

DaVE - A Curated Database of Visualization Examples

Jens Koenen*[†]
RWTH Aachen University, Germany

Marvin Petersen*
RPTU Kaiserslautern-Landau, Germany

Christoph Garth
RPTU Kaiserslautern-Landau, Germany

Tim Gerrits
RWTH Aachen University, Germany




Figure 1: Schematic overview of a potential workflow with DaVE: A domain scientist acquires data, e.g., by running a simulation on an HPC cluster (↗). Based on this data, they enter search terms into DaVE, which provides a list of appropriate visualization techniques from its curated and extensible database for different data types and domains (↗). The user can freely explore the examples until a suitable visualization is chosen. Each example comes with a description, explanation of techniques, additional references, and adaptable containerized code. This allows the visualization to be calculated on suitable hardware using the respective data without the potentially cumbersome setup process (↗).

ABSTRACT

Visualization, from simple line plots to complex high-dimensional visual analysis systems, has established itself throughout numerous domains to explore, analyze, and evaluate data. Applying such visualizations in the context of simulation science where High-Performance Computing (HPC) produces ever-growing amounts of data that is more complex, potentially multidimensional, and multi-modal, takes up resources and a high level of technological experience often not available to domain experts. In this work, we present *DaVE* - a curated database of visualization examples, which aims to provide state-of-the-art and advanced visualization methods that arise in the context of HPC applications. Based on domain- or data-specific descriptors entered by the user, *DaVE* provides a list of appropriate visualization techniques, each accompanied by descriptions, examples, references, and resources. Sample code, adaptable container templates, and recipes for easy integration in HPC applications can be downloaded for easy access to high-fidelity visualizations. While the database is currently filled with a limited number of entries based on a broad evaluation of needs and challenges of current HPC users, *DaVE* is designed to be easily extended by experts from both the visualization and HPC communities.

Index Terms: Visualization, Curated Database, High-Performance Computing

*equal contributions

[†]email: koenen@vis.rwth-aachen.de

1 INTRODUCTION

Nowadays, developers of research software and simulations in various scientific domains must possess expertise in diverse fields. This includes not only domain-specific knowledge of models as well as their physical and mathematical background. Developers are expected to write highly optimized code and find appropriate ways to analyze their results. The steady advancements in those areas drastically increase the complexity of data and require adapted and new visualization techniques, both on a conceptual level (novel visualization techniques) and on a technological level (efficient HPC implementation of visualization techniques). While such approaches are being actively developed and discussed within the visualization community, domain experts and HPC users often lack the resources to follow and keep up-to-date with these developments. This can lead to undesired effects such as a high overhead introduced by finding, learning, and applying approaches from scratch, or being unaware of new and potentially very helpful visualization techniques. Moreover, on the technological side, while implementations of modern visualization algorithms suitable for HPC use exist, incorporating these into complex HPC application stacks requires substantial expertise and an investment of additional engineering effort. Thus, while domain scientists see the need and are very open to using more advanced visualization techniques, they often lack resources and references, to use them.

In this work, we present *DaVE*¹ - a curated database of visualization examples targeted at novice and experienced HPC users alike, that allows for easy access to state-of-the-art and advanced visualization methods. The database is accessible through a modern and fast web interface and includes detailed descriptions, sample media, references, and example implementations tailored to HPC

¹<https://dave-vis.com>

projects which allows for easy integration in existing projects and application stacks. Users receive guidance not only for approaches they know and want to get more information on, but also recommendations on best practices or alternative visualization methods. The underlying system is developed in such a way that curators and users can easily extend the database with their own information and examples to establish a platform of best practices, current developments, and support. As such, DaVE aims to function as a focal point and community resource for visualization within the national HPC community and beyond. The database has the potential to develop into a general resource for visualization regardless of application.

2 BACKGROUND AND RELATED WORK

Visualization has become an integral part of understanding, analyzing, and communicating complex data in virtually all scientific domains. Therefore, many data management tools such as Microsoft Excel, Matlab, or Tableau, already provide built-in visualization options for some data types, such as data tables, that make it easy for non-expert users to achieve fast and easy visualizations. Several projects have emerged in the past that aim to provide easy access to adaptable visualizations for similar or more complex data, such as the D3 library [6] or Vega-Lite [21]. A significant body of work exists on so-called content-based recommendation [18, 13, 25, 14], where systems try to automatically find the best-fitting visualization given a set of data. Naturally, this includes approaches where recent machine learning technologies are leveraged [12, 20], specific tasks are analyzed [23], or user behavior is evaluated [11, 8]. However, most of these are limited to 2D plots.

Simulations that model physical phenomena often leverage high-performance computing infrastructure for, e.g., efficient computation of computational fluid dynamic (CFD) models or force propagation in complex materials. Such data typically describes spatial, multimodal 3D data that might also be time-dependent. While some of that data can be summarized and described by statistical data, visual exploration usually requires efficient interaction with 3D geometries such as direct volume renderings or derived features of often very large data quantities. Besides efficient handling and processing, such data size might even require in-situ processing or rendering which further increases the complexity of visualization. Here, too, a number of tools and libraries exist, such as VTK [22] or ParaView [1], which provide a multitude of options such as filters and renderers. In the context of in-situ or in-transit visualizations, distinct visualization libraries such as Catalyst [2] or Ascent [16] are available. Learning these tools or integrating them into existing simulation code, however, requires time and knowledge. For libraries, such as D3, it has been shown that providing a database of examples and templates can considerably reduce a user's effort to design visualizations [4, 3]. VTK, too, has example code available on a dedicated website, and approaches such as EZ-ISAV [24] seek to make the integration of visualization in complex software stacks easier by providing containerized solutions. However, to the best of our knowledge, there is no publicly available resource that allows non-visualization experts to easily find suitable techniques primarily for HPC-based simulation data.

3 DAVE

For us, the need for an easily accessible resource for HPC users to discover visualization examples became apparent after collaboration with several simulation scientists within a large national research project. An internal user evaluation [10] revealed that HPC users across domains all appreciated more sophisticated state-of-the-art visualization techniques, such as in-situ/in-transit capabilities, visual analysis and exploration systems, and immersive visualizations. However, all lacked resources, guidance, and necessary knowledge to find and integrate suitable techniques. Recognizing this gap, we believe that a curated database of such techniques and

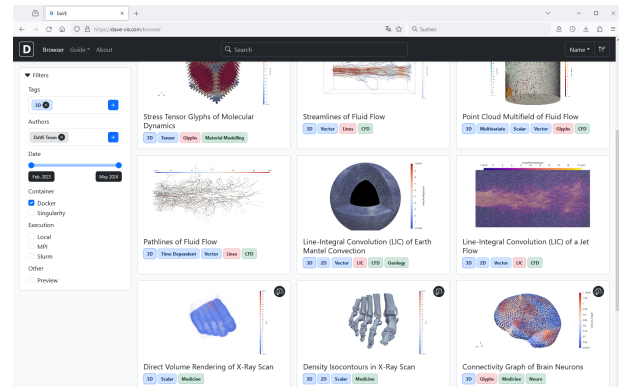


Figure 2: Through the adaptive web interface, DaVE provides a gallery of visualization examples that can be searched and filtered through tags and requirements.

examples can be a way to address these issues. Users also expressed concerns that testing and integrating such techniques requires substantial effort as well as changes to existing and often complicated code bases. After several discussions at workshops, conferences, and personal meetings, we identified three essential requirements for a successful resource. These can be summarized as follows:

- R1** A non-visualization expert-friendly system to find domain- or data-specific visualization techniques
- R2** Easy integration of visualization examples into existing HPC infrastructure
- R3** Easy extensibility of the database with new and relevant information

Based on these, we decided to create DaVE - a curated database of visualization examples. In the following, we will describe and justify the individual components and functionalities.

3.1 Web Platform for Visualization Exploration

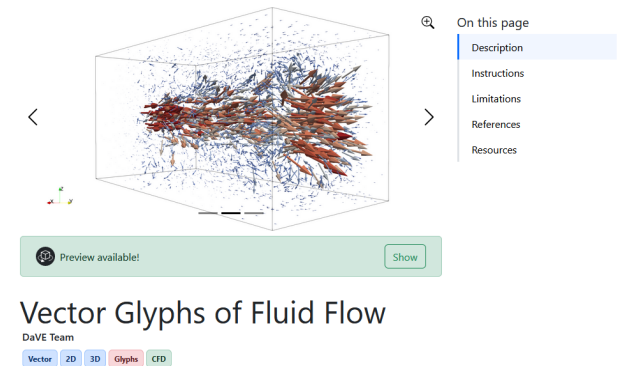


Figure 3: Tags describing data type, visualization technique, and domain of the example are included in the search. Example images, optional interactive previews, and text descriptions provide further information on the technique and usage of provided resources.

DaVE's main point of interaction is a responsive website. Users are prompted to enter a search term that describes their data or needs. This search is aided by suggestions of suitable keywords. From there the user is provided with a gallery overview of the visualization examples that fit the search description. For each example, a preview image and tags for categorization are displayed. In a

sidebar, these items can be further filtered by tags, date, author, and other parameters (see Figure 2). Navigating to one of the example pages (see Figure 3) shows an introductory text, covering the visualization methods, its use cases, and basic functionality. Further, short instructions and materials for reproducing the example are provided. Also, examples may offer an interactive preview of the visualization (see Figure 4). With this design, users directly interact with the system, to help them find solutions for their visualization problem before acquainting them with the needed technology (R1). A complete package for executing the example can be configured and downloaded, which is executed by calling a single script. Configuration options include the path for own datasets, the usage of MPI and Slurm, and the container technology, either Docker or Singularity/Apptainer [15], that is available on the target machine. Dedicated guide pages contain general information about DaVE, the structure of an example, and instructions on how to adapt existing or create custom examples in general. These pages also include references to the documentation of used software for further customization. Additionally, the guide pages specify how adapted and new examples can be added to DaVE, simplifying the contribution process (R3). Providing the executable templates and information allows for exploratory tinkering, and acquiring knowledge about the underlying technologies in the process. Each example has the option to open an issue on the project github that allows to provide feedback or questions if clarifications are needed.

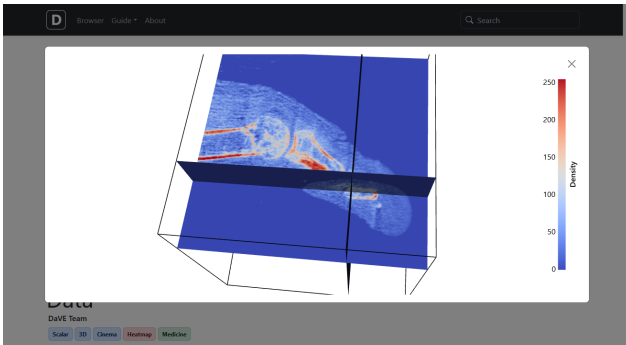


Figure 4: DaVE supports interactive examples on the websites through VTK.js³, allowing users to directly test visualizations.

3.2 Integration Through Containerized Examples

To be able to allow users working on a large variety of hardware configurations, from personal computers to large high-performance computing clusters, to test and explore the examples, DaVE provides containerized environments. This has the following advantages: A containerized environment can easily be reused for multiple examples, reducing the download time if multiple similar examples are explored. Furthermore, these containers offer high portability. Most examples that are currently available via DaVE can be executed locally, with MPI, or via Slurm on HPC clusters (R2). Additionally, extending containerized environments is straightforward. The containers can be extended by using them as a base for a new custom image or users can change the build recipe in the original directly (R3). However, and most importantly, the provided containers do not have to be extended or changed at all to enable using the given examples. This allows users from different domains with varying background knowledge to explore visualization in HPC environments from fully functional examples and start from there instead of having to build a reference use case themselves (R1).

³<https://github.com/Kitware/vtk-js>

With respect to container technologies, DaVE relies on Docker for creating containers. The images can be uploaded to a container image library, like Docker Hub. From there, the executable package of an example in DaVE retrieves the image either via Docker itself or with singularity. In order to build maintainable environments inside the containers, the required software and its dependencies are configured and built with spack [9]. It is a flexible package manager that streamlines the process of managing software dependencies, making it easier to install, configure, and maintain projects, which, in turn, reduces time and effort for creating examples.

3.3 Extending DaVE

DaVE supports the easy addition of new techniques and examples. Entries of the database are built upon a straightforward file and folder structure. Each example is defined by a folder on DaVE’s workspace which contains text files including textual explanations, as well as descriptive tags and meta information. Subfolders contain further downloadable resources and files describing options for container creation. Adding a new example is as simple as adding a new folder at the respective location. This can either be done by the original authors by request, or by submitting a merge request on the project github. This allows and encourages an active community to maintain the quality and correctness of the entries.

As mentioned in Section 3.1, a dedicated page on DaVE further provides all necessary information, instructions, and downloadable templates to create own examples or alter existing containers to work with additional software. There are several levels of complexity for adding new entries, allowing users with different backgrounds and varying familiarity with the technologies used to contribute (R1,R3). The simplest option is to adapt an existing example to new data. This requires no changes in the containerized environment and only minor changes in the visualization procedure. Creating custom visualizations is the next option. Again no changes in the environment are needed, but more knowledge about the visualization software is required. In the case of ParaView, users need to know how to create and manipulate ParaView’s Python trace files. If no available container provides the necessary tools for a new example, existing containers can be extended either via docker or reconfigured and built with spack. Since managing software for use with HPC infrastructure can be a daunting task for non-experts, spack provides easy configuration capabilities for the existing containers. This setup allows users to acquire knowledge on specific areas of visualization in HPC in a step-by-step manner. Our provided examples serve as a starting point to dive deeper into the respective techniques, like Docker, spack, or ParaView. For the greatest flexibility, fully custom docker images can be used. As such, we not only hope to establish DaVE as a standard resource for HPC users but also to encourage active feedback and development by the respective communities.

4 EVALUATION AND EXPERT FEEDBACK

To get an indication, of whether our current design of DaVE fulfills the requirements outlined in Section 3, we invited five simulation scientists from different research institutions working on various domains to evaluate DaVE. None of the participants had been involved in the project before but all were frequent users of HPC infrastructure with research experience ranging from 3 to 10 years. Domains of expertise included computational fluid dynamics, material design, artificial intelligence, and energy conversion. All but one had worked with containers before and considered themselves familiar with containerized code. Participants were asked to freely explore DaVE on their own devices, either under observation at their workplace or via Zoom recording. Sessions typically lasted between 15 and 40 minutes, during which users were asked to verbalize their thoughts and actions using a think-aloud strategy [17]. We further prepared a short questionnaire based on the System Us-

ability Score (SUS) [7] and additional questions that participants received after they had finished their exploration of DaVE. At the time of the evaluation, the database contained 17 curated examples across various domains added by the authors. Participants used their desktop work environment and standard modern web browsers such as Firefox, Chrome, or Safari.

Users started on the landing page and understood the core functionality of DaVE from reading the provided overviews and the use of a standard search bar. The simplicity, modern design, and clear focus on the search functionality were mentioned and appreciated by all. While one user opted for entering a domain-specific search term in the search bar, others chose to explore all examples. Entering the gallery, participants quickly understood the tagging/filtering system, added and removed tags to test it out, and mentioned them as an intuitive and very useful method to browse the database. Users would scan the preview images and example names to look for familiar data before using the tags to narrow down results. Some users found the limited number of examples within their specific domain unfortunate but viewed this as a temporary limitation that could be overcome through active community contributions. Once an example was selected, users especially pointed out the usefulness of interactive previews and focused descriptions. Some suggested more interactive previews, e.g., allowing to apply filters or change color mappings. After experimenting with different options, users would return to the gallery to look for more examples. Only one user configured, downloaded, and started a containerized example, but did not need any instructions to do so. Questions that users had during interaction with DaVE were always answered by either experimenting with features or consulting the provided information on the website. Minor comments on design decisions as well as various feature ideas were verbalized throughout the process. After users claimed they were done exploring, all stated they had enjoyed exploring DaVE. They appreciated the overall idea and effort invested in the system, and recognized its potential, especially with the addition of more examples.

Evaluating the questionnaire, DaVE received a mean system usability score of 83.5, classified as “*excellent*” according to Bangor et al. [5]. Users agreed that DaVE made it *very easy* to *extremely easy* to find domain-data-specific visualization techniques, indicating fulfillment of **R1**. All but one reported discovering new visualization approaches, previously unknown to them. Evaluating **R2** proved less conclusive: users were *fairly confident* that techniques could be integrated into their own work but reported needing more time to test it. It was highly appreciated that containerized environments were provided, given participants’ familiarity with them, and using ParaView was mentioned to be a reasonable choice, due to its recognition. However, some hoped for more examples using alternative tools to accommodate other workflows. Finally, users reported that extending DaVE (**R3**) appears to be *fairly easy*, but considered themselves more as users of DaVE than as contributors. They all reported that the level of detail in the information section should suffice. Yet, one participant saw the potential in DaVE to replace the often fragmented and decentralized collections of code and information within different research groups.

5 DISCUSSION

We carefully designed DaVE to address each of the three requirements comprehensively. Domain- and data-specific search options and examples directly showcase applicable visualization techniques for non-experts (**R1**). Tailored configuration options for containerized environments specifically designed for HPC ensure portability and easy integration into existing infrastructure (**R2**). A simple and intuitive example structure, with human-readable configuration files, annotations, and concise instructions for adaptation and contribution, ensures the extensibility of the system (**R3**). While we are aware, that the number of participants in the expert interview is

limited and provides only initial insights, the results are promising. Due to the nature of free exploration of DaVE, some features were less thoroughly explored than others. Ideally, we hope to gather more feedback once the resource is available to a broader audience and has a larger number of examples added, enabling long-term refinement. Currently, our examples are based on ParaView implementations, but we aim to cover a broader variety of approaches in the future. The current strategy for establishing DaVE as a general resource for many communities consists of growing and validating the approach inside collaborations in a national project before expanding its scope to use cases from interested groups outside. Maintaining such a resource is a multifaceted challenge. Ensuring continued relevance of information in DaVE involves updating outdated examples and adding new ones for emerging techniques. With a growing database, this task becomes insurmountable for a small group of maintainers but should be manageable with active community participation. With increasing user and contributor base, the diversity of target environments is likely to increase as well. Testing the entries of DaVE against every possible use case seems unpractical, but would greatly improve the relevancy and usefulness of the resource. Currently, we rely on a community-driven approach, where problems with examples can be reported in the issue tracker of the accompanying repository. Furthermore, it has to be ensured that the technical infrastructure, including the website and repository, is kept up-to-date and working while highly requested features have to be integrated. The longevity of this resource is ensured through a national long-term project dedicated to building and maintaining such infrastructure.

6 CONCLUSION AND FUTURE WORK

In this work, we introduced DaVE - a curated database of visualization examples, which aims to provide domain scientists working on HPC infrastructure a simple way to find, explore, and apply state-of-the-art visualization techniques in their projects. After evaluating challenges and needs of users within a large national HPC research collaboration, we designed DaVE based on three requirements: A user-friendly system to browse and find visualizations, adaptable containerized examples for easy integration across diverse hardware configurations, and the option to easily extend the database, such that it can become a community-driven project. It is accessible via a simple yet easy-to-navigate web interface and provides information and resources based on user search queries. While the current example size is limited, a preliminary evaluation with expert users yielded promising results, pointing toward the utility of DaVE. Now that DaVE is publicly available, we hope that it results in more requests from users as well as support from the visualization community to help it grow.

In the future, we will continue working on DaVE by incorporating additional and extending existing examples. While the expert evaluation already led to a list of potential improvements, we plan to incorporate further feedback from users for possible design and usability enhancements as well as new features: Automation processes might be used to regularly find and update entries similar to Oppermann et al. [19]. A search wizard might guide users to more precise intention statements, allowing for more appropriate search terms and technique selections. We will also explore how machine-learning-based recommendation systems might benefit the system. Further, we are exploring the idea of developing an API that allows developers to get access to search results and resources, allowing the integration of DaVE results into existing software stacks.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the German Federal Ministry of Education and Research (BMBF) and the state government of NRW and RP for supporting this work/project as part of the NHR funding.

REFERENCES

- [1] J. Ahrens, B. Geveci, and C. Law. *Visualization Handbook*. Elsevier Inc., Burlington, MA, USA, 2005. 2
- [2] U. Ayachit, A. Bauer, B. Geveci, P. O’Leary, K. Moreland, N. Fabian, and J. Mauldin. Paraview catalyst: Enabling in situ data analysis and visualization. In *Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*, ISAV2015, p. 25–29. Association for Computing Machinery, New York, NY, USA, 2015. doi: 10.1145/2828612.2828624 2
- [3] H. Bako, A. Varma, A. Faboro, M. Haider, F. Nerrise, B. Kenah, and L. Battle. Streamlining visualization authoring in d3 through user-driven templates. In *2022 IEEE Visualization and Visual Analytics (VIS)*, pp. 16–20, 2022. doi: 10.1109/VIS54862.2022.00012 2
- [4] H. K. Bako, X. Liu, L. Battle, and Z. Liu. Understanding how designers find and use data visualization examples. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1048–1058, 2023. doi: 10.1109/TVCG.2022.3209490 2
- [5] A. Bangor, P. Kortum, and J. Miller. Determining what individual sus scores mean: adding an adjective rating scale. *J. Usability Studies*, 4(3):114–123, may 2009. 4
- [6] M. Bostock, V. Ogievetsky, and J. Heer. D³: Data-driven documents. *IEEE transactions on visualization and computer graphics*, 17:2301–9, 12 2011. doi: 10.1109/TVCG.2011.185 2
- [7] J. Brooke et al. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996. 4
- [8] W. EPPerson, D. Jung-Lin Lee, L. Wang, K. Agarwal, A. G. Parameswaran, D. Moritz, and A. Perer. Leveraging analysis history for improved in situ visualization recommendation. vol. 41, pp. 145–155, 2022. doi: 10.1111/cgf.14529 2
- [9] T. Gamblin, M. LeGendre, M. R. Collette, G. L. Lee, A. Moody, B. R. de Supinski, and S. Futral. The spack package manager: bringing order to hpc software chaos. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC ’15. Association for Computing Machinery, New York, NY, USA, 2015. doi: 10.1145/2807591.2807623 3
- [10] T. Gerrits. Bringing Visualization to National HPC Infrastructure: A Prologue, June 2022. doi: 10.5281/zenodo.7715663 2
- [11] D. Gotz and Z. Wen. Behavior-driven visualization recommendation. In *Proceedings of the 14th International Conference on Intelligent User Interfaces*, p. 315–324. Association for Computing Machinery, New York, NY, USA, 2009. doi: 10.1145/1502650.1502695 2
- [12] K. Hu, M. A. Bakker, S. Li, T. Kraska, and C. Hidalgo. Vizml: A machine learning approach to visualization recommendation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, p. 1–12. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3290605.3300358 2
- [13] P. Kaur. and M. Owonibi. A review on visualization recommendation strategies. In *Proceedings of the 12th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2017) - IVAPP*, pp. 266–273. INSTICC, SciTePress, 2017. doi: 10.5220/0006175002660273 2
- [14] P. Kubernátová, M. Friedjungová, and M. van Duijn. Knowledge at first glance: A model for a data visualization recommender system suited for non-expert users. In *Proceedings of the 7th International Conference on Data Science, Technology and Applications*, p. 208–219. SCITEPRESS - Science and Technology Publications, Lda, Setubal, PRT, 2018. doi: 10.5220/0006851302080219 2
- [15] G. M. Kurtzer, V. Sochat, and M. W. Bauer. Singularity: Scientific containers for mobility of compute. *PLOS ONE*, 12(5):1–20, 05 2017. doi: 10.1371/journal.pone.0177459 3
- [16] M. Larsen, J. Ahrens, U. Ayachit, E. Brugger, H. Childs, B. Geveci, and C. Harrison. The alpine in situ infrastructure: Ascending from the ashes of strawman. In *Proceedings of the In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*, ISAV’17, p. 42–46. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3144769.3144778 2
- [17] C. Lewis. *Using the “thinking-aloud” method in cognitive interface design*. IBM TJ Watson Research Center Yorktown Heights, NY, 1982. 3
- [18] M. Oppermann, R. Kincaid, and T. Munzner. Vizcommender: Computing text-based similarity in visualization repositories for content-based recommendations. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):495–505, 2021. doi: 10.1109/TVCG.2020.3030387 2
- [19] M. Oppermann and T. Munzner. Vizsnippets: Compressing visualization bundles into representative previews for browsing visualization collections. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):747–757, 2022. doi: 10.1109/TVCG.2021.3114841 4
- [20] X. Qian, R. A. Rossi, F. Du, S. Kim, E. Koh, S. Malik, T. Y. Lee, and J. Chan. Learning to recommend visualizations from data. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, p. 1359–1369. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3447548.3467224 2
- [21] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-lite: A grammar of interactive graphics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):341–350, 2017. doi: 10.1109/TVCG.2016.2599030 2
- [22] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit (4th ed.)*. Kitware, 2006. 2
- [23] L. Shen, E. Shen, Z. Tai, Y. Song, and J. Wang. TaskVis: Task-oriented Visualization Recommendation. In M. Agus, C. Garth, and A. Kerren, eds., *EuroVis 2021 - Short Papers*. The Eurographics Association, 2021. doi: 10.2312/evs.20211061 2
- [24] M. Will, Q. Wofford, J. Patchett, D. Rogers, J. Lukasczyk, and C. Garth. Developing and evaluating in situ visualization algorithms using containers. In *ISAV’21: In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*, ISAV’21, p. 6–11. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3490138.3490141 2
- [25] Z. Zeng, P. Moh, F. Du, J. Hoffswell, T. Y. Lee, S. Malik, E. Koh, and L. Battle. An evaluation-focused framework for visualization recommendation algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):346–356, 2022. doi: 10.1109/TVCG.2021.3114814 2