

Semantic Screen-Space Occlusion for Multiscale Molecular Visualization

Thomas Bernhard Koch¹  and David Kouřil¹  and Tobias Klein¹  and Peter Mindek¹  and Ivan Viola¹ 

¹TU Wien

Abstract

Visual clutter is a major problem in large biological data visualization. It is often addressed through the means of level of detail schemes coupled with an appropriate coloring of the visualized structures. Ambient occlusion and shadows are often used to improve the depth perception. However, when used excessively, these techniques are sources of visual clutter themselves. In this paper we present a new approach to screen-space illumination algorithms suitable for use in illustrative visualization. The illumination effect can be controlled so that desired levels of semantic scene organization cast shadows while other remain flat. This way the illumination design can be parameterized to keep visual clutter, originating from illumination, to a minimum, while also guiding the user in a multiscale model exploration. We achieve this by selectively applying occlusion shading based on the inherent semantics of the visualized hierarchically-organized data. The technique is in principle generally applicable to any hierarchically organized 3D scene and has been demonstrated on an exemplary scene from integrative structural biology.

CCS Concepts

• **Human-centered computing** → **Scientific visualization**;

1. Introduction

In biology, illustrations are often used to communicate knowledge and novel discoveries. For this purpose, illustrators often use various shading methods to highlight particular phenomena or to visually discriminate individual objects from each other. The illumination effects and shadows are often used for highlighting a particular object of interest. This way, intuitive visual guidance is traded for a physically plausible illumination.

Computer graphics can be utilized to create informative illustrations and animations describing biological processes on atomic level. However, for interactive visualization, extremely large number of graphics primitives, such as spheres representing individual atoms, require the usage of screen-space shading methods like screen-space ambient occlusion to create depth cues. The problem with these methods is that they act on image, depth, and normal buffer pixels and as such are not aware of the scene hierarchy, inherently present in models of the biological structures. Shading effects, such as ambient occlusion or shadows, are indiscriminately and equally applied to all graphical elements without any notion of objects or user's focus. Additionally, complex scenes consisting of thousands of molecules often suffer from excessive visual clutter introduced by the screen-space shading methods, as the shading effects are applied to all individual molecules resulting in unnecessarily detailed shading.

To circumvent this problem, we propose a method to extend

screen-space shading methods, such as screen-space ambient occlusion (SSAO) [BSD08, SA07] so that they are aware of the hierarchy of the displayed mesoscale models. In this way, higher-level structures can be lit to cast shadows for illustrative purposes while their own detail remains flat shaded. Using our method, illustrators can efficiently highlight the structure of integrative biological models, which are inherently multiscale, while maintaining the real-time performance of the visualization.

The contribution of this paper is a new approach extending screen-space shading methods to take the hierarchy of the displayed multiscale models into account with the goal to highlight certain features and to reduce visual clutter. The new technique is demonstrated by implementing a hierarchy-aware extension of screen-space directional occlusion (SSDO) [RGS09] and hierarchy-aware screen-space ray traced shadows into the Marion visualization framework [MKS*17], an interactive rendering system to visualize biological data.

2. Related Work

Ambient Occlusion is a well-known technique and is used to enhance the perception of the overall 3D shape of complex cluttered scenes, where standard techniques with direct lighting often fail. Specifically in molecular visualization, the work of Tarini et al. [TCM06] is one of the first, which applies a more sophisticated shading model with ambient occlusion on molecular scenes.

Grottel et. al. [GKSE12] present an object-space ambient occlusion method for large, dynamic particle-based data sets that reaches interactive frame rates to improve depth perception and reduce visual clutter. This is achieved by storing density information in a coarse resolution approximation of the data set. Wahle et. al. [WW15] use SSAO with varying sphere diameters to highlight different features of atomic structures.

Many screen-space ambient occlusion techniques [BSD08, SA07, MOBH11] have in common that first an ambient occlusion factor is determined which is then multiplied by the calculated illumination. This separation leads to the problem that the lighting is ignored when computing AO, thus resulting in equally darkened areas. The SSDO technique by Ritschel et al. [RGS09] removes this decoupling of the occlusion and the illumination calculation. On top of that, they describe a method to add one indirect bounce of light in screen space. Occluders are approximated in screen-space. For this, samples points have to be generated. Sampling techniques generating uniformly distributed samples are favorable. Compared to simpler techniques like random sampling, Halton sequences [WLH97] give a very uniformly distributed sampling pattern. Screen-space ambient occlusion techniques typically produce noise artifacts due to an insufficient number of sample points. In order to reduce noise, smoothing filters are usually used. Instead of a simple Gaussian low-pass filter, which would also blur edges, bilateral filtering [TM98] is often used.

Technological advances in biology and proteomics facilitate the generation of large-scale models [JAAA*15, KAK*18] of whole viruses or bacteria in atomic detail. There are several tools and platforms that support the visualization of molecular multiscale data. VMD [HDS96] is a comprehensive tool designed to display and analyze large biomolecular systems, especially the results of molecular dynamics simulations. The prototyping framework Megamol [GKM*15] is a visualization tool that focuses on fast rendering of large particle systems. The tool cellVIEW [MAPV15] is an illustrative approach for large scale rendering of biomolecular datasets integrated in the game engine Unity. The semantic screen-space occlusion approach of this work is implemented in a rendering system called Marion [MKS*17]. Marion supports biologists in creating multiscale visualizations of biological data for efficient public dissemination. The system derives a visual representation of the data, which is inspired by traditional scientific illustrations, and renders it in real-time. In Marion, dynamic multiscale color mapping [WMW*16] is used in order to reduce visual clutter. The idea is to reuse color space and redistribute it to currently visible elements. Chroma and luminance are used to encode structural and hierarchical properties and to reduce visual clutter. Furthermore, a level of detail scheme is used for the displayed geometry in order to maintain clarity when changing the zoom level.

3. Semantic Shading Approach

Biological multiscale models are often hierarchically organized into several compartments, typically separated by lipid membranes. For instance, the HIV virion consists of three compartments. The innermost compartment is enclosed by a capsid containing the RNA genome. The capsid itself is enclosed by another compartment containing several proteins, which are encapsulated by a lipid

membrane. Surrounding blood plasma can be considered as the outermost compartment. Previous shading approaches do not take this hierarchical information into account when rendering biological scenes. We propose a shading approach that considers the semantic structuring of the data and thus improving algorithms' flexibility in conveying the desired shape.

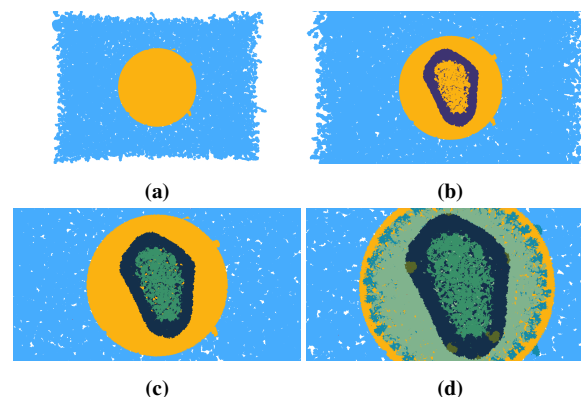


Figure 1: Regions buffer at increasing scale levels i.e. zooming in. Each color represents a semantic region. With large distance from the camera, the scene is subdivided into two regions (a), representing the blood plasma compartment in blue and the whole HIV virion in orange. When zooming in, the virion is further subdivided (b) resulting in a separate region for the capsid, depicted in dark blue. This continues in the same fashion (c, d) until also the innermost compartment is represented with a region.

Based on the Labels on Levels approach by Kouřil et. al. [KvK*18], the scene is subdivided into different semantic regions depending on the scene objects' hierarchy. The subdivision is controlled using user-defined depth thresholds that divide the depth of the camera into several segments. The resulting regions buffer is shown in Figure 1. Differently to Labels on Levels, the subdivision does not depend on the distance of individual objects to the camera but solely on the distance of the camera to the center of the scene. This way artifacts are avoided as explained later. This behavior is shown in Figure 2.




Figure 2: The difference in semantic regions subdivision: (a) The subdivision into semantic regions in Labels on Levels depends on the distance of objects to the camera. (b) Here, the subdivision depends on the distance of the camera to the center of the scene. The subdivision in (a) causes shadowing artifacts as semantic regions are gradually subdivided when the distance to the camera changes i.e. when zooming and, for example, shadows are falsely cast from the blue onto the pink region.

The regions buffer is then used during the occlusion calculation so that there is just ambient occlusion between different semantic regions. Occluders are approximated in screen-space, as explained by Ritschel et. al. [RGS09]. First, sample points for each pixel at 3D position x within a normal oriented hemisphere are chosen. The sample points are back-projected into screen-space. Now screen-space coordinates are known. The sample points are then projected onto the surface by reading the world-space position from a position buffer using the obtained screen-space coordinates. Also, the semantic region at the position of the sample point is fetched from the regions buffer. A sample point is then classified as an occluder, if it moved closer to the viewer by the projection onto the surface and if it belongs to a different semantic region than the current pixel. This way, there is just ambient occlusion between objects of different regions (see Figure 3).




Figure 3: The condition when a sample point is classified as an occluder is extended, compared to Ritschel et. al. [RGS09], taking also the semantic regions the sample points belong to into account. The illumination of x is calculated from the directions of sample point B and C (green arrows) because they belong to the same region (represented by the color of the dot on the surface) as x . The light is occluded from the direction of sample point A because it belongs to a different region.

4. Algorithmic Pipeline

To demonstrate our approach, we implemented a hierarchy-aware extension of SSDO and screen-space ray traced shadows into Marion. Figure 4 gives an overview of the visualization pipeline. The pipeline consists of multiple passes that each render their output into a texture which is similar to G-buffer [ST90]. In the following, the notation \otimes is used to refer to the output images in Figure 4.

The scene containing HIV in blood plasma is generated using cellPACK [JAAA*15]. In the first step, the scene is rendered using the approach of Le Muzic et. al. [MAPV15]. This step outputs a G-buffer containing world space coordinates $\otimes 1$ and normals $\otimes 2$. In the *label regions* step, the scene is subdivided into semantic regions $\otimes 4$ using the approach of Koufil et. al. [KvK*18] depending on the depth segment the camera currently is in. Also, the subdivision into regions of the subsequent depth segment is stored in a buffer $\otimes 5$. In the next step, the occlusion factor is calculated and stored in a buffer $\otimes 6$ by taking the semantic regions $\otimes 4$ and $\otimes 5$ from

the previous step into account (see Section 3). $\otimes 4$ and $\otimes 5$ are linearly interpolated depending on the camera position between two depth-thresholds. In order to capture occlusion between atoms but also between larger, more distant structures like molecules, the occlusion of a point is determined using a small and a large radius for the hemisphere. The radii are user-defined and scene dependent. To get the total occlusion of a point, both the occlusion factor caused by near and by far occluders, are multiplied. The occluders are determined like in SSDO. The hemisphere is sampled with samples generated by Halton sequences which are mapped onto a hemisphere using *Malley's method* [PJH16]. In this step also the light bounce $\otimes 7$ is calculated according to Ritschel et. al. [RGS09]. To produce color bleeding between adjacent atoms, occluders are searched along each sampling direction using a fixed step size. Both outputs, the occlusion factor $\otimes 6$ and the buffer containing the light bounce $\otimes 7$, are blurred using an edge preserving bilateral filter [TM98] to reduce noise. The bilateral blur is separated to reduce the computational complexity [PvV05].

Before combining the results of the previous steps in the last pass, shadows are rendered into a buffer $\otimes 9$ which is then also blurred using the bilateral blur to get softer shadows. To determine if a point is in shadow, screen-space ray tracing, which is often used for screen-space reflections [Val14, Sai16], is used. The depth buffer is marched in the direction of the key light source to determine if a point is in shadow. If an intersection is found, the point is in shadow. In Marion, a three point light setup is used, utilizing a key, a fill and a back light. Here, just the key light source casts shadows.

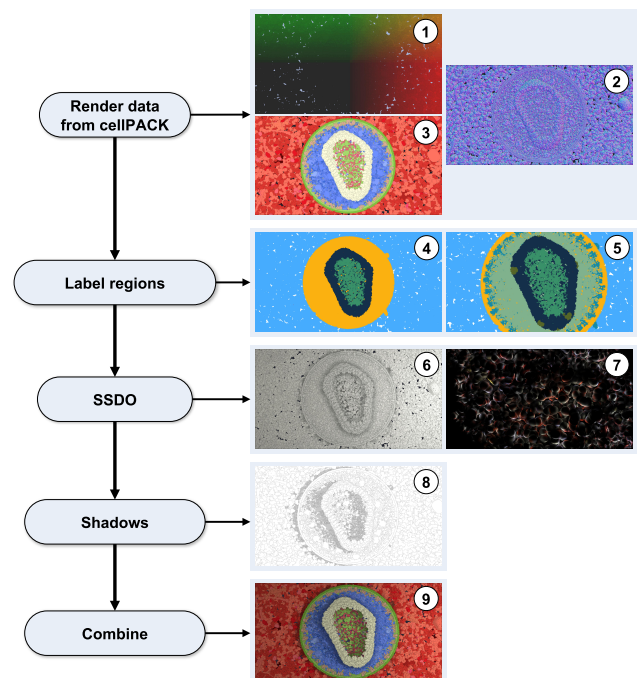


Figure 4: Overview of the implemented hierarchy-aware SSDO pipeline. The output of each step is given on the right side.

5. Results

In this section results are presented and discussed, showing the HIV scene used throughout the paper. Compared to previous work [GKSE12, WW15], our work focuses on camera dependent highlighting of whole semantic regions by selectively applying ambient occlusion, based on information about the structure of the visualized molecular model. Visual clutter becomes apparent when viewing the dataset in its entirety and is caused by ambient occlusion between a large number of small objects on the screen. This has been improved by incorporating information about the semantics of the molecular model into the occlusion calculation thus restricting ambient occlusion just to certain areas. When zooming into the scene, the structure of the molecular model is gradually emphasized by slowly appearing soft shadows caused by ambient occlusion. Figure 5 shows this behavior. It can be observed that the semantic ambient occlusion approach supports the level of detail system that is used for the geometry in hiding unnecessary detail. The smaller the objects causing ambient occlusion become on the screen, the more visual clutter is introduced and the image starts to look noisy. This is demonstrated in Figure 6. With semantic ambient occlusion the image looks much calmer because there is just ambient occlusion around larger structures.

In order to assess the amount of visual clutter with and without our semantic shading approach, we have used an image-based clutter metric presented by Rosenholtz et. al. [RLN07] where color, orientation and luminance contrast are taken into account at multiple sizes of the image. For our measurements we have used the publicly available MATLAB implementation. The amount of visual clutter is measured at four different zoom levels of the HIV scene rendered with Marion used throughout this document. We have measured the visual clutter of the scene with screen-space ambient occlusion without taking semantics into account and with our implemented semantic SSDO approach with and without screen-space ray traced shadows. The measured results are shown in Figure 5. It can be seen that the visual clutter steadily increases when zooming in. Our semantic shading approach reduced the visual clutter, especially when areas consisting of a large number of objects of similar color, like the surrounding blood plasma, were visible (first three zoom levels in Figure 5). The last and further zoom levels show a similar amount of visual clutter because at these zoom levels ambient occlusion between individual objects, even if they belong to the same region, is rendered. Additional shadows did not help to reduce the visual clutter.

Based on our performance measurements of the implemented hierarchy-aware SSDO approach (Intel Core i5-6600K CPU and NVIDIA GeForce GTX 1070 GPU), we expect the impact on rendering speed of making a screen-space shading method hierarchy aware to be negligible. This is especially the case when current hardware is used.

6. Future Work and Conclusion

We presented an approach which allows to make existing screen-space shading methods aware of the hierarchy of the visualized multiscale models. Especially interactive rendering systems, where large biological datasets are visualized resulting in scenes with

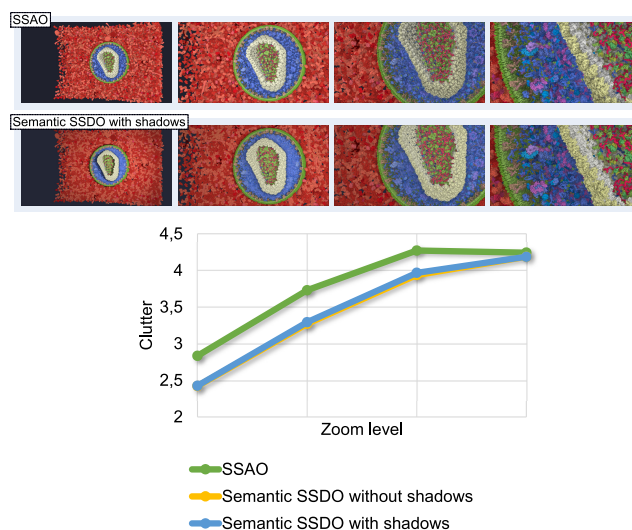


Figure 5: Visual clutter measured [RLN07] at four zoom levels. The top row shows zooming into the HIV scene with standard SSAO without taking semantics into account while the second row is rendered with our semantic SSDO approach. It can be seen that soft shadows caused by ambient occlusion are slowly appearing gradually emphasizing the structure of the cell. First, shadows are just visible around and on the inside of the capsid (white) and around the outermost lipid membrane (green). When zooming further into the scene, soft shadows on the inside of the outermost membrane become visible, revealing further details. The shadows cast by the capsid and the lipid membrane are in turn faded out, preventing the scene from getting too dark at the atom level. Zooming in even further reveals soft shadows between molecules of the same region (third zoom level), before individual atoms are revealed.




Figure 6: Comparison of the HIV scene with (a) and without (b) semantic SSDO. With semantic SSDO the image looks calmer and less important areas like the surrounding blood plasma attract less attention. Without semantic SSDO the scene looks noisy and cluttered due to the small size of objects on the screen that cause ambient occlusion, distracting the user from the HIV virion.

thousands of molecules, can benefit from our method as they typically utilize screen-space shading methods.

We demonstrated our method by implementing hierarchy-aware SSDO and screen-space ray traced shadows into Marion, a state-of-the-art interactive molecular visualization system. As the algorithm

operates in screen space, it can be effortlessly integrated with other rendering frameworks that expose multiscale character of the scene. Our method could also be used to actively guide users through the scene by employing the illumination for highlighting the essential structures.

In future work, we would like to investigate what other traditional and well-known screen-space shading methods could be extended in order to take advantage of additional hierarchy information that is often provided with meso-scopic molecular models. Additionally, reusing sampled pixel neighborhood in future temporal steps, as well as progressive refinement schemes, could improve the performance and temporal coherency of the presented technique.

7. Acknowledgements

The authors would like to thank the anonymous reviewers for their helpful comments and suggestions on the paper.

References

- [BSD08] BAVOIL L., SAINZ M., DIMITROV R.: Image-space horizon-based ambient occlusion. In *ACM SIGGRAPH 2008 Talks* (New York, NY, USA, 2008), SIGGRAPH '08, ACM, pp. 22:1–22:1. URL: <http://doi.acm.org/10.1145/1401032.1401061>, doi:10.1145/1401032.1401061. 1, 2
- [GKM*15] GROTTTEL S., KRONE M., MÜLLER C., REINA G., ERTL T.: MegaMol – a prototyping framework for particle-based visualization. *IEEE Transactions on Visualization and Computer Graphics* 21, 2 (Feb 2015), 201–214. doi:10.1109/TVCG.2014.2350479. 2
- [GKSE12] GROTTTEL S., KRONE M., SCHARNOWSKI K., ERTL T.: Object-space ambient occlusion for molecular dynamics. In *2012 IEEE Pacific Visualization Symposium* (Feb 2012), pp. 209–216. doi:10.1109/PacificVis.2012.6183593. 2, 4
- [HDS96] HUMPHREY W., DALKE A., SCHULTEN K.: VMD: Visual molecular dynamics. *Journal of Molecular Graphics* 14, 1 (1996), 33 – 38. URL: <http://www.sciencedirect.com/science/article/pii/0263785596000185>, doi:https://doi.org/10.1016/0263-7855(96)00018-5. 2
- [JAAA*15] JOHNSON G. T., AUTIN L., AL-ALUSI M., GOODSSELL D. S., SANNER M. F., OLSON A. J.: cellPACK: A virtual mesoscope to model and visualize structural systems biology. *Nature methods* 12, 1 (2015), 85. 2, 3
- [KAK*18] KLEIN T., AUTIN L., KOZLÍKOVÁ B., GOODSSELL D., OLSON A., GRÖLLER M., VIOLA I.: Instant construction and visualization of crowded biological environments. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (1 2018), 862–872. 2
- [KvK*18] KOUŘIL D., ČMOLÍK L., KOZLÍKOVÁ B., WU H., JOHNSON G., GOODSSELL D. S., OLSON A., GRÖLLER M. E., VIOLA I.: Labels on levels: Labeling of multi-scale multi-instance and crowded 3d biological environments. *IEEE Transactions on Visualization and Computer Graphics* (2018), 1–1. doi:10.1109/TVCG.2018.2864491. 2, 3
- [MAPV15] MUZIC M. L., AUTIN L., PARULEK J., VIOLA I.: celVIEW: a tool for illustrative and multi-scale rendering of large biomolecular datasets. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (Sept. 2015), Bühler K., Linsen L., John N. W., (Eds.), EG Digital Library, The Eurographics Association, pp. 61–70. URL: https://www.cg.tuwien.ac.at/research/publications/2015/cellVIEW_2015/. 2, 3
- [MKS*17] MINDEK P., KOUŘIL D., SORGER J., TOLOUDIS D., LYONS B., JOHNSON G., GRÖLLER M. E., VIOLA I.: Visualization multi-pipeline for communicating biology. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017). URL: <https://www.cg.tuwien.ac.at/research/publications/2017/mindek-2017-marion/>. 1, 2
- [MOBH11] MCGUIRE M., OSMAN B., BUKOWSKI M., HENNESSY P.: The alchemy screen-space ambient obscurity algorithm. In *Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics* (New York, NY, USA, 2011), HPG '11, ACM, pp. 25–32. URL: <http://doi.acm.org/10.1145/2018323.2018327>, doi:10.1145/2018323.2018327. 2
- [PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016. 3
- [PvV05] PHAM T. Q., VAN VLIET L. J.: Separable bilateral filtering for fast video preprocessing. In *2005 IEEE International Conference on Multimedia and Expo* (July 2005), pp. 4 pp.–. doi:10.1109/ICME.2005.1521458. 3
- [RGS09] RITSCHHEL T., GROSCH T., SEIDEL H.-P.: Approximating dynamic global illumination in image space. In *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2009), I3D '09, ACM, pp. 75–82. URL: <http://doi.acm.org/10.1145/1507149.1507161>, doi:10.1145/1507149.1507161. 1, 2, 3
- [RLN07] ROSENHOLTZ R., LI Y., NAKANO L.: Measuring visual clutter. *Journal of Vision* 7, 2 (2007), 17. URL: <http://dx.doi.org/10.1167/7.2.17>, arXiv:/data/journals/jov/932848/jov-7-2-17.pdf, doi:10.1167/7.2.17. 4
- [SA07] SHANMUGAM P., ARIKAN O.: Hardware accelerated ambient occlusion techniques on gpus. In *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2007), I3D '07, ACM, pp. 73–80. URL: <http://doi.acm.org/10.1145/1230100.1230113>, doi:10.1145/1230100.1230113. 1, 2
- [Sai16] SAIKIA S.: *Screen Space Reflections in Killing Floor 2*, 2016. <https://sakibsaikia.github.io/graphics/2016/12/25/Screen-Space-Reflection-in-Killing-Floor-2.html> (accessed May 27, 2018). 3
- [ST90] SAITO T., TAKAHASHI T.: Comprehensible rendering of 3-d shapes. *SIGGRAPH Comput. Graph.* 24, 4 (Sept. 1990), 197–206. URL: <http://doi.acm.org/10.1145/97880.97901>, doi:10.1145/97880.97901. 3
- [TCM06] TARINI M., CIGNONI P., MONTANI C.: Ambient occlusion and edge cueing for enhancing real time molecular visualization. *IEEE Transactions on Visualization and Computer Graphics* 12 (2006). 1
- [TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)* (Jan 1998), pp. 839–846. doi:10.1109/ICCV.1998.710815. 2, 3
- [Val14] VALIENT M.: *Reflections and volumetrics of Killzone Shadow Fall. Presentation at SIGGRAPH Advances in Real-Time Rendering in Games course*, 2014. http://advances.realtimerendering.com/s2014/valient/Valient_Siggraph14_Killzone.pptx (accessed May 27, 2018). 3
- [WLH97] WONG T.-T., LUK W.-S., HENG P.-A.: Sampling with hammersley and halton points. *J. Graph. Tools* 2, 2 (Nov. 1997), 9–24. URL: <http://dx.doi.org/10.1080/10867651.1997.10487471>, doi:10.1080/10867651.1997.10487471. 2
- [WMW*16] WALDIN N., MUZIC M. L., WALDNER M., GRÖLLER M. E., GOODSSELL D., AUTIN L., VIOLA I.: Chameleon dynamic color mapping for multi-scale structural biology models. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2016). URL: https://www.cg.tuwien.ac.at/research/publications/2016/Waldin_Nicholas_2016_Chameleon/. 2
- [WW15] WAHLE M., WRIGGERS W.: Multi-scale visualization of molecular architecture using real-time ambient occlusion in sculptor. *PLOS Computational Biology* 11, 10 (10 2015), 1–14. URL: <https://doi.org/10.1371/journal.pcbi.1004516>, doi:10.1371/journal.pcbi.1004516. 2, 4