EGVE Symposium (2008)
B. Mohler and R. van Liere (Editors)

Supervision of 3D Multimodal Rendering for Protein-protein Virtual Docking

G. Bouyer and P. Bourdot

VENISE transversal action, LIMSI-CNRS, Université Paris-Sud, France

Abstract

Protein-Protein docking is a recent practice in biological research which involves using 3D models of proteins to predict the structure of complexes formed by these proteins. Currently, the most common methods used for docking are fully computational approaches, combined with molecular visualization tools. However, these approaches are time consuming and provide a large number of potential solutions. Our basic hypothesis is that Virtual Reality (VR) interactions can combine the benefits of multimodal rendering, biologist's expertise in the field of docking with automated algorithms, in order to increase efficiency in reaching docking solutions. To this end, we have designed an immersive and multimodal application for molecular docking. Visual, audio and haptic feedbacks are combined to communicate biological information, help manipulating proteins and exploring possible solutions of docking. Multimodal distribution is supervised by a rule-based software module, depending on the interaction context.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [User Interfaces]: [Ergonomics, User-centered design, Graphical user interfaces (GUI), Auditory feedback, Voice I/O, Auditory (non-speech) feedback, Haptic I/O] ; H.5.1 [Multimedia Information Systems]: [Artificial, augmented, and virtual realities]

1. Introduction

Proteins are essential compounds of living organisms. They are composed of amino acids which are arranged in a linear chain and form a 3D structure. Proteins achieve structural or mechanical functions. Some of these functions are only possible if proteins combine with each other or with other molecules to form stable complexes. The protein-protein docking field tries to understand which proteins can form stable complexes, how these complexes are formed, and what their role in the cell is. The main information that biologists have to consider is: 3D protein models, physico-chemical interactions (e.g. hydrophobic, electrostatic and Van der Waals forces) and bonds between hydrogen or sulphur atoms on the interface of the two proteins. Automatic tools have been developed to partially solve docking problems. However, fully computational approaches are not yet efficient enough to provide accurate and definite results. For this reason, it is important to develop other user interfaces for computer-aided docking. Our approach is to merge the biologist into the loop of simulation, and give him the possibility to naturally interact with virtual proteins, thanks to multimodal VR interactions. The remainder of this paper is as follows. Section 2 describes current automated techniques for protein-protein docking and previous works on the integration of VR interactions in this domain. Our approach is presented in Section 3 and multimodal VR renderings are designed. These feedbacks are structured in a complete supervision process, detailed in Section 4. We conclude the paper in Section 5.

2. Previous work

2.1. Automated Protein-Protein Docking

Many algorithms dealing with the search for protein binding sites are based on an exhaustive approach. They explore the surface of each molecule in order to scan all potential geometric solutions. Indeed, docking is based, first of all, on the complementarity of protein surfaces. The main problem with these approaches is combinatorial explosion. Some sampling

procedures, such as genetic algorithms, or Monte Carlo simulations have also been tested for docking problems. Scanning procedures can also be restricted by manually giving some specific parameters. Yet computing time remains long (from a few minutes to several hours) and algorithms provide a large number of potential configurations for two proteins. These geometric solutions are then evaluated by a scoring function. This function evaluates physico-chemical interactions between proteins. Here again, scoring functions deal with a very large search space. As a consequence, this step can not be completely covered in a suitable time. Another drawback is accuracy. While docking yields conclusive results in many cases, the risk for selecting false positives is high. The result of the process provides a large set of possible solutions, many of which are not biologically valid and can not all be experimentally tested. Therefore, a final evaluation stage has to be conducted by biologists. Thanks to molecular visualization tools and desktop-based input devices, they can observe various biological data for each potential scored solution. However, it is quite difficult to simultaneously manipulate two 3D structures with a mouse and a keyboard, and visual rendering can be easily overloaded by the data. We think Virtual Reality (VR) technologies are adapted to the docking task as they allow natural interaction. Moreover, adaptative stereovision is suitable to understand 3D conformations of proteins, and multimodal feedbacks such as audio and haptic can reduce the complexity of the visual information and decrease user's stress on this channel. The following section presents a short state-of-the-art of docking applications in Virtual Environments (VE).

2.2. VR and multimodal interaction for Protein-Protein Docking

Taking into account the limits of the current protein docking approaches, the use of new possibilities offered by VR to support molecular docking has gathered some interest in past years. Early work focused mainly on technical needs for visual rendering (e.g. STALK [LFH*97]). Many projects have also explored the potential of multimodal rendering to support molecular docking. A first point is that such projects often use haptic or auditory feedback, rarely both. Projects such as GROPE [BOYJBK90], IVPS [MCET05] and SenSitus [WB03] have studied the combination of haptic and visual cues to manipulate molecules. GROPE and IVPS aim to provide haptic means to perceive force fields or more generally volumetric datasets. SenSitus is a haptic plug-in for VMD (Visual Molecular Dynamics) [HDS96] which provides a haptic feedback about the suitability of the current docking location. Haptic feedback was found on the whole to reduce error (e.g. energetic score of the final configuration) or trajectory lengths, but at the price of increasing completion time (due to necessary local and sequential exploration).

Audio-oriented projects aim to provide the user with clues about molecular properties, in particular binding properties, using earcons (abstract sounds, e.g. a pure tone), auditory icons (everyday sounds) or sonification (mapping of sound parameters in relation to scientific data) [GRGP06]. Early work involved sonification of sequential data, e.g. DNA or amino acids sequences. In contrast, docking involves processing data based on interaction between several hundreds of atoms, and applications which use sonification to study biomolecules are rare.

In conclusion, various work has been carried out regarding immersive and multimodal docking, but the results obtained are still too specific or not convincing enough. Besides, they place the user more in the role of an observer or controller of an automatic process than in the role of an actor. Our approach is more human-centered and addresses the design of a system to carry out protein-protein docking with immersive and multimodal interactions. Our view is that the intelligent management of all the available modalities (i.e. all the means to transmit information through human sensori-motor channels) can improve biologists' manipulations and allow simultaneous rendering of complex physico-chemical data.

3. Our approach for immersive and multimodal protein-protein docking

3.1. Exploring potential solutions

Our main goal is to enable the user to act directly on the docking process. Firstly, the user strongly reduces the search space by selecting a small number of interesting areas on the surface of the proteins whereupon computations are executed. Secondly, the user selects protein complexes that are deemed most probable. This architecture allows two levels of examination: a large scale observation and another level focused on the interface between the two proteins. From this method, we hope (1) to exploit prior knowledge and hypotheses of the biologist about the interaction throughout his manipulation, rather than only while configuring the docking algorithm, and (2) to reduce the risk of "false positives" since the user heads for the best solution step by step rather than chooses the best alternative in a panel of computer-generated complexes.

For this purpose, our system involves three stages (figure 1). During the first stage, the multimodal VR framework, will allow users to manipulate proteins and place them in interesting configurations. Visual stereoscopy, 3D audio and haptic feedback are combined and will lead to a reduction of the search space, using the abilities and specific knowledge of the biologist regarding protein-protein docking. The second stage will be similar to classical automatic processes for protein-protein docking but applied on the restricted space chosen by the user. Finally, the third stage will allow the user to explore the solution space provided by the previous stage, within the VR framework. During the exploration, the

user receives multimodal feedback on the energetic stability of each possible docking configuration. This finally leads to the extraction of relevant configurations that are then tested in biological experimental conditions.

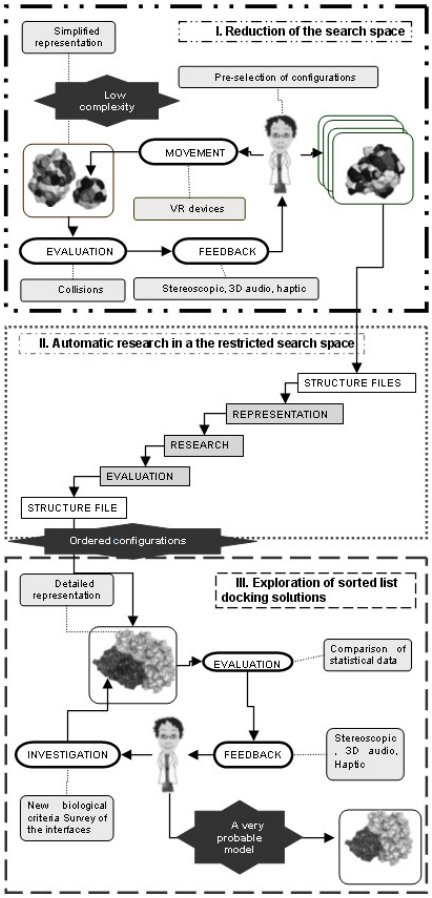

Figure 1: *Our approach of protein-protein docking with VR interactions [FNB*08].*

The following sections only detail the first stage: the reduction of the search space via multimodal interactions with the proteins.

3.2. Informational needs and design principles

We have conducted a preliminary study on the use of existing desktop-based docking tools, as well as the biologists' impressions on how VE equipped with multimodal devices might affect their work [FNB*08].

The first result of this task analysis was the identification of four informational needs, i.e. protein properties or knowledge highly considered by the biologist during the docking task. These four main data are: the electrical charges, the hydrophobic residues, the topological factors and the "hot spots", i.e. residues which are likely to belong to the interface.

Our analysis also put the emphasis on four constraints related to the docking task: the risk of sensory overload (especially on the visual channel), the temporal properties of data (static vs. dynamic), the possible semantic aspect of interactions (e.g. electrostatic complementary) and the need for two levels of exploration (molecular vs. atomic). These constraints have led us to elaborate the following design principles for the four informational needs we have identified:

1. Display molecular surfaces at least visually to support easy manipulation of molecular models and to respect the higher spatial resolution of vision compared to other channels;
2. Display all data involved in the computation of binding energies simultaneously, using a combination of modalities on different channels;
3. When possible, stay congruent with familiar sensory experiences and task-related data semantics, e.g. render physical collisions and electrostatic forces with haptic modalities;
4. Auditory signals can render time-dependent information such as scores.

Some of these ergonomic principles (e.g. 1, 3 and 4) are totally consistent with existing psychophysical results, which state that the most appropriate channel with respect to a given task is the channel that dominates the perception in the context of this task [SS01].

3.3. Multimodal interactions

A complete multi-sensory VR platform is available for our application. Active stereoscopy for visual rendering is provided on a two-face CAVE-like system. 3D audio rendering is generated, spatialised, and displayed using 8 loudspeakers or headphones. Our initial experimental scenario provides bimanual direct manipulation of two proteins (*barnase* and *barstar*). One is manipulated with a tracked 3D mouse and the other is attached to a 6 DOF haptic device as shown in figure 2. These devices, as well as gesture and voice recognition systems (to command the application), are managed by a dedicated software platform.

3.3.1. Visual rendering

Our immersive docking application is centered around Pymol [DeL02], an existing molecular visualization software which is extensively used by biologists. Using an existing platform saved us from developing necessary visual modalities such as atomic, surface or structural representations, onto which physico-chemical information can be projected (atoms type, charges, hydrophobic properties, etc.). Additional visual feedbacks can also be drawn via OpenGL objects, such as electrostatic fields or surface complementarity.

Pymol also offers a Python scripting API which makes

it compatible with an immersive C++ architecture. It allows to program new natural input functionalities useful for biologists. For example, within the immersive VR application, the user can manually select specific atoms on the surface of the protein with a pointing device. This command allows to make the most of biologist's prior knowledge to define "hotspots" and interesting docking areas.

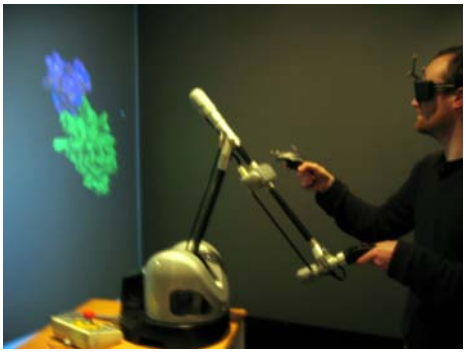

Figure 2: *Bimanual direct manipulation of proteins using a 6 DOF haptic device, a tracked 3D mouse and active visual stereoscopy.*

3.3.2. Haptic interactions

Haptic feedback is known to improve the quality of local interactions in an immersive environment, and to enhance the perception of virtual objects. In our docking application, we intend to use haptic feedback to communicate complementarities (topology, collisions) between proteins and force fields (hydrophobic, electrostatic and Van der Waals) induced by manipulations. In the current prototype, haptic is used to render collisions according to a "lock-and-key" or "LEGO" paradigm, and to render the global electrostatic field by attractive, repulsive and torque force feedback.

A major technical constraint of applications using haptic devices is to maintain a refresh rate between 200 Hz to 1 kHz. This is essential to guarantee stiff and stable rendering. This entails that computations within the application (collisions, scores, etc.) should be carried out in real time, and time latencies should be reduced as much as possible. For that reason, and because our first goal was to validate our multimodal interactions in VR, our application only considers rigid proteins.

In order to render collisions between proteins, we represent surfaces with triangle meshes. For each protein, the mesh is computed from a Protein Data Bank (PDB) model using MSMS [SOS96]. Then the RAPID library is used to detect collisions on the triangle mesh during docking interactions. The result of the collision detection is a list of triangles, needed to compute a repulsion force feedback, which is then sent to the haptic device. The norm of the repulsion force is computed from the number of triangles involved in the collision. The direction of this force comes from the current position of the protein and its last recorded position.

Concerning the electrostatic force feedback, the electrostatic field around the receptor (the larger protein) is first computed off-line with APBS (Adaptive Poisson-Boltzmann Solver) [BSJ*01]. The result is a 3D grid of electrostatic potential. Then, the second protein is immersed in this 3D electrostatic field, leading to calculate a force for each of its charged amino-acids. This set of forces results in a global force and torque for the protein. This method allows us to compute feedback for the resulting global electrostatic force, in linear time, depending on the number of amino-acids in the second protein (figure 3.1).

We are also studying the perceptualization of topologic complementarity. It consists in verifying, before the collision, that the surfaces of the two proteins are geometrically compatible, i.e. that bumps meet holes. This test results in distances calculation between pairs of atoms, and the goal is to have identical distances. Otherwise, a haptic torque is computed to guide the user to a better orientation, and to shorten or extend the appropriate distances (figure 3.2).

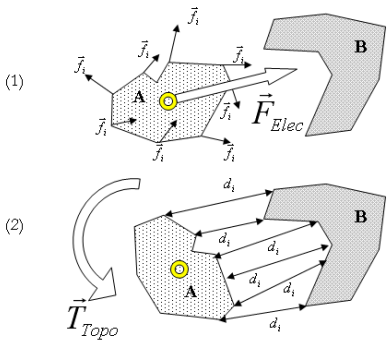

Figure 3: *Principle of haptic interactions designed for electrostatic (1) and topologic (2) complementarities.*

3.3.3. Audio rendering

Two pieces of information seem interesting to be rendered on the audio channel: electrostatic and topologic complementarities. In both cases, these complementarities are the result of numerical data: either a global numerical score based on a biological formula or a set of distances/forces (between some pairs of atoms). These numerical data can be sonified, for example by changing the frequency of a sound. In the case of several variables to sonify, we are currently studying the spatialisation of several audio feedbacks, to match the spatial distribution of the atoms that are concerned.

Audio feedback is also used to render collisions. Two ways have been chosen: the broadcasting of recorded collision sounds to enhance the realism and the variation of the

frequency of a beep, depending on the surface complementarity score.

The following table (figure 4) sums up haptic and audio interactions designed for these three informational needs.

Information	Electrostatic complementarity	Topologic complementarity	Collisions
Available data	N forces computed on N points of the second protein	N distances between N pairs of atoms	Triangles in contact in a mesh
Goal of the biologist	Minimize electrostatic forces	N equal distances	Feel the collision
Haptic rendering	1 Force + 1 Torque (attraction)	1 Torque (attraction) No translation force	1 Force + 1 Torque (penalty)
Audio rendering	Communicate N variations of forces with N spatialized sonifications	Communicate N variations of distances with N spatialized sonifications	Sounds of collision
	Communicate the variation of 1 global score with 1 sonification		1 sonification : frequency / score

Figure 4: *Haptic and audio interactions designed for three biological complementarities.*

4. Multimodal Supervision

Multimodal rendering choices presented up to then in this paper are static, i.e. decided before the execution of a docking application. However, concrete achievement of the renderings must be controlled all along the use of the application. These feedbacks depend on the unpredictable context of the interactions, induced by: the user, the state of the virtual scene, the real environment and the system. In order to manage this dynamic control of multimodal rendering, we have modelled and developed a supervision process [BBA07]. This multimodal supervision also ensures that the communications between the user and the system are consistent and respond to the same logical, ergonomic and psychophysical principles.

4.1. Architecture

Our supervision process is built on four components (figure 5):

Real World The biologist commands the system and interacts with the proteins using a 3D mouse and a haptic device. He/she perceives multimodal feedbacks through screens, loudspeakers and the haptic device. His/her hand and head are also tracked via infrared sensors.

Docking Application It is represented in our model by a triplet. The VE contains a virtual representation of the user, and the whole data which can be observed or manipulated: proteins, charges, hotspots, scores, numerical variable, etc. This VE is handled by a hardware and software architecture (e.g. Pymol). Input and output interfaces are run by drivers, and feedbacks are created by visual, audio and haptic engines. Lastly, an interaction manager interprets the commands of the user, asks the supervisor the appropriate rendering modalities and starts the rendering engines.

Supervisor It is in charge of deciding the most appropriate multimodal rendering, thanks to a knowledge base of logical distribution rules and a context base which contains the current state of all the elements that could influence the rendering.

Observer/Interpreter This module manages the communications between the application and the supervisor, by translating exchanged information. Its second role is the continuous observation of the application, in order to isolate context elements, and to dynamically provide them to the supervisor. These elements can be: tracking data (e.g. position of the user), virtual scene data (e.g. relative position of the proteins), state of the rendering capacities, etc.

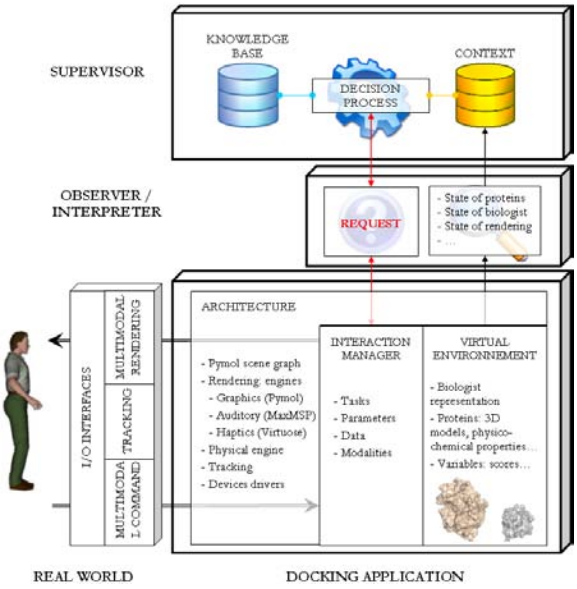

Figure 5: *Architecture of multimodal rendering supervision for our docking application.*

In order to communicate, these four components need a common language. This is why we have set up a model of multimodal VR interactions.

4.2. Model of multimodal VR interactions

Each possible interaction in the docking application is represented by three semantical and generic elements: the type of the task that is performed, the parameters specifying the task and the data concerned by the task (figure 6). This model allows us to represent numerous concrete interactions of our docking application. For example, when a biologist analyses

the global topology of a protein, this interaction is modelled by:

```
task = manipulation
parameter = global
data = topological object = protein
```

When the user wants to begin a new interaction, this one is modelled by the interaction manager using the triplet (`task`, `parameter`, `data`). The supervisor can then exploit its knowledge on these generic three elements to decide the most effective multimodal rendering.

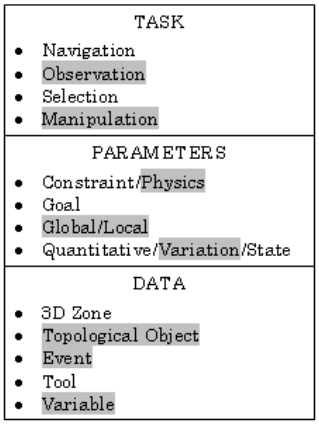

Figure 6: *Our model of multimodal VR interactions: type of task, parameters and data. Grey elements are currently used in our docking application.*

4.3. Decision process

The process between the request to the supervisor and the result of the decision is detailed in figure 7. This process is programmed in Prolog, in order to benefit from logical languages proof capacities.

First, the knowledge base and the elements of the context (media, modalities, user, etc.) are all represented by Prolog predicates. Then, the requested interaction is described by an ID and three predicates corresponding to the triplet (`task`, `parameter`, `data`).

The principle of the decision process is the elaboration of a score for each available modality in the application. This calculation relies on logical rules which apply our design principles of section 3.3 on each element which could influence the rendering: type of data, type of parameter (e.g. `global`, `local`), type of task (e.g. `observation`, `manipulation`), user's preferences, default rendering chosen by the designer of the application, etc. Specific treatments can also be done depending on elements of the context base: current interactions, media loads, user's position, etc. Then the supervisor returns the best modalities for the current interaction and context to the interaction manager.

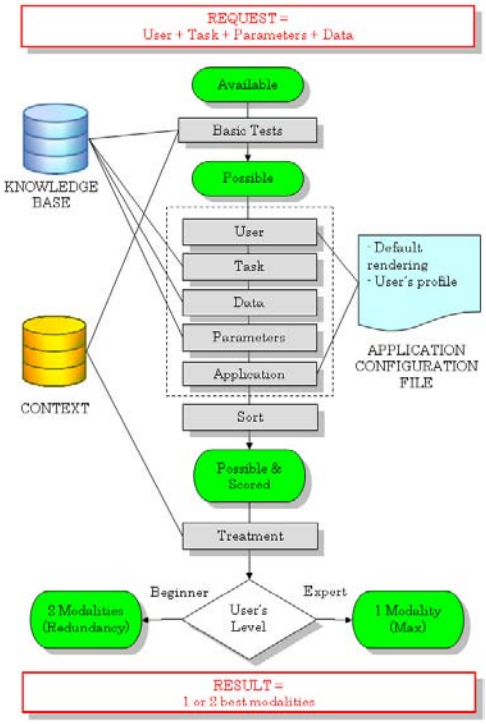

Figure 7: *Sequence of filters and scores leading to the choice of the rendering modalities.*

4.4. Interactive scenario

In order to evaluate our multimodal VR interactions and the decisions of the supervisor during docking tasks, we have elaborated a complete interactive scenario, in collaboration with biologists and ergonomists. This scenario comes in the form of successive or simultaneous basics tasks to accomplish under clear conditions. Figure 8 gives a selection of these tasks. For each of them, the figure shows their Prolog representation as well as the results of the supervision with the current knowledge base and under "ideal" conditions (non disturbing real environment, expert user, no technical latencies, all renderings presented in section 3.3 available). This part of scenario corresponds to a typical beginning of a manual docking session. First the biologist observes each isolated protein in a `global` way (only one protein appears in the figure). Then he/she wants to analyse some physicochemical properties of the protein (interactions `2` and `3`), which can be disable or requested later (`5`). Secondly, he/she begins to manipulate proteins `X` and `Y`, and tries to find the best configuration while feeling complementarity scores (`4`, `6`). Finally, complementarity renderings are stopped and he/she can try to fit the two proteins.

In front of these tasks, we can see that the multimodal supervisor decides the following renderings:

- Biological properties which concern one protein (interac-

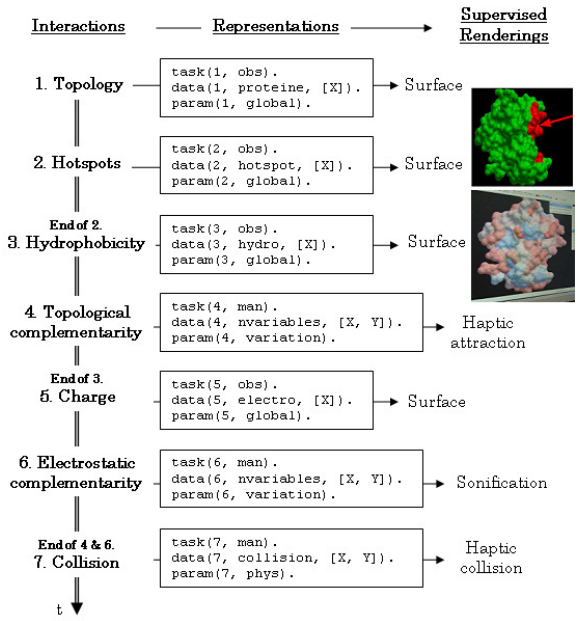

Figure 8: *Example of interactive scenario of docking with supervised renderings.*

tions 1, 2, 3 and 5) are rendered along the visual channel. This result respects constraints identified in section 3.3. The `global` parameter leads to a surface representation, whereas `local` would have led to an atomic visualization. Moreover there is a specific rule to respect the consistency of visual representations (atoms, surface, structure) for protein properties that are projected on the protein model.

- Topological complementarity is haptized because the knowledge base contains a rule which favours this natural channel for `manipulation` tasks, and when one wants to feel the `variation` of a data.
- Electrostatic complementarity is sonified. Indeed, we have chosen for the moment to limit the number of simultaneous haptic modalities to one, and the haptic channel is already occupied by interaction 4. Otherwise, the decision would have been the same as topological complementarity, as these interactions are modelled by the same triplet of generic elements.
- Collision is haptised. This channel is favoured by the `manipulation` task, and also by the `phys` parameter (which means it has to respect some physical model). This decision is possible since interactions 4 and 6 have been stopped.

This example of interactive scenario shows (1) the respect for multimodal design principles and (2) the adaptability of the decision depending on the dynamic context. This context is mainly composed of current and previous interactions, which dynamically change media availability and loads.

4.5. Multimodal situations in protein-protein docking

In addition to contextual elements that appear in the previous example, other multimodal situations in protein-protein docking can happen, during which a rule-based supervision is useful. These problems can be handled by our model and our architecture. Yet there remains some prior evaluations to precise the knowledge base.

First, whereas one of the goals of our multimodal system is to offload the visual channel by distributing information on the other ones, we have to be sure that these new renderings are compatible with each other and do not confuse the informational contents. Existing studies have identified cross-modal effects under very precise conditions. For example, Shimojo remind us that vision can alter other channels and that sound can alter some aspects of vision [SS01]. The pseudo-haptic effect [LCK*00] also shows that haptic perception can be modified by relevant visual modalities. In the docking context, we have to study the psychophysical or ergonomic validity of some concrete combinations of feedbacks. Once such combinations will be evaluated, we will easily integrate the results in the knowledge base, for instance to control the number of simultaneous visual, audio and haptic data.

Secondly, an interesting spatial problem is the relative positions of objects with each other, and of these objects with user's hand and viewpoint. For example, manipulating two proteins can lead to configurations where some data are masked, requiring specific visual modalities (such as visual transparency of one protein or even change of the global point of view). These situations can be tracked by the context observer and requested to the supervisor, leading to the appropriate update of the renderings via "spatial rules" (figure 5).

A third problem we are currently experimenting is latency. Indeed, visual, audio and haptic technologies require various computation and transmission time performances. When two devices are used to render the same information, this can cause informational bias or even physiological disturbance. One way to address this issue is to synchronize the two renderings from the initial stage. The idea is to add "temporal rules" within the supervisor's knowledge base, in order to manage the date and length of each rendering. On the other hand this needs to know in advance the technical latency of each rendering.

5. Conclusions and Future Work

In this work we have tackled the use of multimodal rendering to enhance the quality of protein-protein docking tasks. We have proposed some visual, audio and haptic interactions

to convey meaningful biological information to users. We have structured these immersive feedbacks in a global supervision process, centered on an intelligent module which controls the appropriateness of the designed renderings with the dynamic context.

We are currently enhancing our multimodal VR docking application by settling the "spatial" and "temporal" rules for specific multimodal situations we have discussed in previous section. Simultaneously, we are developing the context observer to handle these new elements. The next step is to conduct psychophysical and usability experiments in order to validate our interactions, and to expand the knowledge base, e.g. on the combination of audio and haptic feedback for two different biological complementarities.

6. Acknowledgments

This work was conducted as a part of the CoRSAIRe project (Combination of sensori-motor renderings for the immersive analysis of results). This project is supported by the French National Research Agency (ANR). We want to thank the academic partners of this project: LIMSI-CNRS (UPR 3251), EBGM-INSERM (UMRS 726) and the Ergonomics Unit (ECI) of Paris 5 René Descartes University. We especially thank Jean-Marie Burkhardt and Julien Nelson from ECI for conducting very useful ergonomic analysis, and Nicolas Férey, Christine Martin and Abdelhamid Drif from LIMSI-CNRS for their contribution to the development of the docking application.

References

[BBA07] BOUYER G., BOURDOT P., AMMI M.: Supervision of task-oriented multimodal rendering for vr applications. In *Eurographics Symposium on Virtual Environments (EGVE)* (2007). 15-18th July 2007, Weimar, Germany.

[BOYJBK90] BROOKS F., OUH-YOUNG M., J.J. BATTER J. J., KILPATRICK P. J.: Project grope: Haptic displays for scientific visualization. In *SIGGRAPH Comput. Graph.* (1990), vol. 24(4), pp. 177–185.

[BSJ*01] BAKER N. A., SEPT D., JOSEPH S., HOLST M. J., MCCAMMON J. A.: Electrostatics of nanosystems: application to microtubules and the ribosome. *National. Academy of. Science.of the USA 98* (2001), 10037–10041.

[DeL02] DELANO W.: The pymol molecular graphics system. http://www.pymol.org, 2002.

[FNB*08] FÉREY N., NELSON J., BOUYER G., MARTIN C., BOURDOT P., BURKHARDT J.-M.: User needs analysis to design a 3d multimodal protein-docking interface. In *IEEE Symposium on 3D User Interfaces* (2008). 8-9th March 2008, Reno, Nevada, USA.

[GRGP06] GARCIA-RUIZ M., GUTTIEREZ-PULIDO J.: An overview of auditory display to assist comprehension of molecular information. *Interacting with Computers 18(4)* (2006), 853–868.

[HDS96] HUMPHREY W., DALKE A., SCHULTEN K.: Vmd - visual molecular dynamics. *Journal of Molecular Graphics and Modeling 14* (1996), 33–38.

[LCK*00] LÉCUYER A., COQUILLART S., KHEDDAR A., RICHARD P., COIFFET P.: Pseudo-haptic feedback : Can isometric input devices simulate force feedback? IEEE Int. Conf. on Virtual Reality, pages 83-90, New Brunswick, US, 2000.

[LFH*97] LEVINE D., FACELLO M., HALLSTROM P., REEDER G., WALENZ B., STEVENS F.: Stalk: an interactive system for virtual molecular docking. In *Proceedings of IEEE Computational Science and Engineering* (1997), vol. 4(2), pp. 55–65.

[MCET05] MACIEJEWSKI R., CHOI S., EBERT D. S., TAN H. Z.: Multi-modal perceptualization of volumetric data and its application to molecular docking. In *WHC '05: Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems* (Washington, DC, USA, 2005), IEEE Computer Society, pp. 511–514.

[SOS96] SANNER M., OLSON A., SPEHNER J.: Reduced surface: An efficient way to compute molecular surfaces. *Biopolymers 38* (1996), 305–320.

[SS01] SHIMOJO S., SHAMS L.: Sensory modalities are not separate modalities: plasticity and interactions. *Current Opinion in Neurobiology 11* (2001), 505–509.

[WB03] WRIGGERS W., BIRMANNS S.: Interactive fitting augmented by force-feedback and virtual reality. *Journal of Cellular Biology 144* (2003), 123–131.