proceedings of the 2009 IEEE Pacific Visualization Symposium*, pages 169–176, Washington, DC, USA, 2009. IEEE Computer Society.
- [10] M. Falk, M. Krone, and T. Ertl. Atomistic visualization of mesoscopic whole-cell simulations using ray-casted instancing. *Computer Graphics Forum*, pages 195–206, 2013.
- [11] T. D. Goddard and T. E. Ferrin. Visualization software for molecular assemblies. *Current Opinion in Structural Biology*, 17(5):587–595, 2007. Carbohydrates and glycoconjugates / Biophysical methods.
- [12] S. Grottel, P. Beck, C. Muller, G. Reina, J. Roth, H.-R. Trebin, and T. Ertl. Visualization of electrostatic dipoles in molecular dynamics of metal oxides. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2061–2068, Dec. 2012.
- [13] S. Grottel, M. Krone, C. Muller, G. Reina, and T. Ertl. Megamol – a prototyping framework for particle-based visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 21(2):201–214, Feb 2015.

- [14] S. Grottel, M. Krone, K. Scharnowski, and T. Ertl. Object-space ambient occlusion for molecular dynamics. In *IEEE Pacific Visualization Symposium*, pages 209–216, 2012.
- [15] S. Grottel, G. Reina, C. Dachsbaecher, and T. Ertl. Coherent Culling and Shading for Large Molecular Dynamics Visualization. *Computer Graphics Forum*, 29(3):953–962, 2010.
- [16] D. Gunther, R. A. Boto, J. Contreras-Garcia, J.-P. Piquemal, and J. Tierny. Characterizing molecular interactions in chemical systems. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):2476–2485, 2014.
- [17] M. D. Hanwell, D. E. Curtis, D. C. Lonie, T. Vandermeersch, E. Zurek, and G. R. Hutchison. Avogadro: An advanced semantic chemical editor, visualization, and analysis platform. *J. Cheminformatics*, 4:17, 2012.
- [18] P. Hermosilla, V. Guallar, A. Vinacua, and P. Vázquez. High quality illustrative effects for molecular rendering. *Computers & Graphics*, 54:113–120, 2016. Special Issue on CAD/Graphics 2015.
- [19] P. Hermosilla, V. Guallar, A. Vinacua, and P.-P. Vázquez. Instant Visualization of Secondary Structures of Molecular Models. In K. Böhler, L. Linsen, and N. W. John, editors, *Eurographics Workshop on Visual Computing for Biology and Medicine*. The Eurographics Association, 2015.
- [20] M. E. Hodsdon, J. W. Ponder, and D. P. Cistola. The NMR solution structure of intestinal fatty acid-binding protein complexed with palmitate: application of a novel distance geometry algorithm. *Journal of molecular biology*, 264(3):585–602, 1996.
- [21] W. Humphrey, A. Dalke, and K. Schulten. VMD: visual molecular dynamics. *Journal of molecular graphics*, 14(1):33–38, 1996.
- [22] N. A. Khazanov and H. A. Carlson. Exploring the composition of protein-ligand binding sites on a large scale. *PLoS Computational Biology*, 9(11):e1003321, 2013.
- [23] A. Knoll, I. Wald, P. A. Navrátil, M. E. Papka, and K. P. Gaither. Ray tracing and volume rendering large molecular data on multi-core and many-core architectures. In *Proceedings of the 8th International Workshop on Ultrascale Visualization*, UltraVis '13, pages 5:1–5:8, New York, NY, USA, 2013. ACM.
- [24] B. Kozlikova, M. Krone, N. Lindow, M. Falk, M. Baaden, D. Baum, I. Viola, J. Parulek, and H.-C. Hege. Visualization of Biomolecular Structures: State of the Art. In R. Borgo, F. Ganovelli, and I. Viola, editors, *Eurographics Conference on Visualization (EuroVis) - STARs*, pages 061–081. The Eurographics Association, 2015.
- [25] M. Krone, K. Bidmon, and T. Ertl. GPU-based Visualisation of Protein Secondary Structure. In I. S. Lim and W. Tang, editors, *Theory and Practice of Computer Graphics*. The Eurographics Association, 2008.
- [26] M. Krone, K. Bidmon, and T. Ertl. Interactive visualization of molecular surface dynamics. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1391–1398, Nov. 2009.
- [27] O. D. Lampe, I. Viola, N. Reuter, and H. Hauser. Two-level approach to efficient visualization of protein dynamics. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1616–1623, 2007.
- [28] R. A. Laskowski and M. B. Swindells. Ligplot+: multiple ligand–protein interaction diagrams for drug discovery. *Journal of chemical information and modeling*, 51(10):2778–2786, 2011.
- [29] K. Lawonn, M. Krone, T. Ertl, and B. Preim. Line Integral Convolution for Real-Time Illustration of Molecular Surface Shape and Salient Regions. *Computer Graphics Forum*, 33(3):181–190, 2014.
- [30] M. Le Muzic, J. Parulek, A. K. Stavrum, and I. Viola. Illustrative visualization of molecular reactions using omniscient intelligence and passive agents. *Computer Graphics Forum*, 33(3):141–150, 2014.
- [31] N. Lindow, D. Baum, and H.-C. Hege. Interactive Rendering of Materials and Biological Structures on Atomic and Nanoscopic Scale. *Computer Graphics Forum*, 31(3pt4):1325–1334, 2012.
- [32] A. Madadkar-Sobhani and V. Guallar. PELE web server: atomistic study of biomolecular systems at your fingertips. *Nucleic acids research*, 41(W1):W322–W328, 2013.
- [33] K. Momma and F. Izumi. VESTA: a three-dimensional visualization system for electronic and structural analysis. *Journal of Applied Crystallography*, 41(3):653–658, 2008.
- [34] K. Momma and F. Izumi. VESTA 3 for three-dimensional visualization of crystal, volumetric and morphology data. *Journal of Applied Crystallography*, 44(6):1272–1276, 2011.
- [35] J. Parulek, D. Jansson, T. Ropinski, S. Bruckner, A. Ynnerman, and I. Viola. Continuous Levels-of-Detail and Visual Abstraction for Seamless Molecular Visualization. *Computer Graphics Forum*, 33(6):276–287, 2012.
- [36] K. Reda, A. Febretti, A. Knoll, J. Aurisano, J. Leigh, A. Johnson, M. Papka, and M. Hereld. Visualizing large, heterogeneous data in hybrid-reality environments. *IEEE Computer Graphics & Applications*, 33(4):38–48, July 2013.
- [37] G. Reina and T. Ertl. Hardware-accelerated glyphs for mono- and dipoles in molecular dynamics visualization. In *Proceedings of the Seventh Joint Eurographics / IEEE VGTC Conference on Visualization*, EuroVis'05, pages 177–182, Aire-la-Ville, Switzerland, Switzerland, 2005. Eurographics Association.
- [38] S. Salentin, S. Schreiber, V. J. Haupt, M. F. Adasme, and M. Schroeder. Plip: fully automated protein–ligand interaction profiler. *Nucleic acids research*, 43(W1):W443–W447, 2015.
- [39] A. Sarikaya, D. Albers, J. Mitchell, and M. Gleicher. Visualizing validation of protein surface classifiers. *Computer Graphics Forum*, 33(3):171–180, 2014.
- [40] K. Scharnowski, M. Krone, G. Reina, T. Kulschewski, J. Pleiss, and T. Ertl. Comparative visualization of molecular surfaces using deformable models. *Computer Graphics Forum*, 33(3):191–200, 2014.
- [41] N. Schneider, S. Hindle, G. Lange, R. Klein, J. Albrecht, H. Briem, K. Beyer, H. Claußen, M. Gastreich, C. Lemmen, et al. Substantial improvements in large-scale redocking and screening using the novel hyde scoring function. *Journal of computer-aided molecular design*, 26(6):701–723, 2012.
- [42] Schrödinger, LLC. Schrödinger release 2015-4: Maestro version 10.4. <https://www.schrodinger.com/Maestro/>, 2015.
- [43] C. Sigg, T. Weyrich, M. Botsch, and M. Gross. Gpu-based ray-casting of quadratic surfaces. In *Proceedings of the 3rd Eurographics / IEEE VGTC Conference on Point-Based Graphics*, SPBG'06, pages 59–65, Aire-la-Ville, Switzerland, Switzerland, 2006. Eurographics Association.
- [44] R. K. Singh, A. Ethayathulla, T. Jabeen, S. Sharma, P. Kaur, and T. P. Singh. Aspirin induces its anti-inflammatory effects through its specific binding to phospholipase a2: Crystal structure of the complex formed between phospholipase a2 and aspirin at 1.9 Å resolution. *Journal of drug targeting*, 13(2):113–119, 2005.
- [45] R. Skanberg, P. Vázquez, V. Guallar, and T. Ropinski. Real-time molecular visualization supporting diffuse interreflections and ambient occlusion. *IEEE transactions on visualization and computer graphics*, 22(1):718–727, 01 2016.
- [46] K. Stierand, P. C. Maaß, and M. Rarey. Molecular complexes at a glance: automated generation of two-dimensional complex diagrams. *Bioinformatics*, 22(14):1710–1716, 2006.
- [47] M. Tarini, P. Cignoni, and C. Montani. Ambient occlusion and edge cueing for enhancing real time molecular visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1237–1244, 2006.
- [48] M. Wahle and S. Birmanns. Gpu-accelerated visualization of protein dynamics in ribbon mode. In *IS&T/SPIE Electronic Imaging*, pages 786805–786805. International Society for Optics and Photonics, 2011.
- [49] J. Weber. ProteinShader: illustrative rendering of macromolecules. *BMC Structural Biology*, 9(1), 2009.
- [50] G. Wolber and T. Langer. Ligandscout: 3-d pharmacophores derived from protein-bound ligands and their use as virtual screening filters. *Journal of chemical information and modeling*, 45(1):160–169, 2005.
- [51] J. Word, S. C. Lovell, T. H. LaBean, H. C. Taylor, M. E. Zalis, B. K. Presley, J. S. Richardson, and D. C. Richardson. Visualizing and quantifying molecular goodness-of-fit: small-probe contact dots with explicit hydrogen atoms. *Journal of Molecular Biology*, 285(4):1711–1733, 1999.