

# scenery: Flexible Virtual Reality Visualization on the Java VM

Ulrik Günther\*  
Technische Universität Dresden  
Center for Systems Biology Dresden  
MPI-CBG, Dresden

Kyle I.S. Harrington  
University of Idaho  
Howard Hughes Medical Institute,  
Janelia Research Campus

Tobias Pietzsch  
Center for Systems Biology Dresden  
MPI-CBG, Dresden

Pavel Tomancak  
MPI-CBG, Dresden  
IT4Innovations, VŠB - Technical  
University of Ostrava

Ivo F. Sbalzarini†  
Technische Universität Dresden  
Center for Systems Biology Dresden  
MPI-CBG, Dresden

Aryaman Gupta  
Technische Universität Dresden  
Center for Systems Biology Dresden  
MPI-CBG, Dresden

Stefan Gumhold  
Technische Universität Dresden

## ABSTRACT

Life science today involves computational analysis of a large amount and variety of data, such as volumetric data acquired by state-of-the-art microscopes, or mesh data from analysis of such data or simulations. Visualization is often the first step in making sense of data, and a crucial part of building and debugging analysis pipelines. It is therefore important that visualizations can be quickly prototyped, as well as developed or embedded into full applications. In order to better judge spatiotemporal relationships, immersive hardware, such as Virtual or Augmented Reality (VR/AR) headsets and associated controllers are becoming invaluable tools. In this work we introduce *scenery*, a flexible VR/AR visualization framework for the Java VM that can handle mesh and large volumetric data, containing multiple views, timepoints, and color channels. *scenery* is free and open-source software, works on all major platforms, and uses the Vulkan or OpenGL rendering APIs. We introduce *scenery*'s main features and example applications, such as its use in VR for microscopy, in the biomedical image analysis software Fiji, or for visualising agent-based simulations.

**Index Terms:** Human-centered computing—Visualization—Visualization systems and tools Human-centered computing—Virtual reality

## 1 INTRODUCTION

Recent innovations in biology, like lightsheet microscopy [13], or Serial Block-Face Scanning Electron Microscopy [7] are now making large, spatiotemporally complex volumetric data available. However, the data acquired by the microscope is only a means to an end: researchers need to extract results from it, and for that efficient tools are needed. This includes tools that not only enable the researcher to visualize their data, but to interact with it, and to enable them to use the tool in ways the original designer had not anticipated.

For this purpose, we introduce *scenery*, a flexible, open-source visualization framework for the Java Virtual Machine (JVM) that can handle mesh data (e.g. from triangulated surfaces), and multi-channel, multi-timepoint, multi-view

\*e-mail: guenther@mpi-cbg.de

†e-mail: ivos@mpi-cbg.de

volumetric data of large size<sup>1</sup>. Our main contribution with *scenery* is to combine all of the following design goals into one reusable, open-source framework:

- G<sub>1</sub> **Virtual/Augmented Reality support:** The framework should make the use of VR/AR in an application possible with minimal effort. Distributed systems, such as CAVEs or Powerwalls, should also be supported.
- G<sub>2</sub> **Out-of-core volume rendering:** The framework should be able to handle datasets that do not fit into graphics memory and/or main memory, contain multiple channels, views, and timepoints. It should be possible to visualize multiple such datasets in a single scene.
- G<sub>3</sub> **User/Developer-friendly API:** The framework should have a simple API that makes only limited use of advanced features, such as generics, so the user/developer can quickly comprehend and customize it.
- G<sub>4</sub> **Cross-platform:** The framework should run on the major operating systems: Windows, Linux, and macOS.
- G<sub>5</sub> **JVM-native and embeddable:** The framework should run natively on the JVM, and be embeddable, such that it can be used in popular biomedical image analysis tools like Fiji [29,30], Icy [5], and KNIME [1].

## 2 RELATED WORK

A particularly popular *framework* in scientific visualization is VTK [10]: VTK offers rendering of both geometric and volumetric data, using an OpenGL 2.1 renderer. However, VTK's complexity has also grown over the years and its API is becoming more complex, making it difficult to change internals without breaking existing applications (G<sub>3</sub>). A more recent development is *VTK.js*, which brings VTK to web browsers. *ClearVolume* [28] is a visualization toolkit tailored to high-speed, volumetric microscopy and supports multi-channel/multi-timepoint data, but focuses solely on volumetric data and does not support VR/AR. *MegaMol* [9] is a special-purpose framework focused on efficient rendering of a large number of discrete particles that provides a thin abstraction layer over the graphics API for the developer. *3D Viewer* [31] does general-purpose image visualization tasks, and supports multi-timepoint data, but no out-of-core volume rendering, or VR/AR.

<sup>1</sup>Out-of-core data is stored in tiles, with 64 bit tile indices, and each tile comprising up to 2<sup>31</sup> voxels. Therefore the theoretical limit for a single volume is 2<sup>94</sup> voxels, roughly corresponding to a cube with 2.1 billion voxels edge length, equal to 20000 Yottabyte. The largest tested dataset was an 8TB multi-angle time-series, with 7GB per timepoint.

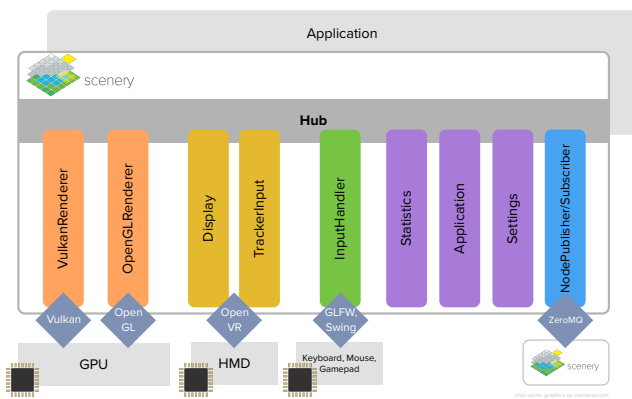


Figure 1: Overview of scenery's architecture.

In *out-of-core rendering* (OOCR), the rendering of volumetric data that does not fit into main or graphics memory, existing software packages include *Vaa3D/Terafly* [4, 22], which is written with applications like neuron tracing in mind, and *BigDataViewer* [25], which performs by-slice rendering of large datasets, powered by the *ImgLib2* library [24]. The *VR neuron tracing tool* [33] supports OOCR, but lacks support for multiple timepoints and is not customizable. *Invivo* [15] supports OOCR and interactive development, but does not support overlaying multiple volumetric datasets in a single view.

In the field of biomedical image analysis, various *commercial packages* exist: *Arivis*, *Amira*, and *Imaris*<sup>2</sup>, and *syGlass* [23] support out-of-core rendering, and are scriptable by the user. *Arivis*, *Imaris*, and *syGlass* offer rendering to VR headsets, while *Amira* can run on CAVE systems. *Imaris* provides limited Fiji and Matlab integration. Due to being closed-source, the flexibility of these packages is ultimately limited (*e.g.*, changing rendering methods, or adding new input devices).

### 3 SCENERY

With *scenery*, we provide a flexible framework for developing visualization prototypes and applications, on systems ranging from desktop screens, VR/AR headsets (like the Oculus Rift or HTC Vive), to distributed setups. *scenery* is written in Kotlin, a language for the JVM that requires less boilerplate code and has more functional constructs than Java itself. This increases developer productivity, while maintaining 100% compatibility with existing Java code. *scenery* runs on Windows, Linux, and macOS ( $G_4$ ). *scenery* uses the low-level Vulkan API for fast and efficient rendering, and can fall back to an OpenGL 4.1-based renderer<sup>3</sup>.

*scenery* is designed around two concepts: A scene graph for the scene organisation into nodes, and a hub organizing all subsystems — *e.g.* rendering, input, statistics — and enables communication between them. *scenery*'s application architecture is depicted in Fig. 1. *scenery*'s subsystems are only loosely coupled, meaning they can work fully independent of each other. The loose coupling enables isolated testing of the subsystems, and thereby we can reach 65% code coverage at the moment (the remaining 35% is mostly code that requires additional hardware and is therefore harder to test in an automated manner).

<sup>2</sup>See [arivis.com/en/imaging-science/imaging-science](http://arivis.com/en/imaging-science/imaging-science), [fei.com/software/amira/](http://fei.com/software/amira/), and [imaris.oxinst.com](http://imaris.oxinst.com)

<sup>3</sup>The Vulkan renderer uses the LWJGL Vulkan bindings (see [lwjgl.org](http://lwjgl.org)), while the OpenGL renderer uses JOGL (see [jogamp.org](http://jogamp.org)).

## 4 HIGHLIGHTED FEATURES

### 4.1 Realtime rendering on the JVM — $G_5$

Historically, the JVM has not been the go-to target for realtime rendering: For a long time, the JVM had the reputation of being slow and memory-hungry. However, since the HotSpot VM has been introduced in Java 6, this is less true, and state-of-the-art just-in-time compilers like the ones used in Java 12 have become very good at generating automatically vectorized code<sup>4</sup>. The JVM is widely used, provides excellent dependency management via the *Maven* or *Gradle* build tools, and efficient, easy-to-use abstractions for, *e.g.*, multithreading or UIs on different operating systems. Additionally, with the move to low-overhead APIs like Vulkan, pure-CPU performance is becoming less important. In the near future, *Project Panama*<sup>5</sup> will introduce JVM-native vectorization primitives to support CPU-heavy workloads. These primitives will work in a way similar to those provided by .NET.

Another convenience provided by the JVM is scripting: Via the JVM's scripting extensions, *scenery* can be scripted using its REPL with third-party languages like Python, Ruby, and Clojure. In the future, *GraalVM*<sup>6</sup> will enable polyglot code on the JVM, *e.g.* by ingesting LLVM bytecode directly [3]. *scenery* has already been tested with preview builds of both GraalVM and Project Panama.

### 4.2 Out-of-core volume rendering — $G_2$

*scenery* supports volume rendering of multiple, potentially overlapping volumes that are placed into the scene via arbitrary affine transforms. For out-of-core direct volume rendering of large volumes ( $G_2$ ) we develop and integrate the *BigVolumeViewer* library, which builds on the pyramidal image data structures and in-memory caching of large image data from *BigDataViewer* [25]. We augment this by a GPU cache tier for volume blocks, implemented using a single large 3D texture. This cache texture is organized into small (*e.g.*,  $32^3$ ) uniformly sized blocks. Each texture block stores a particular block of the volume at a particular level in the resolution pyramid, padded by one voxel on each side to avoid bleeding from neighboring blocks during trilinear interpolation [2]. The mapping between texture and volume blocks is maintained on the CPU.

To render a particular view of a volume, we determine a base resolution level such that screen resolution is matched for the nearest visible voxel. Then, we prepare a 3D lookup texture in which each voxel corresponds to a volume block at base resolution. Each voxel in this lookup texture stores the coordinates of a block in the cache texture, as well as its resolution level relative to base, encoded as a RGBA tuple. For each (visible) volume block, we determine the optimal resolution by its distance to the viewer. If the desired block is present in the cache texture, we encode its coordinates in the corresponding lookup texture voxel. Otherwise, we enqueue the missing cache block for asynchronous loading through the CPU cache layer of *BigDataViewer*. Newly loaded blocks are inserted into the cache texture, where the cache blocks to replace are determined by a least-recently-used strategy that is also maintained on the CPU. For rendering, currently missing blocks are substituted by lower-resolution data if it

<sup>4</sup>For this project, we have measured the timings of performance-critical parts of code, such as  $4 \times 4$  matrix multiplication. Compared to hand-tuned, vectorized AVX512 code, the native code generated by the JVM's JIT compiler is about a factor of 3-4 slower.

<sup>5</sup>See [openjdk.java.net/projects/panama](http://openjdk.java.net/projects/panama).

<sup>6</sup>See [graalvm.org](http://graalvm.org).

is available from the cache. Intermittently, tiles may render at lower resolution or be missing completely. We prioritize maintaining interactive framerates over rendering the most complete data. Our technique is a combination of hierarchical blocking [2, 20] and the missing data scheme of [25].

Once the lookup texture is prepared, volume rendering proceeds by raycasting and sampling volume values with varying step size along the ray, adapted to the viewer distance. To obtain each volume sample, we first downscale its coordinate to fall within the correct voxel in the lookup texture. A nearest-neighbor sample from the lookup texture yields a block offset and scale in the cache texture. The final value is then sampled from the cache texture with the accordingly translated and scaled coordinate. With this approach, it is straightforward to raycast through multiple volumes simultaneously, simply by using multiple lookup textures. It is also easy to mix in smaller volumes which are simply stored as 3D textures and do not require indirection via lookup textures. To adapt to varying number and type of visible volumes, we generate shader sources dynamically at runtime. Blending of volume and mesh data is achieved by reading scene depth from the depth buffer for early ray termination, thereby hiding volume values that are behind rendered geometry.

### 4.3 Code-shader communication and reflection — $G_3$

In traditional OpenGL (before version 4.1), parameter data like vectors, matrices, etc. are communicated to shaders via uniforms, which are set one-by-one. In scenery, instead of single uniforms, Uniform Buffer Objects (UBOs) are used. UBOs lead to a lower API overhead and enable variable update rates. Custom properties defined for node classes that need to be communicated to the shader are annotated in the class definition with the `@ShaderProperty` annotation, scenery picks up annotated properties automatically, and serializes them. See Listing 1 for an example of how properties can be communicated to the shader, and Listing 2 for the corresponding GLSL code for UBO definition in the shader. Procedurally-generated shaders can use a hash map storing these properties.

For all values stored in shader properties a hash is calculated, and they are only communicated to the GPU when the hash changes. Currently, all elementary types (ints, floats, etc.), as well as matrices and vectors thereof, are supported.

Listing 1: Shader property example

```
// Define a matrix and an integer property
@ShaderProperty var myMatrix: GLMatrix
@ShaderProperty var myIntProperty: Int
// For a dynamically generated shader: Store ←→
// properties as hash map
@ShaderProperty val shaderProperties = HashMap<←→
String, Any>()
```

Listing 2: GLSL code example for shader properties

```
layout(set = 5, binding = 0)
uniform ShaderProperties {
    int myIntProperty;
    mat4 myMatrix;
};
```

Determination of the correct memory layout required by the shader is done by our Java wrapper for the shader reflection library *SPIRV-cross* and the GLSL reference compiler *glslang*<sup>7</sup>. This provides a user- and developer-friendly API ( $G_3$ ).

<sup>7</sup>See [github.com/KhronosGroup/SPIRV-cross](https://github.com/KhronosGroup/SPIRV-cross) and [github.com/scenerygraphics/spirvcrossj](https://github.com/scenerygraphics/spirvcrossj) for our wrapper, *spirvcrossj*.



Figure 2: A scientist interactively explores a 500 GiB multi-timepoint dataset of the development of an embryo of the fruit fly *Drosophila melanogaster* in the CAVE at the CSBD using a scenery-based application. Dataset courtesy of Loïc Royer, MPI-CBG/CZI Biohub, and Philipp Keller, HHMI Janelia Farm [27].

Furthermore, scenery supports shader factories — classes that dynamically produce shaders to be consumed by the GPU — and use them, e.g., when multiple volumetric datasets with arbitrary alignment need to be rendered in the same view.

### 4.4 Custom rendering pipelines — $G_3$

In scenery, the user can use custom-written shaders and assign them on a per-node basis in the scene graph. In addition, scenery allows for the definition of fully customizable rendering pipelines. The rendering pipelines are defined in a declarative manner in a YAML file, describing render targets, render passes, and their contents. Render passes can have properties that are adjustable during runtime, e.g., for adjusting the exposure of a HDR rendering pass. Rendering pipelines can be exchanged at runtime, and do not require a full reload of the renderer — e.g., already loaded textures do not need to be reloaded.

The custom rendering pipelines enable the user/developer to quickly switch between different pipelines, thereby enabling rapid prototyping of new rendering pipelines. We hope that this flexibility stimulates the creation of custom pipelines, e.g., for non-photorealistic rendering, or novel applications, such as Neural Scene (De)Rendering [21, 34].

### 4.5 VR and preliminary AR support — $G_1$

Recent reviews, e.g., [32], summarize how the use of VR/AR can lead to improved perception and navigation. Motivated by these observations, scenery supports rendering to VR headsets via the OpenVR/SteamVR library and rendering on distributed setups, such as CAVEs or Powerwalls — addressing  $G_1$ . The modules supporting different VR devices can be exchanged quickly and at runtime, as all of these implement a common interface. The use of Vulkan in turn enables improved rendering performance compared to older APIs.

In the case of distributed rendering, one machine is designated as master, to which multiple clients can connect. We use the same hashing mechanism as described in Section 4.3 to determine which node changes need to be communicated over the network, use Kryo<sup>8</sup> for fast serialization of the changes, and finally ZeroMQ for low-latency and resilient network communication. A CAVE usage example is shown in Fig. 2.

We have also developed an experimental compositor that enables scenery to render to the Microsoft HoloLens.

<sup>8</sup>See [github.com/EsotericSoftware/Kryo](https://github.com/EsotericSoftware/Kryo).




Figure 3: Agent-based simulation with 10,000 agents collectively forming a sphere.

#### 4.6 Remote rendering and headless rendering

To support downstream image analysis and usage settings where rendering happens on a powerful, but non-local computer, scenery can stream rendered images out, either as raw data or as H264 stream, which can be saved to disk or streamed over the network via RTP. All produced frames are buffered and processed in a separate coroutine, such that rendering performance is not impacted.

scenery can run in headless mode, creating no windows, enabling both remote rendering on machines that do not have a screen, e.g., in a cluster setup, or easier integration testing. Most examples provided with scenery can be run automatically (see the `ExampleRunner` test) and store screenshots for comparison. In the future, broken builds will be automatically identified by comparisons against known good images.

### 5 EXAMPLE APPLICATIONS

#### 5.1 VR control of microscopes

We have used scenery to study VR control and visualization for state-of-the-art volumetric microscopes. In one study with 8 microscopy experts we found that users reported an improvement compared to current 2D controls, due to enhanced perception in VR. We also investigated whether users tend to suffer from motion sickness during use of our interfaces. We found an average SSQ score [16] of  $6.2 \pm 6.7$ , which is very low, indicating that users tolerated VR rendering of live microscopy data and interaction with it well. The interface used in the study is shown in Supplementary Video 1.

#### 5.2 Agent-based simulations

We have utilized scenery to visualize agent-based simulations with large numbers of agents. By adapting an existing agent- and physics-based simulation toolkit [12], we have increased the number of agents that can be efficiently visualized by a factor of 10. This performance improvement enables previous studies of swarms with evolving behaviors to be revisited under conditions that may enable new levels of emergent behavior [8, 11]. In Figure 3, we show 10,000 agents using flocking rules inspired by [26] to collectively form a sphere.

#### 5.3 sciview

On top of scenery, we have developed a plugin for embedding in Fiji/ImageJ2 [29] — sciview, fulfilling  $G_5$ . We hope it will boost the use of VR technology in the life sciences, by enabling the user to quickly prototype visualizations and add new functionality. In sciview, many aspects of the UI are automatically generated, including the node property inspector




Figure 4: Out-of-core dataset of a *D. melanogaster* embryo visualised with scenery/sciview. The image is a composite of three different volumetric views, shown in different colors. The transfer function on the left was adjusted to highlight volume boundaries. Dataset courtesy of Michael Weber, Huisken Lab, MPI-CBG/Morgridge Institute.

and the list of Fiji plugins and commands applicable to the currently active dataset. sciview has been used in a recent lightsheet microscopy pipeline [6]. In Supplementary Video 2, we show sciview rendering three overlaid volumes from a fruit fly embryo, a still frame of that is shown in Figure 4.

### 6 CONCLUSIONS AND FUTURE WORK

We have introduced scenery, an extensible, user/developer-friendly rendering framework for geometric and large volumetric data and demonstrated its applicability in several use cases. Compared to previous solutions, scenery combines the aspects of virtual reality rendering and control, with out-of-core rendering of multiple volumetric datasets in the same view, and enables the user to design their own prototypes and applications. To our knowledge, scenery is also the first framework using Vulkan on the JVM. Although scenery has undergone significant development, it is still relatively early in development compared to more mature tools and does not possess the breadth of features that are present in some alternative frameworks. However, this limitation is relaxed by compatibility with Fiji/ImageJ2, which provides a wide range of image processing capabilities.

In the future, we will introduce better volume rendering algorithms (e.g. [14, 17]) and investigate their applicability to VR settings. Furthermore, we are looking into providing support for out-of-core mesh data, e.g. using sparse voxel octrees [18, 19]. On the application side, we are driving forward projects in microscope control (see Section 5.1) and VR/AR augmentation of laboratory experiments.

### 7 SOFTWARE AND CODE AVAILABILITY

scenery, its source code, and a variety of examples are available at [github.com/scenerygraphics/scenery](https://github.com/scenerygraphics/scenery) and are licensed under the LGPL 3.0 license. A preview of the Fiji plugin sciview is available at [github.com/scenerygraphics/sciview](https://github.com/scenerygraphics/sciview).

### ACKNOWLEDGEMENTS

The authors thank C. Rueden, M. Weigert, R. Haase, V. Ulman, P. Hanslovsky, W. Büschel, V. Leite, and G. Barbieri for additional contributions, L. Royer, P. Keller, N. Maghelli, and M. Weber for allowing use of their datasets, and I. Tsakpinis and K. Burjack from the LWJGL community for their support. This work was supported by the European Regional Development Fund, project number CZ.02.1.01/0.0/0.0/16.013/0001791.

## REFERENCES

- [1] M. R. Berthold, N. Cebon, F. Dill, T. R. Gabriel, T. Kttr, T. Meinel, P. Ohl, C. Sieb, K. Thiel, and B. Wiswedel. KNIME: The Konstanz Information Miner. *Springer Berlin Heidelberg*, pp. 319–326, 2001.
- [2] J. Beyer, M. Hadwiger, T. Möller, and L. Fritz. Smooth mixed-resolution gpu volume rendering. In *Proceedings of the Fifth Eurographics / IEEE VGTC Conference on Point-Based Graphics, SPBG'08*, pp. 163–170. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2008. doi: 10.2312/VG/VG-PBG08/163-170
- [3] D. Bonetta. GraalVM: metaprogramming inside a polyglot system (invited talk). In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Meta-Programming Techniques and Reflection*, pp. 3–4, 2018. doi: 10.1145/3281074.3284935
- [4] A. Bria, G. Iannello, L. Onofri, and H. Peng. TeraFly: real-time three-dimensional visualization and annotation of terabytes of multidimensional volumetric images. *Nature Methods*, 13:192–194, 2016. doi: 10.1038/nmeth.3767
- [5] F. d. Chaumont, S. Dallongeville, N. Chenouard, N. Herv, S. Pop, T. Provoost, V. Meas-Yedid, P. Pankajakshan, T. Lecomte, Y. L. Montagner, T. Lagache, A. Dufour, and J.-C. Olivo-Marin. Icy: an open bioimage informatics platform for extended reproducible research. *Nature Methods*, 9:690, 2012. doi: 10.1038/nmeth.2075
- [6] S. Daetwyler, U. Günther, C. D. Modes, K. Harrington, and J. Huisken. Multi-sample spim image acquisition, processing and analysis of vascular growth in zebrafish. *Development*, 146(6):dev173757, 2019.
- [7] W. Denk and H. Horstmann. Serial Block-Face Scanning Electron Microscopy to Reconstruct Three-Dimensional Tissue Nanostructure. *PLoS Biology*, 2:e329, 2004. doi: 10.1371/journal.pbio.0020329
- [8] J. Gold, A. Wang, and K. Harrington. Feedback control of evolving swarms. In *Artificial Life Conference Proceedings 14*, pp. 884–891. MIT Press, 2014. doi: 10.7551/978-0-262-32621-6-ch145
- [9] S. Grottel, M. Krone, C. Müller, G. Reina, and T. Ertl. MegaMol - A Prototyping Framework for Particle-Based Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 21:201–214, 2014. doi: 10.1109/tvcg.2014.2350479
- [10] M. D. Hanwell, K. M. Martin, A. Chaudhary, and L. S. Avila. The Visualization Toolkit (VTK): Rewriting the rendering code for modern graphics cards. *SoftwareX*, 1:1–4, 2015. doi: 10.1016/j.softx.2015.04.001
- [11] K. Harrington and L. Magbunduku. Competitive dynamics in eco-evolutionary genetically-regulated swarms. In *Proceedings of the European Conference on Artificial Life 14*, vol. 14, pp. 190–197. MIT Press, 2017. doi: 10.7551/eca.a.034
- [12] K. I. S. Harrington and T. Stiles. kephale/brevis 0.10.4, July 2017. doi: 10.5281/zenodo.822902
- [13] J. Huisken, J. Swoger, F. D. Bene, J. Wittbrodt, and E. H. Stelzer. Optical sectioning deep inside live embryos by selective plane illumination microscopy. *Science*, 305:1007–1009, 2004. doi: 10.1126/science.1100035
- [14] O. Igouchkine, Y. Zhang, and K.-L. Ma. Multi-Material Volume Rendering with a Physically-Based Surface Reflection Model. *IEEE Transactions on Visualization and Computer Graphics*, 24:3147–3159, 2017. doi: 10.1109/tvcg.2017.2784830
- [15] D. Jonsson, P. Steneteg, E. Sunden, R. Englund, S. Kottravall, M. Falk, A. Ynnerman, I. Hotz, and T. Ropinski. Inviwo - A Visualization System with Usage Abstraction Levels. *IEEE Transactions on Visualization and Computer Graphics*, PP:1–1, 2019. doi: 10.1109/tvcg.2019.2920639
- [16] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness. *The International Journal of Aviation Psychology*, 3:203–220, 1993. doi: 10.1207/s15327108ijap0303\_3
- [17] T. Kroes, F. H. Post, and C. P. Botha. Exposure Render: An Interactive Photo-Realistic Volume Rendering Framework. *PLoS ONE*, 7:e38586, 2012. doi: 10.1371/journal.pone.0038586
- [18] V. Kmpe, E. Sintorn, and U. Assarsson. High resolution sparse voxel DAGs. *ACM Transactions on Graphics*, 32:1–8, 2013. doi: 10.1145/2461912.2462024
- [19] S. Laine and T. Karras. Effective Sparse Voxel Octrees - Analysis, Extensions and Implementation.
- [20] E. LaMar, B. Hamann, and K. I. Joy. Multiresolution techniques for interactive texture-based volume visualization. In *Proceedings of the Conference on Visualization '99: Celebrating Ten Years, VIS '99*, pp. 355–361. IEEE Computer Society Press, Los Alamitos, CA, USA, 1999.
- [21] O. Nalbach, E. Arabadzhiyska, D. Mehta, H. Seidel, and T. Ritschel. Deep Shading: Convolutional Neural Networks for Screen Space Shading. pp. 65–78, 2017. doi: 10.1111/cgf.13225
- [22] H. Peng, A. Bria, Z. Zhou, G. Iannello, and F. Long. Extensible visualization and analysis for multidimensional images using Vaa3D. *Nature Protocols*, 9:193–208, 2014. doi: 10.1038/nprot.2014.011
- [23] S. Pidhorskyi, M. Morehead, Q. Jones, G. Spirou, and G. Doretto. syGlass: Interactive Exploration of Multidimensional Images Using Virtual Reality Head-mounted Displays. 2018.
- [24] T. Pietzsch, S. Preibisch, P. Tomancak, and S. Saalfeld. ImgLib2: generic image processing in Java. *Bioinformatics*, 28:3009–3011, 2012. doi: 10.1093/bioinformatics/bts543
- [25] T. Pietzsch, S. Saalfeld, S. Preibisch, and P. Tomancak. Big-DataViewer: visualization and processing for large image data sets. *Nature Methods*, 12:481–483, 2015. doi: 10.1038/nmeth.3392
- [26] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '87*, pp. 25–34. ACM, New York, NY, USA, 1987. doi: 10.1145/37401.37406
- [27] L. A. Royer, W. C. Lemon, R. K. Chhetri, Y. Wan, M. Coleman, E. W. Myers, and P. J. Keller. Adaptive light-sheet microscopy for long-term, high-resolution imaging in living organisms. *Nature Biotechnology*, 34:1267–1278, 2016. doi: 10.1038/nbt.3708
- [28] L. A. Royer, M. Weigert, U. Günther, N. Maghelli, F. Jug, I. F. Sbalzarini, and E. W. Myers. ClearVolume: open-source live 3D visualization for light-sheet microscopy. *Nature Methods*, 12:480–481, 2015. doi: 10.1038/nmeth.3372
- [29] C. T. Rueden, J. Schindelin, M. C. Hiner, B. E. DeZonia, A. E. Walter, E. T. Arena, and K. W. Eliceiri. ImageJ2: ImageJ for the next generation of scientific image data. *BMC Bioinformatics*, 18:529, 2017. doi: 10.1186/s12859-017-1934-z
- [30] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J.-Y. Tinevez, D. J. White, V. Hartenstein, K. Eliceiri, P. Tomancak, and A. Cardona. Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9:676, 2012. doi: 10.1038/nmeth.2019
- [31] B. Schmid, J. Schindelin, A. Cardona, M. Longair, and M. Heisenberg. A high-level 3D visualization API for Java and ImageJ. *BMC Bioinformatics*, 11:274, 2010. doi: 10.1186/1471-2105-11-274
- [32] M. Slater and M. V. Sanchez-Vives. Enhancing our lives with immersive virtual reality. *Frontiers in Robotics and AI*, 3:74, 2016. doi: 10.3389/frobt.2016.00074
- [33] W. Usher, P. Klacansky, F. Federer, P.-T. Bremer, A. Knoll, J. Yarch, A. Angelucci, and V. Pascucci. A Virtual Reality Visualization Tool for Neuron Tracing. *IEEE Transactions on Visualization and Computer Graphics*, 24:994–1003, 2017. doi: 10.1109/tvcg.2017.2744079
- [34] J. Wu, J. B. Tenenbaum, and P. Kohli. Neural Scene De-rendering. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2017. doi: 10.1109/cvpr.2017.744