Part IV
Collaboration with Domain Experts

Visualization requires three components to work: data, tasks, and an audience. Any foundation on data visualization will deal with all three. Furthermore, as in other computer science disciplines, there is basic and applied visualization research. Basic visualization research is driven by generalized visualization tasks, evaluation issues, or theoretical questions about visualization. Applied visualization starts from an application (most likely outside visualization) including data, task, and audience, and tries to find the best visualization solution for the given case. Of course, applied visualization research may (and should) lead to new visualization tasks which triggers new basic visualization research. Also, basic visualization research offers new possibilities for applications.

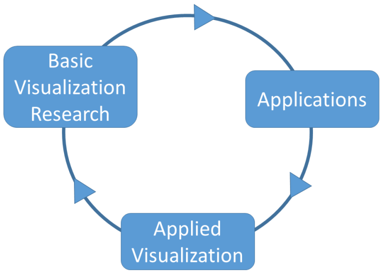

In addition, current and past experience shows that there is a substantial difference between visualizations for domain experts knowing data, underlying model, and assumptions as well as tasks very well, and visualizations for a broad audience with very different backgrounds and much less or even no knowledge about data, models, and tasks. The latter case is covered in another part, so this part of the book focuses only on domain experts. As domains, the collaboration models, data, and tasks vary substantially, this part starts with a chapter that contains seven successful case studies. Visualization experts describe in these cases how they approached the cooperation, what was important, and which lessons they took out of these projects. The domains cover biological, medical, and engineering examples which are the most often seen domains in cooperations between visualization experts and domain experts. While these cooperations mainly concern the cooperation within academia, the second chapter describes experiences and advice from industry practioneers on the collaboration between university and commercial companies. The final chapter of this part takes a somewhat more abstract point of view. It looks at the process of actually selecting a domain expert as collaboration partner, and how to create impact by the research. The authors indicate clearly that just starting a collaboration by chance might work, but more mindful strategies provide better chances to lead to success. This includes some thoughts on measures of success where it becomes clear that this is, at least partly, a subjective question that any visualization researcher (and domain expert) has to define for him- or herself.

References

Gerik Scheuermann

Chapter 14
Case Studies for Working with Domain Experts

JOHANNA BEYER, *Harvard University, USA*
CHARLES HANSEN, *University of Utah, USA*
MARIO HLAWITSCHKA, *University of Applied Sciences Leipzig, Germany*
INGRID HOTZ, *Linköping University, Sweden*
BARBORA KOZLIKOVA, *Masaryk University, Czech Republic*
GERIK SCHEUERMANN[1], *Leipzig University, Germany*
MARKUS STOMMEL, *TU Dortmund, Germany*
MARC STREIT, *Johannes Kepler University Linz, Austria*
JOHANNES WASCHKE, *University of Applied Sciences Leipzig, Germany*
THOMAS WISCHGOLL, *Wright State University, USA*
YONG WAN, *University of Utah, USA*

The collaboration with domain experts concentrates always on an application domain where the experts work. Usually, they provide the data, and directions of research that require visualization support. This chapter presents seven successful cases of such collaborations. The domain varies from biology and medicine to mechanical engineering. There are examples of long time cooperation as well as smaller short term projects. The description concentrates on the process, output, and especially on the lessons learnt from these cooperations. The scientific work is described to understand the context and goals of the cooperation but many details can only be found in the references. The reason for this unusual writing is the wish on one hand to describe various aspects of collaboration with domain experts which is an important part of the foundations of data visualization. On the other hand, the text should not become lengthy and filled with too many details of individual cases that can be found elsewhere.

1 Case Study: FluoRender

Yong Wan and Charles Hansen

FluoRender is a software package for visualizing and analyzing 3D and 4D (3D over time) fluorescence microscopy data. FluoRender has become an established system with many features driven by collaborations with biologists delivering visualization, segmentation, measurement, and tracking functions with an emphasis on accuracy, interactivity, and intuitiveness. Originally developed for the Zebrafish community, it has extended to other biological applications. FluoRender has been deployed as a standard tool in research labs both domestically and internationally, facilitating research on cell movements, neuronal circuitry, and tissue development during conventional analysis of wild-type and mutant embryos of popular model

---

[1] e-mail: `scheuermann@informatik.uni-leipzig.de`

species. The increasing popularity and growing user base of FluoRender have given rise to new visualization and analysis challenges from both general and specialized workflows. Close collaborations between biologists and computer scientists have provided a systematic insight into the workflows in real-world biological research. A data analysis workflow is indeed a far cry from a rigid pipeline; it has to be highly adaptable and easily customizable for varying data analysis needs. In practice, user interactions and decisions are involved through the entire data analysis workflow. Interactive visualization and analysis functions work hand in hand, allowing exploring and iterative investigations, as well as progressive improvement to the refined results that lead to biological discoveries.

One example of the expanded user base is our collaboration with Professor Gabrielle Kardon and her interest in developing a mouse atlas [27]. Such atlases are important for understanding normal anatomy and the development and function of structures, and for determining the etiology of congenital abnormalities. Although the focus of FluoRender was the analysis and visualization of confocal microscopy data, the atlas required not only volume rendering and segmentation but also polygonal modeling, for muscles and bones, as well as advanced texturing which captured the anisotropy of muscle tissue reflecting what was seen in the confocal scan Figure 37.

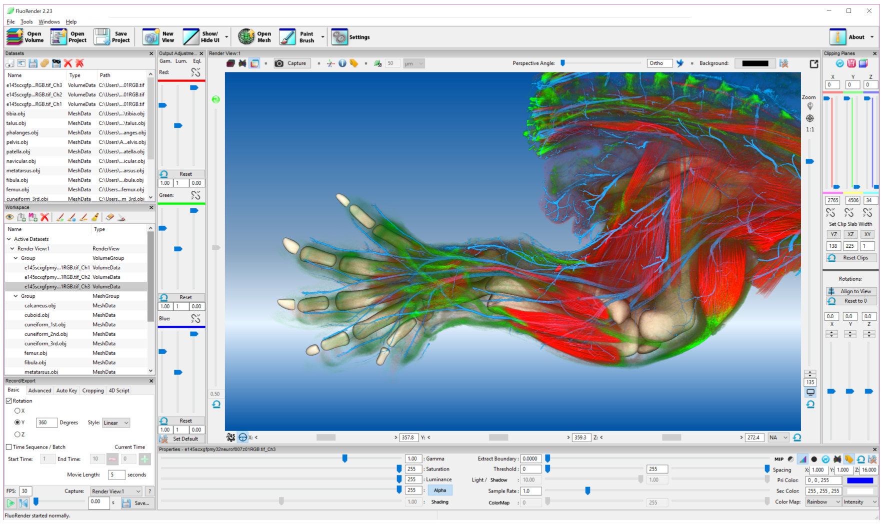

Fig. 37 Mouse Hind Limb Atlas, Y. Wan, C. Hansen, SCI Institute and A. Kelsey Lewis and G. Kardon, Human Genetics, University of Utah.

1.1 Lessons Learned

In the development of FluoRender, we learned several lessons from the integral collaboration with biologists. The first lesson was communication. Domain scientists use terminology from their domain. Visualization scientists use terminology from the visualization and computer science domain. It is critical that communications find the common ground so that there are no misunderstandings. This takes time for detailed explanation from the visualization researchers of what their ideas and methods accomplish and how they are accomplished. It also is incumbent on the domain scientists to explain their ideas and methods in terms that are easily understood. While the visualization scientists may not have the same detailed knowledge that the domain scientists have of a particular biological or domain process, the visualization scientists should have sufficient application knowledge to understand the biology and processes being investigated. Once this is accomplished, collaborative research is enhanced, and advances are more easily made.

Another lesson we have learned is that all participating collaborators have science to accomplish. Of course, the domain scientists have science research questions they seek to answer in their particular domain. They are seeking answers to questions and testing hypotheses in their particular biological domain. At the same time, the visualization scientists should be advancing the field of visualization. This delicate balance of both advancing the field and providing solutions to needs and problems in biology is critical to a successful collaboration. The visualization researchers should not be simply serving the needs of the domain collaborators but should focus on both advancing visualization and providing solutions to requirements from the domain scientists.

Do not simply ask domain scientists what needs solving or which desired features are missing in a visualization system. By understanding the domain workflow, productive progress can be made. This often requires working with collaborators in their lab and having collaborators spend time in the visualization lab. It is important to not simply meet and discuss the domain problem but to actually work with collaborators in their research setting and observe what data anlysis tasks are easily accomplished while others can be improved using an updated workflow. This leads to better understanding of the practical domain problems in greater detail.

Lastly, it is important to be creative. By providing creative solution to biological problems, advances in both biology and visualization science can be made. Such creative solutions should be enjoyable and fun for all sides of the collaboration. Success can be measured by typical means such as publications in domain journals or in visualization journals. It can also be measured by recognition of those in the respective fields. For example, our collaboration with the Kardon Lab was recognized by the Director of NIH in 2017 [9]. Dr. Francis Collins recognized the advances in muscle and soft tissue development research by Professor Kardon in finding that when two muscles are fused and indistinct, forelimb function is limited. This is due to a mutation in a gene called Tbx3. It is already known that the mutation of Tbx3 in human is the cause of a rare condition called ulnar-mammary syndrome (UMS). However, because of the lack of detailed examination and visual-

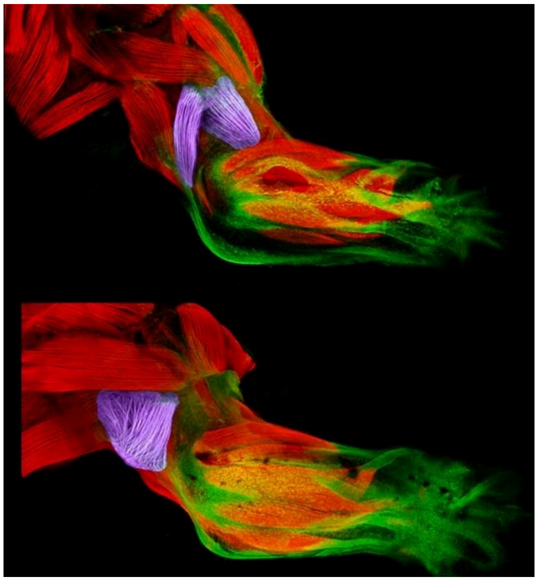

Fig. 38 Developing forelimb of a healthy mouse strain (top) compared to that of a mutant mouse strain with a stiff, abnormal gait (bottom). Lateral triceps in brachialis muscles (purple), other types of muscle (red), and tendons (green). Note that in the top image (wildtype mouse), the lateral triceps and brachialis muscles are distinct, while in the bottom image (mutant mouse), the two muscles are fused, limiting the forelimb's function.

ization on human patients, muscle anomalies of the UMS patients were overlooked in the original research. Researchers initially declined the idea of fused muscles in UMS patients because of the anatomical differences between mice and humans. Interestingly, at Kardon's urging, a similar pattern of missing muscles was confirmed in the reexamination of a UMS patient. This research demonstrates the astounding similarity between human and mouse genetics, which provides an excellent application stage for visualization tools, as such details in Figure 38 can be prohibitive to obtain for human patients. Further improvements to the clarity are achieved by coloring muscles using the interactive segmentation tools in FluoRender, which are like 3D paint brushes. They have to be intuitive and enjoyable to use for researchers, as the operations can be repeated for from several tens to hundreds of samples in an investigation. The Director of NIH also recognized FluoRender [9]:

> '... there's one more NIH connection to this work. Kardon's team produced this image, featured in the University of Utah's 2016 Research as Art competition, using a free software program, called FluoRender that was developed by another NIH-supported team at the University of Utah. FluoRender enables researchers to take a series of 2D photos from a scanning confocal microscope and turn them into amazingly informative 3D imagery.'

2 Case Study: Connectomics

Johanna Beyer

This chapter describes an on-going collaboration between the Visual Computing Group at Harvard University and neuroscientists working in the field of Connectomics at the Harvard Center for Brain Science. The collaboration focuses on the visualization and analysis of large-scale connectomics data and has spanned over the last eight years.

2.1 Domain Problem

Connectomics aims to reconstruct the detailed neural connectivity in the mammalian brain, containing billions of interconnected nerve cells, at the resolution of individual connections (i.e., synapses). Determining this "wiring diagram" or so-called *connectome* is one of the grand challenges of modern neuroscience and will allow scientists to better understand how the brain functions and develops, and how mental illnesses and neural pathologies manifest themselves on the connectivity level. Recent advances in high-resolution electron microscopy and sample preparation have made it possible to acquire data at the speed and resolution necessary to reconstruct the brain's connectivity at the level of individual synapses. However, the acquired image stacks are typically hundreds of terabytes to petabytes in size, exhibit severe noise and imaging artifacts, and can contain tens of thousands of complex neural structures. A lot of effort has gone into developing novel methods for data acquisition, volume registration, and (semi-)automatic segmentation, resulting in large labeled volumes of brain tissue. The main goal of our collaboration was to enable the next logical step: supporting scalable and interactive volume exploration and visual analysis of the collected data.

2.2 Process and Output

This case study encompasses several sub-projects that were all developed within the same collaboration over the last eight years. Projects always started with initial meetings and interviews with the neuroscientists. Most neuroscientists we talked to, although in the same group, often had very different visualization and analysis tasks and requirements. Therefore, the initial project phase always focused on finding a main collaborator for the next project and quickly coming up with an initial prototype for further discussion. Throughout the development and implementation of the project we kept a tight feedback loop with the domain scientists, to ensure that our project a) solved an actual problem; and b) solved the problem that was relevant for our collaborators. To achieve this, we held regular in-person meetings, video

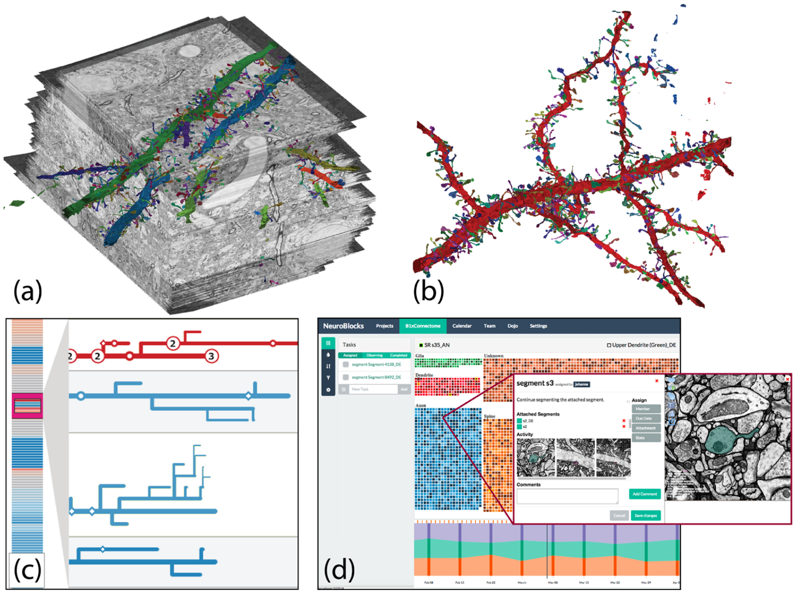

Fig. 39 ***Exploration, visualization, and analysis of connectomics data.*** (a) Volume rendering of a segmented terabyte electron microscopy volume. (b) A segmented dendrite. (c) Connectivity analysis with *Neurolines*. (d) Visual tracking of a segmentation project with a pop-up for visual proof-reading.

conferences, and visited their labs to observe our collaborators at their routine data collection and processing tasks. After finishing each software prototype, we made sure to demonstrate it to our collaborators and encouraged them to use the software on their own. We evaluated the usefulness of our projects based on expert feedback and specific use-cases that were developed together with our collaborators.

In the initial phase of our collaboration, we focused on developing basic visualization and data management infrastructure for large-scale segmented (i.e., labeled) neuroscience datasets. Having built that initial framework, in later years we shifted our focus on visual analysis and integrating domain-knowledge into our data exploration framework. Figure 39 shows some of the different projects we have developed over time.

Our first major project was a scalable volume rendering framework [12] for exploring petascale microscopy data streams. In a second step, we extended the framework to support interactive volume visualization of labeled volumes [4]. Using this framework, scientists could interactively explore their raw image data, as well as their segmentation data. However, a more in-depth quantitative evaluation was still difficult. To allow our collaborators to explore the data based on their domain knowledge, we developed ConnectomeExplorer [3], a tool for interactive domain-specific queries of neuroscience data. These queries allowed the first glimpse into how different segmented neural structures were connected. As connectomics is ultimately interested in the neural connectivity of the brain, our next project Neurolines [1] solely focused on visualizing and exploring the connectivity of axons and dendrites. In this project, we went beyond the initial volumetric data visualization and focused on an abstract 2D view that enabled users to focus on connectivity rather than the detailed 3D morphology of their data. Therefore, we abstracted the topology of 3D brain tissue data into a multi-scale, relative distance-preserving subway map visualization where each neurite is represented as a tree structure based on its real, but adaptively simplified, anatomy, and its branches. During the development of those projects, it became clear to us that the major bottleneck of our collaborators was not the analysis of their data, but the actual segmentation process, and tracking the segmentation status of a volume over time. Therefore, we developed tools for visual proofreading of segmentations [13], as well as for visual segmentation tracking and management [2].

What this list of different projects demonstrates, is that the challenges and goals of domain experts often evolve and change over time. Therefore, it is vital to meet with them regularly, observe how they work, and to make an effort to understand their current set of challenges. That includes not just the challenges that are stated explicitly but also the implicit challenges that scientists might not even think of mentioning.

2.3 Lessons Learned

Collaborating with domain scientists has its own set of challenges, however, it is also incredibly rewarding. Here are some lessons we learned during our collaboration with neuroscientists:

Understand the domain problem. While this hint seems obvious, make sure to meet with domain scientists regularly. Visit their lab, follow them around and observe their work. What do they spend the most time on, what are the difficult and/or annoying tasks? Domain scientists typically do not have a background in visualization, so they might not know where and how visualization could be most useful. For example, our collaborators would have never thought about a visual tool for tracking the segmentation process over time, even though they considered this one of the most time consuming and difficult tasks in their everyday work.

Define your roles and expectations. Make sure that everyone on the project is on the same page regarding each others roles and responsibilities. Is the end goal of the project a scientific publication, a useful software framework, or, ideally, both? Make sure to address these concerns early on.

Have a user adoption strategy. We have observed a much better adoption of web-based systems as compared to stand-alone applications that require local installation of software. Neuroscientists typically do not have a desktop PC but prefer laptops, as they tend to move around a lot between wet labs and other lab spaces, which makes adoption of software that requires specific hardware (e.g., GPUs, large monitors) a lot more difficult.

Go where the domain challenge takes you, not where your previous research has positioned you. Sometimes you will discover interesting problems and challenges that you did not expect or foresee, and that will require you to branch out into a different area or sub-area of research. Yet, if those are the most pressing challenges of your collaborators, embrace them and do not shy away from them. Start with the nail (i.e., the domain problem), then find the hammer, not the other way around.

3 Case Study: Mechanical Part Design

Gerik Scheuermann, Markus Stommel, Ingrid Hotz

lead: GS, Markus Stommel, Ingrid

Component design is a major task in mechanical engineering. There is a well-defined work flow including structural mechanics simulation and analysis using the finite element method. We looked for possibilities to use tensor visualization of stress fields to leverage the full tensor information for the design. This is in strong contrast to the usual reduction to scalar fields like the von Mises stress which is done in all engineering post processing tools. We tested nearly all available tensor visualization techniques until we finally found a way to show directional information using tensor lines that actually led to better component design. We describe the problem, the cooperation process, the success of the method and the learned lessons.

3.1 Domain Problem

In this case study, we look at a standard problem in mechanical engineering. Engineers have to design a mechanical part with defined functional and qualitative properties which can be produced by standard methods. This is an important part of the product development process in mechanical engineering and follows a clearly defined work flow. First, a manual sketch is created, followed by a 2D, and finally a 3D CAD model. This model is the basis for a virtual mechanical test using the Finite Element Method (FEM). The FEM result is interpreted using visualization. If the result is sound, a rapid prototype will undergo physical tests before a classical prototypes is finally tested. If the virtual test shows problems or has to be further optimized, the design is altered and undergoes again a FEM analysis.

As example, we use a break lever of a bike which is currently made from metal and shall be replaced by a plastic component. The example is still fictitious in the sense that the actual manufacturer is not involved but realistic enough to show the potential in engineering terms. The specific engineering goal is to optimize the plastic rib support structures where the baseline is a textbook design using straight ribs. Rib structures are the most often used reinforcement structure for injection-molded plastic parts. Their design means definition of position, number and shape of the ribs while considering given boundary conditions that follow from the manufacturing process or the part appearance. Even though there are algorithmic optimization methods for some design steps, rib design is still a manual process driven by the engineer's experience. Therefore, this is a typical example of trying to support a domain expert's daily task by enforcing his/her intuition through visualization.

3.2 Process and Output

The modern product design process in mechanical engineering is a completely virtual process that leads to a physical prototype by 3D printing. It consists of a number of improvements cycles. Each cycle contains 3D CAD design, FEM, and analysis using visualization. The design criteria include part stiffness, maximum stress peaks, weight, geometrical or functional boundary conditions and also practical aspects of manufacturing. For material stressing, this comparison is performed so far on the basis of a couple of scalar key metrics.

The idea of this case study (for more details see [19]) is to use the complete FEM result instead, especially all stress tensor information to obtain an optimal design. The visualization partners offered the domain experts (i.e. the engineers) a framework of almost all tensor visualization methods ever invented in visualization research. Especially, we tested multiple linked views and linking-and-brushing for stress tensors, as well as several different tensor line methods. The engineers tested these methods and discussed their meaning in mechanics with visualization researchers. The question was always what does this visualization mean for the mechanics and how can the visual information be used to improve the design. After a number of visualization methods that did not deliver insight for the design process, we finally ended up with tensor lines and fabric textures [14]. Here, we showed planar cuts through the stress field. Thicker and thinner lines showed the eigen directions. This led to an intuitive design of rib support structures, see Figure 41.

This intuitive design meant to follow the tensor lines of the stress field for rib support design. In a first step, we compared the standard textbook design with three different rib designs following different tensor lines. In a second step, we verified the results experimentally by 3D printed brake levers. The results confirmed the hypothesis that tensor lines are good guidelines for rib support structures and lead to stiffer designs without additional material or production costs. As can be seen in figure 41, all three test designs performed substantially better than the textbook design. Honestly, we did not try any other designs, so basically every informed rib design was better than the standard.

For the engineers, a major outcome has been a hypothesis that can substantially support the design process of technical parts. The results of the finite element simulations and the experiments give evidence that tensor lines are valuable for the design of rib patterns.

Fig. 40 Tensor fabric of the stress tensor in a planar cut through the CAD model. The engineer draw some lines manually to design rib support structures aligned with the tensor lines. These manual lines were the basis for the first new design. Two other drawings based on the same visualization led to the other two designs in our case study.

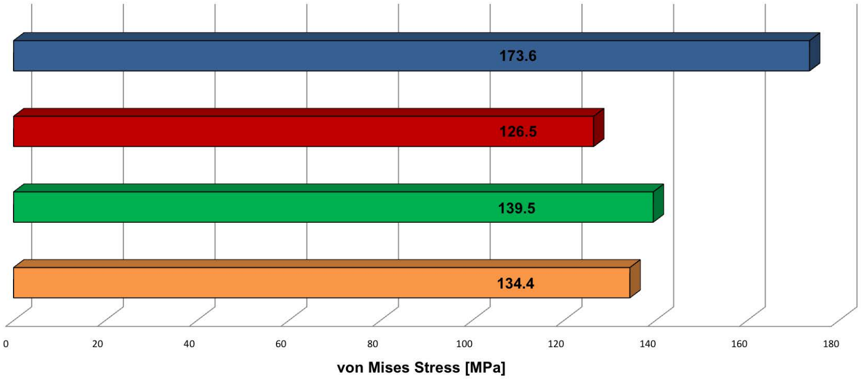

Fig. 41 Maximal von Mises stress in the four different rib structure designs. The top blue bar is the text book result. It shows substantially higher maximal stress, i.e. the design is much worse than the three new designs.

3.3 Lessons Learned

In this case study, it can be clearly seen that a close cooperation between domain experts and visualization scientists is needed for success. The engineers do not know about many modern visualization methods, so they have no access to them or no understanding of them. The visualization scientists do not have enough understanding of the creative tasks of the domain expert and his/her thinking that leads from the problem over the visual input to better solutions like a better design in this case. Also, it is difficult and sometimes nearly impossible to find the best visual metaphors for insight on the domain expert side without classical try and error. Testing the different methods, discussing their meaning in terms of the problem at hand is the key. In this example, the engineers need to derive insight about the stress transport in the part, and even more important, an idea how to place the rib support structures from the visualization. Therefore, the lessons are:

- Try to present more data than before.
- Look for information that helps the domain expert with his/her task.
- Try out many possibilities of visual representations and data.
- Close cooperation is key - present your possibilities and let the domain expert explains their thoughts.

4 Case Study: Drug Target Prioritization

Marc Streit

This section describes the process, outcome, and lessons learned from a research collaboration between the visualization group at Johannes Kepler University Linz and a computational biology group at the pharmaceutical company Boehringer Ingelheim. The goal of the collaboration was to develop visual analysis solutions that help researchers to identify new drug targets for cancer therapy. A drug target constitutes the basis for the development of next generation drugs.

4.1 Domain Problem

Discovering new drug targets is a challenging process because domain experts need to take into account a rich spectrum of data sources. The data sources that need to be incorporated in the exploratory analysis include experimental data from patients, animals, and cell lines, but also publicly available knowledge of what we know about biological processes and diseases.

4.2 Process and Output

Together with partners from Harvard University we started to work with a public cancer genomics dataset from *The Cancer Genome Atlas (TCGA)* project. TCGA was a large US-based initiative that followed the goal of collecting and analyzing biomolecular data from cancer patients for all major tumor types. Based on continuous feedback from cancer genomics researchers, we developed StratomeX [20], a visual analysis tool for comparing patient subsets in large-scale heterogeneous genomics data. We initially published the work as a design study paper at EuroVis'12.

Later on, we extended the tool with guided exploration techniques that support users in picking potentially interesting data subsets during the exploration [26]. In contrast to our earlier work on StratomeX [20] that appeared at a visualization conference, the guided exploration technique was published in Nature Methods.

A core component of the guided exploration workflow was a visualization technique for ranking genes and other entities based on statistical scores and meta-attributes. As ranking problems appear in many different contexts, we generalized the solution and developed the LineUp visualization technique [1]. LineUp allows users to flexible create and explore multi-attribute rankings. The technique was published at IEEE InfoVis'13 where it won the Best Paper Award. LineUp was later on also integrated as a component in the Microsoft PowerBI software. LineUp is available as an open source JavaScript library (`https://lineup.js.org`) that can be flexibly used as a component in various environments, such as Jupyter Notebooks and R Notebooks. Making the library publicly available not only increases the reproducibility of the visualization research but also increases the potential for adoption of the technique.

Publishing our visual analysis tool for genomics data in Nature Methods helped us to gain interest from pharmaceutical companies, which finally led to a three year research collaboration with Boehringer Ingelheim. As part of this collaboration, we created the Ordino drug target discovery tool [25] that at its heart also integrates the interactive ranking technique LineUp. To increase the impact in the fast progressing life science community, we uploaded the paper to `biorxiv.org` at the time of the initial paper submission and made the source code available on Github.

Making a research prototype ready for productive use goes far beyond what a research collaboration is able to cover. To be able to deploy, maintain, and extend the platform, we founded a spin-off company that goes the extra mile required in terms of software development. Only by being able to demonstrate that we can transform innovative visualization solutions developed as part of a research collaboration, into a stable and feature-rich software, we were able to acquire additional funding for the next phase of the collaboration.

As another positive side-effect of having a collaborator that actively uses our tools, we had access to a growing provenance graph containing automatically recorded visualizations and user interactions from the visual exploration sessions. To make the provenance information accessible to the users, we developed the KnowledgePearls search and retrieval solution for querying and exploring similar analysis states, which we again published in the visualization community [24].

4.3 Lessons Learned

One of the most important lessons we have learned is that it is essential to have collaborators who acknowledge visualization as a scientific field that goes beyond creating pretty pictures. Try to find out at the very beginning if your collaborators are convinced that visualization can contribute to solving their domain problems. If

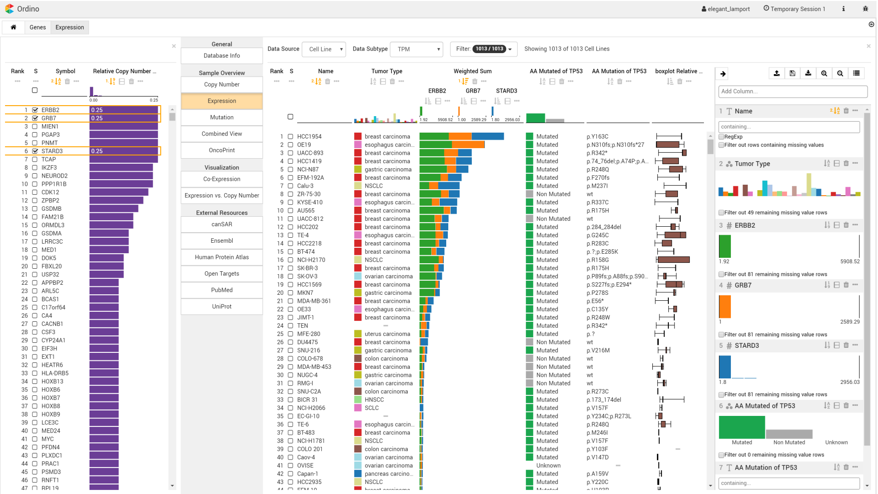

Fig. 42 ***The Ordino visual cancer analysis tool for ranking and exploring genes and other entities based on statistical scores and additional meta-attributes.***

this is not the case, convincing them in the course of the project is extremely difficult and frustrating.

Make it clear at the beginning of the collaboration what can be expected as output and—even more importantly—what is out of scope. The typical output of a research collaboration are visualization prototypes and publications targeted at the visualization community. Promising more than that will likely fall on your head later on.

The most critical and valuable resource domain experts can contribute is their own time. The more added value they see in the visual analysis solutions, the more time they will contribute. The more time they contribute, the more valuable the outcome will be for their own work as well.

The longer the collaboration lasts, the more productive it becomes. Understanding the domain problem is key to being able to contribute. However, learning about a target domain and understanding the domain-specific language is time-consuming and can take months or even years. One success strategy is to stick with one or few problem domains, if possible.

Having a liaison person on board is highly beneficial. In our project, the actual users of our tools are biologists and other life science experts, while our direct collaborators are bioinformaticans. The value and role of a liaison person is discussed in Chapter 15.

5 Case Study: In-Situ Simulation Visualization of Parameter Spaces

Thomas Wischgoll

This chapter describes an ongoing collaboration between the Advanced Visual Data Analysis group at Wright State University and researchers at Wright Patterson Air Force Base. This research collaboration resulted in various different research projects, including the simulation and visualization of a dragonfly during take-off [16]. The following sections will describe a visualization solution that addresses the need for being able to visualize parameter spaces for models from the cognitive science realm.

The ability to do rapid visual assessments of parameter spaces has the potential to change the work-flow for both model simulation and model fitting/parameter recovery. It enables the rapid identification of input parameters that result in similar output data or model behaviors. This allows researchers to eliminate redundant input parameters for more efficient use of modeling and simulation computational resources. For example, should two parameters exhibit a strong correlation, one might be held constant while the other varied in order to capture all the unique model behaviors. Further, early visual assessment of the parameter space means that ineffective or incorrect models may be rapidly identified and eliminated from study. This again results in effective use of both experimenter and computational time. Finally, parameter space visualizations can reveal unexpected relationships between the parameters and model behavior. If the behavior is incorrect, errors in model design or in model may be more easily found. If the behavior is novel, parameter space visualization will have resulted in new hypotheses or expanded research findings.

This approach [10] is a web-based solution that is capable of handling larger data sets compared to other commonly available solutions. At the same time, the described solution is directly integrated into the server structure that is used to run the simulations for the models of the cognitive science researchers. As such, it is readily available within the interface for starting and controlling. Hence, the researchers can immediately run the visualization on the simulation data that was calculated so far and make any adjustments to the simulation as necessary.

5.1 Domain Problem

Web-based visualizations are of interest in this application area as they can be directly integrated into the high-performance computing (HPC) environment. At the same time, this approach eliminates the need to install additional software on the researchers' computers beyond a browser as security limitations may not allow installation of any software of any kind. The potential for interacting with the data and feeding any resulting visually-identified parameter constraints directly into the modeling and simulation process would further improve the modeling work-flow.

The data sets typically are larger than many common JavaScript-based tools, such as D3 [5] or Plotly can handle. Such common tool-kits are not capable of handling data sets that contain more than half a million data points. At the same time, downloading data sets of that size takes a considerable time as well. It is therefore more desirable to generate the visualization on the server directly where the data is computed and stored. Then, only the visual results need to be transferred which is typically a lot less data compared to the entire data set.

5.2 Process and Output

This project evolved out of a close relationship between researchers at Wright State University and the 7/11 Human Performance Wing at Wright Patterson Air Force base. We were fortunate enough to have representatives of our collaborators be present at regular research meetings to discuss specific approaches for visualizing high-dimensional parameter space data. Due to the restrictive environment disallowing our collaborators to install software on their own, it was quickly identified that a web-based solution for our collaborator's visualization needs was a web-based approach. However, conventional tools or services, such as D3 or Plotly, quickly failed to handle the size of the data sets. With more than half a million data points, the browser typically ran out of memory so that a visualization could not be achieved with those tools. In addition, having to download each data set would take too long to be acceptable. Instead, a server side visualization approach was chosen. This avoids the need for downloading the data and at the same time allows the researchers to visualize their data while it is still being generated for in-situ visualization. This server-side approach still utilizes the D3 library. However, it uses node.js to render the results into an image, which is then transferred to the client. Any interactive features, such as axis, are still drawn on the client side to preserve the full interactivity of the visualization approach. The server renders the visualization results in parallel on as many nodes as are available or allocated combined with with additional performance enhancements resulting in faster rendering times. In addition, the parallel approach allows us to handle significantly larger data sets compared to the original implementation. Overall, this approach enables our collaborators to visualize their data sets quickly. It is integrated within their web-based scheduling mechanism and hence readily available to them as they track their simulation progress. The in-situ capabilities allow them to adjust parameters on the fly based on our visualization.

5.3 Lessons Learned

There are several lessons learned from this project that can be useful in general. The fact that our collaborators actively participated in regular research meetings was a great benefit to the outcomes of the project. It helped make the generated tools better suited for their needs. Due to the fact that the visualization tools were directly integrated into their work-flow made those tools directly accessible to the researchers using that high-performance computing platform for their simulations making it as easy as clicking a button on their web-based scheduling interface. Unfortunately, the project ended as the lead team of researchers at Wright Patterson Air Force Base moved to to a different national lab. However, the visualization tools are still accessible for the computing platform.

The fact that this project uses a web-based platform also has the additional advantage of keeping track of utilization within that server environment. Hence, this provides yet another way of evaluating the visualization tool based on our collaborators voting with their mouse by electing whether to use out tool or not.

Overall, this project resulted in a successful implementation of a visualization approach that enabled our collaborators to directly visualize their results. It was very well received among the users of that high-performance computing platform by allowing them to immediately investigate their simulations while they were being computed.

6 Case Study: Protein analysis and Visualization (CAVER)

Barbora Kozlikova

Understanding the structure and behavior of protein molecules is crucial in many biological and biochemical fields, such as drug design and protein engineering. This process requires studying the proteins from many aspects, including their constitution, physico-chemical properties, temporal behavior, or interactions with other molecules. These properties and their combination is very hard to perceive and understand using the traditionally used visual representations of molecules and animations of their behavior over time. Therefore, the biochemists require specifically designed visualizations which help them to explore and understand the proteins in more convenient and faster way. This creates very tight connection between the biochemical and visualization fields.

The case described in this chapter captures the interesting aspects of our long-term collaboration with protein engineers from Loschmidt Laboratories at the Masaryk University in Brno, focusing namely on the exploration of the void space inside proteins and its connection with the protein surface. Such paths, connecting the inner voids with the surface, are denoted as tunnels in literature. There are already several existing algorithms and tools available for tunnel calculation. One of the first tools for tunnel detection was the CAVER tool, whose first version was developed in 2007 [22]. In the same year, the authors of this tool contacted us with the request to improve their original algorithm and to enable them to get insight into the detected tunnels. At that point our collaboration was established, which lasts until now. Of course, over the years, the research tasks of the biochemists have changed. At the very beginning, they focused on detection of tunnels in static molecules, which was further extended to molecular dynamics simulations [8]. The possibility to simulate longer and longer trajectories of protein movements resulted in the situation when the biochemists are not capable of observing such simulations frame-by-frame. Therefore, new visual abstractions, enabling the domain experts to drive their focus only on the interesting parts of the simulation, became a necessity [6, 7]. Moreover, recent advances in computational capabilities enabled the biochemists to generate large ensembles of molecular dynamics simulations. This leads to new challenges for visual guidance and comparative visualization, which is our current topic of common interest.

6.1 Domain Problem

As already stated, the main focus of our collaborators from the protein engineering group is the detection and analysis of tunnels in proteins. The presence of these void paths significantly influences the reactivity of proteins with small ligands entering the protein inner space and performing a chemical reaction in the protein active site. This specific site is capable of reacting with the ligand and the product of such a reaction can be, for example, a basis of a new drug. On the other hand, the goal of protein engineers is to change the properties of the protein by mutating selected amino acids, i.e., by replacing one amino acid by another. The protein engineers proved that the mutations of amino acids in the close vicinity of tunnels have large impact on protein properties [17], such as its stability in normal temperature or activity towards ligands and other molecules [21].

With the increasing possibilities to capture large molecular dynamics simulations, currently spanning to hundreds of thousands of timesteps, the domain experts urgently needed help with the exploration of behavior of tunnels in them. They were interested namely in the development of the shape and properties of the tunnel narrowest site, denoted as the tunnel bottleneck. However, their interest was driven towards the overall behavior of the whole tunnel as well. Here they were interested namely in the changes of tunnel shape and its constitution, i.e., the movements of amino acids forming the tunnel boundary. For better understanding of the conformational changes of these amino acids, they also had to understand their physico-chemical properties, such as their hydrophobicity or charges of atoms.

6.2 Process and Output

The starting phases of our collaboration were mostly about finding the common language with the protein engineers and understand their needs. In the first stage, we were focusing on the improvement of the original grid-based CAVER algorithm for the detection of tunnels and we came with the approach utilizing Voronoi diagrams. The next step was to visualize the resulting tunnels and their surrounding amino acids so the biochemists could get proper insight. The first straightforward solution was to create a plugin for CAVER to the commonly used PyMOL tool for molecular visualization [23]. This enabled us to get our algorithm to the domain experts worldwide. However, as PyMOL was not designed specifically for visualization of tunnels, it could show the resulting tunnel only in a very basic way which was not sufficient for proper exploration. Therefore, with the protein engineers from the Loschmidt Laboratories we decided to design and create a new tool for visualization and visual exploration of protein tunnels. This standalone tool, called CAVER Analyst, intensified our collaboration even more, as we had to closely discuss not only the functions of the tool, but also the user interface, layout, and interaction. This enabled us to get more insight into the daily workflow of the protein engineers and on the other hand, the protein engineers also had to look at their research problems from a different viewpoint. The development of the first published version of CAVER Analyst took several years. There were several reasons for that. First, we did not have enough experience with designing such a robust tool which led to several bad design choices which took us significant time to fix. Second, the fluctuation of students at the university made the development complicated, as approximately every three years the development team had changed. And getting new students to the interdisciplinary topic and knowing the tool always took several months. Before releasing the Beta version, intensive testing by the biochemists had to be performed. Therefore, CAVER Analyst 1.0 was released in 2014, after almost 8 years of its development [18]. The release helped not only with getting the tool designing specifically for tunnel exploration to the community, but created a stable platform for further prototyping of new visualizations. Since 2014, we were working namely on designing specific visualization methods for visual exploration of large molecular dynamics simulations, which are not anymore observable by traditional animation. We designed very abstracted representations of changes of the shape and surrounding amino acids of the tunnel bottleneck [6] (see Figure 43 a) and a method for exploration of tunnel changes along its centerline [7] (see Figure 43 b). Both representations helped the biochemists to design proper mutations of amino acids surrounding a given tunnel. The application of these mutations had a significant desired impact on the function and properties of the corresponding proteins.

In 2018, we released the 2.0 version of CAVER Analyst [15], which contains these techniques for visual exploration of single trajectories of molecular dynamics simulations. Currently we continue in the successful collaboration in the same manner as it proved to be worthwhile for both partner sides. We collaborate on new techniques for visual exploration of ensembles of trajectories and their comparison and again, we are using CAVER Analyst as the prototyping environment.

6.3 Lessons Learned

This project taught us several important things about successful collaboration which are worth to share.

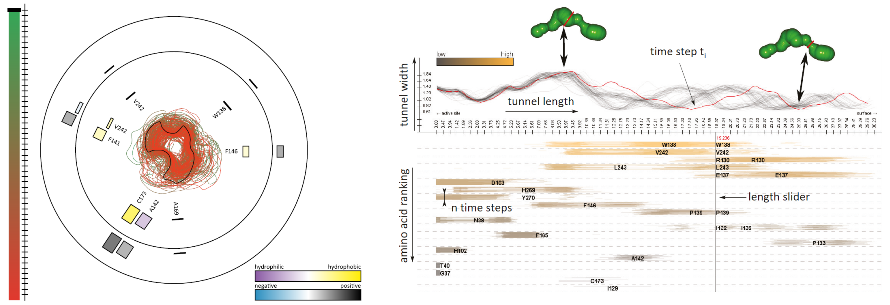

a) b)

Fig. 43 a) Visualization of the shape of tunnel bottleneck and its changes over time. b) Visualization of tunnel width along the centerline over time and the surrounding amino acids, with their amount of contribution to the tunnel.

First, building a trust between two research groups, having the research interest in completely different fields, is a long-term run. Except for speaking the same language, we had to clarify our research goals and expectations from each other. The protein engineers had to understand that our basic research is based on designing new visualization methods and publishing them on visualization venues and we had to keep in mind the real usability of the designed methods and their benefit in the biochemical research.

Second, designing a tool for prototyping and for public release is completely different. However, creating the tool with keeping in mind the actual users makes also the future prototyping much easier for the developers. In our experience, the prototyping tool is more sustainable if it is paid attention to its usability as well.

The biochemists can also participate on the visualization publications which makes the collaboration even stronger. We decided to include the protein engineers to our publications by helping us with designing, performing, and describing the case studies, demonstrating the usefulness of our newly developed methods.

To conclude, this still ongoing project is already for several years resulting in interesting visualization methods which gave our collaborators the necessary insight to protein tunnels and their behavior and properties. Moreover, the CAVER tools are well accepted by the community, which was one of the initial goals of our collaboration.

7 Small-Scale Visualization Projects

Johannes Waschke and Mario Hlawitschka

Visualization as a field of research aims to support the evaluation and presentation of data delivered from domain experts. Some domains work with highly complex data, which in their native form might be too challenging for the human understanding. This is for example the case with multi-dimensional data like diffusion MRI data, which heavily rely on visualization and therefore attract much research interest. In the shadow of these projects, however, many smaller challenges are waiting for the visualization community. These challenges span time ranges of several working weeks or months, up to a single working year. In contrast to larger projects, which heavily depend on data processing and visualization as a means for further understanding, the purpose of short-term projects might remain in the production of beautiful and informative figures. The focus lies on a better presentation of domain experts' result. Basically, scientific visualization in that case extends the work of a graphic designer, but with difficulty increased for two reasons: First, visualization in scientific context requires for scientific validity. Second, complexity of the utilized data types is usually higher and therefore demands for enhanced knowledge in data processing.

7.1 Domain Problem

This book chapter explains the *general* concepts of small-scale visualization projects rather than focusing on a single experience. However, I worked on a number of problems arising around trajectory data, which should serve as an example here. The goal for trajectory data is to present a number (up to tens of thousands) of motion paths or connections, which often overlap, twist, or occlude each other. The presentation of such data should emphasize certain data characteristics. Interesting characteristics could be similarities between subsets of the data, as well as additional properties like speed or direction of trajectories. Evaluation of this data cannot be solved with standard image processing software and thus I see two options how to proceed.

One option for the experts is to do it by themselves. The methods of evaluation and presentation naturally depend on the interests and abilities of the respective researcher and they are biased by both official and "unspoken" rules of the concrete research domain. An example is to track the motion of a surgeon's instrument with the purpose to classify the performance of the surgeon. Given the trajectory data, physicians tend to prefer an evaluation performed by a number of experts which should describe or rate qualities of the motion. Examples for such evaluations are verbal phrases like "very direct" or "many attempts" rather than using standardized tests. Other research fields prefer quantified results, which includes parameters like speed or distance. Hence, the way of evaluation depends on individual knowledge and the domain's standards. All of these results might be good enough to answer the domain's research question, but they lack visual power for a presentation, and they only contain a limited degree of information.

The second option is a collaboration between domain experts and visualization researchers. Of course, scientific visualization also cannot provide all possible information in one single image, but it can extend the perspective and try to maximize the information level. Furthermore, it should be clear that visualization must consider the above-mentioned qualitative and quantitative evaluation steps and therefore should be seen as extension of traditional evaluation steps. As often stated, the open dialogue between both sides, and the interest to understand each other's needs, is crucial.

Besides the knowledge of visualization techniques, as well as experience in possibilities and limitations of visual data presentation, technical factors play a major role. Domains with little relation to computer science are accustomed to work with standard software like Excel and PowerPoint. However, the abilities of these frameworks are limited and their visualization results often are some kind of generic graph. As soon as we leave common data types, which can quickly be the case for individual experimental setups, data processing is a challenge for domain experts. And how can you visualize data that is even hard for you to simply open, read, and store? For many data types, specialized software frameworks exist. Application of them can be, unfortunately, a complicated endeavor, since they might be hard to find, hard to install (or have to be self-compiled), and hard to use.

7.2 Process and Output

The working steps, as I have experienced, are relatively straight-forward and similar compared to descriptions from the previous chapters. It is unquestionable to have a solid relationship between visualizers and domain experts that builds on mutual interest to understand and help each other. On the one hand, domain experts must be open for new ideas and they must sacrifice time to formalize the problem and give feedback. On the other hand, the visualizers should carefully avoid ludicrous visual experiments that distract from the actual work of the domain expert. These small-scale projects are meant to benefit the presentation of another domain's research and not to push visualization research to new limits.

In the beginning the problem should be stated and the needs of the domain experts should be clarified. It is apparently helpful to consider their ideas and previous work, but I also recommend to keep some distance — to avoid a biased perspective on the problem solution. As an example for trajectory data, I want to bring up a question that concerns the level of abstraction of the visualization. For various domains it is interesting to (only) have an abstract view on the data, which for example summarizes trajectories to clusters, or which simply presents features derived from the original data. However, in the medical field, absolute positions often have high meaning and abstractions are less popular. Physicians prefer to see unaltered positions in anatomical context, and points like these have to be learned during the meetings.

In my experience, further development was performed in numerous iteration steps. After a couple of days or weeks, I presented a prototype and proposed some options how to continue. These prototypes gave the domain experts a quick preview on the realistic outcome. Demonstrations of the prototypes regularly gave the experts new ideas that could be considered for further development of the visualization. This is repeated as long as the resources allow it. Finally, a common result is a set of figures that is planned to be used in the domain expert's publication.

7.3 Lessons Learned

First of all there are technical questions concerning the implementation of the visualization algorithm. While we generally aim for the goal to fabricate re-usable visualization tools (which includes a user-friendly interface, documented API, tutorials, and so on), the reality struggles with time pressure and short working periods. In my opinion, a lot of time can be wasted on making the software too "consumer-friendly". Since this is against common software development rules, I want to provide further explanations on that. Small projects often involve very specific data types and the visualization is highly individual as well. The chances of the visualization algorithm to be ever used again — in that concrete implementation — might be very low. Thus, it simply saves time to handle the technical steps by yourself and provide only the result images to the domain experts. Additional features, like a graphical user interface or a beginner's guide to the software, could cost some months of work but are probably never used — and thus a waste of time.

Some visualization problems are too challenging for standard software (and common knowledge), but not interesting enough to be an active field of visualization research. Nevertheless, solving these problems and proposing creative visualizations helps the domain experts to compose better papers with at least improved conveyance of their research results. However, there are negative aspects for the scientific visualizer. Since the quantity of the newly created knowledge is usually small — we are considering projects of several weeks or months —, publication as a full paper might be inappropriate. The chances to be successful within the visualization community are higher for long-term projects (with a higher degree of new results) or at least for multiple accumulative projects with similar challenges. The scientific visualizer is here, up to a certain degree, a service provider for the domain experts.

This raises two questions. First, who is responsible for the funding, since low-grade publications don't help to raise money? One answer are institutes that employ researchers of multiple disciplines. This provides the additional benefit of closely situated working places and therefore good conditions for interdisciplinary collaborations. Outside of these institutes, small-scale visualization challenges might only be worth to keep as a side-project.

The second question is about publications and the scientific reward for the work. Since it seems harder to publish in a high-impact visualization medium, the point of co-authorship in a domain-specific journal should be raised early in the meetings with the domain experts. The visualization researcher could be co-author in a high-level journal of the experts' domain, and — if the methods are not covered in the paper yet — on a small conference or at a poster session. Another promising function of small-scale projects is to use them for teaching and for thesis projects of students. They form a practical problem and thus motivate the students to grow into the research field of visualization — small-scale projects can be a nice starter for a research career.

References

Johannes Waschke and Mario Hlawitschka

1. Al-Awami, A., Beyer, J., Strobelt, H., Kasthuri, N., Lichtman, J., Pfister, H., Hadwiger, M.: NeuroLines: A subway map metaphor for visualizing nanoscale neuronal connectivity. IEEE Trans. on Visualization and Computer Graphics (Proc. IEEE InfoVis ’14) **20**(12), 2369–2378 (2014)
2. Al-Awami, A.K., Beyer, J., Haehn, D., Kasthuri, N., Lichtman, J.W., Pfister, H., Hadwiger, M.: Neuroblocks - visual tracking of segmentation and proofreading for large connectomics projects. IEEE Transactions on Visualization and Computer Graphics **22**(1), 738–746 (2016). DOI 10.1109/TVCG.2015.2467441
3. Beyer, J., Al-Awami, A., Kasthuri, N., Lichtman, J.W., Pfister, H., Hadwiger, M.: ConnectomeExplorer: Query-Guided Visual Analysis of Large Volumetric Neuroscience Data. IEEE Trans. on Visualization and Computer Graphics (Proc. IEEE SciVis ’13) **19**(12), 2868–2877 (2013)
4. Beyer, J., Hadwiger, M., Al-Awami, A., Jeong, W.K., Kasthuri, N., Lichtman, J., Pfister, H.: Exploring the Connectome - Petascale Volume Visualization of Microscopy Data Streams. IEEE Computer Graphics and Applications **33**(4), 50–61 (2013)
5. Bostock, M., Ogievetsky, V., Heer, J.: D3: Data-driven documents. IEEE Transactions on Visualization and Computer Graphics **17**(12), 2301–2309 (2011). DOI 10.1109/TVCG.2011.185
6. Byska, J., Jurcik, A., Groeller, M.E., Viola, I., Kozlikova, B.: Molecollar and tunnel heat map visualizations for conveying spatio-temporo-chemical properties across and along protein voids. Computer Graphics Forum **34**(3), 1–10. DOI 10.1111/cgf.12612. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12612`
7. Byska, J., Muzic, M.L., Groeller, M.E., Viola, I., Kozlikova, B.: AnimoAminoMiner: Exploration of protein tunnels and their properties in molecular dynamics. IEEE Transactions on Visualization and Computer Graphics **22**(1), 747–756 (2016). DOI 10.1109/TVCG.2015.2467434
8. Chovancova, E., Pavelka, A., Benes, P., Strnad, O., Brezovsky, J., Kozlikova, B., Gora, A., Sustr, V., Klvana, M., Medek, P., Biedermannova, L., Sochor, J., Damborsky, J.: CAVER 3.0: a tool for the analysis of transport pathways in dynamic protein structures. PLoS Comput. Biol. **8**(10), e1002,708 (2012)
9. Collins, D.F.: Snapshots of life: Muscling in on development (2017). URL `https://directorsblog.nih.gov/2017/07/27/snapshots-of-life-muscling-in-on-development/`
10. Glendenning, K., Wischgoll, T., Harris, J., Vickery, R., Blaha, L.: Parameter space visualization for large-scale datasets using parallel coordinate plots. Journal of Imaging Science and Technology **60**(1), 10,406–1–10,406–8 (2016)
11. Gratzl, S., Lex, A., Gehlenborg, N., Pfister, H., Streit, M.: LineUp: Visual Analysis of Multi-Attribute Rankings. IEEE Transactions on Visualization and Computer Graphics (InfoVis ’13) **19**(12), 2277–2286 (2013). DOI 10.1109/TVCG.2013.173
12. Hadwiger, M., Beyer, J., Jeong, W.K., Pfister, H.: Interactive volume exploration of petascale microscopy data streams using a visualization-driven virtual memory approach. IEEE Trans. on Visualization and Computer Graphics (Proc. IEEE SciVis’12) **18**(12), 2285–2294 (2012)
13. Haehn, D., Knowles-Barley, S., Roberts, M., Beyer, J., Kasthuri, N., Lichtman, J., Pfister, H.: Design and evaluation of interactive proofreading tools for connectomics. IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE SciVis 2014) **20**(12), 2466–2475 (2014)

14. Hotz, I., Feng, L., Hagen, H., Hamann, B., Joy, K., Jeremic, B.: Physically based methods for tensor field visualization. In: Proceedings of the conference on Visualization'04, pp. 123–130. IEEE Computer Society (2004)
15. Jurcik, A., Bednar, D., Byska, J., Marques, S.M., Furmanova, K., Daniel, L., Kokkonen, P., Brezovsky, J., Strnad, O., Stourac, J., Pavelka, A., Manak, M., Damborsky, J., Kozlikova, B.: CAVER Analyst 2.0: Analysis and Visualization of Channels and Tunnels in Protein Structures and Molecular Dynamics Trajectories. Bioinformatics (2018)
16. Koehler, C., Wischgoll, T., Dong, H., Gaston, Z.: Vortex visualization in ultra low reynolds number insect flight. IEEE Transactions on Visualization and Computer Graphics **17**(3), 2071–2079 (2011)
17. Koudelakova, T., Chaloupkova, R., Brezovsky, J., Prokop, Z., Sebestova, E., Hesseler, M., Khabiri, M., Plevaka, M., Kulik, D., KutaâĂĚSmatanova, I., Rezacova, P., Ettrich, R., Bornscheuer, U.T., Damborsky, J.: Engineering enzyme stability and resistance to an organic cosolvent by modification of residues in the access tunnel. Angewandte Chemie International Edition **52**(7), 1959–1963. DOI 10.1002/anie.201206708. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/anie.201206708`
18. Kozlikova, B., Sebestova, E., Sustr, V., Brezovsky, J., Strnad, O., Daniel, L., Bednar, D., Pavelka, A., Manak, M., Bezdeka, M., Benes, P., Kotry, M., Gora, A., Damborsky, J., Sochor, J.: CAVER Analyst 1.0: graphic tool for interactive visualization and analysis of tunnels and channels in protein structures. Bioinformatics **30**(18), 2684–2685 (2014)
19. Kratz, A., Schoeneich, M., Zobel, V., Burgeth, B., Scheuermann, G., Hotz, I., Stommel, M.: Tensor visualization driven mechanical component design. In: Visualization Symposium (PacificVis), 2014 IEEE Pacific, pp. 145–152. IEEE (2014)
20. Lex, A., Streit, M., Schulz, H.J., Partl, C., Schmalstieg, D., Park, P.J., Gehlenborg, N.: StratomeX: Visual Analysis of Large-Scale Heterogeneous Genomics Data for Cancer Subtype Characterization. Computer Graphics Forum (EuroVis '12) **31**(3), 1175–1184 (2012). DOI 10.1111/j.1467-8659.2012.03110.x
21. Pavlova, M., Klvana, M., Prokop, Z., Chaloupkova, R., Banas, P., Otyepka, M., Wade, R.C., Tsuda, M., Nagata, Y., Damborsky, J.: Redesigning dehalogenase access tunnels as a strategy for degrading an anthropogenic substrate. Nat. Chem. Biol. **5**(10), 727–733 (2009)
22. Petrek, M., Otyepka, M., Banas, P., Kosinova, P., Koca, J., Damborsky, J.: CAVER: a new tool to explore routes from protein clefts, pockets and cavities. BMC Bioinformatics **7**, 316 (2006)
23. Schrödinger, LLC: The PyMOL molecular graphics system, version 1.8 (2015)
24. Stitz, H., Gratzl, S., Piringer, H., Zichner, T., Streit, M.: KnowledgePearls: Provenance-Based Visualization Retrieval. IEEE Transactions on Visualization and Computer Graphics (VAST '18) p. 11 (2018)
25. Streit, M., Gratzl, S., Stitz, H., Wernitznig, A., Zichner, T., Haslinger, C.: Ordino: a visual cancer analysis tool for ranking and exploring genes, cell lines and tissue samples (2019). DOI 10.1093/bioinformatics/btz009. URL `https://doi.org/10.1093/bioinformatics/btz009`
26. Streit, M., Lex, A., Gratzl, S., Partl, C., Schmalstieg, D., Pfister, H., Park, P.J., Gehlenborg, N.: Guided visual exploration of genomic stratifications in cancer. Nature Methods **11**(9), 884–885 (2014). DOI 10.1038/nmeth.3088
27. Wan, Y., Lewis, A., Colasanto, M., van Langeveld, M., Kardon, G., Hansen, C.: A practical workflow for making anatomical atlases in biological research. IEEE Computer Graphics and Applications **32**(5), 70–80 (2012). DOI 10.1109/MCG.2012.64. URL `http://www.sci.utah.edu/publications/wan12/Wan_CGA2012.pdf`