# Visibility Optimization for Direct and Indirect Volume Rendering using Level Set Propagation – Supplemental Material –

Paul Himmler and Tobias Günther

## I. Extension of Viola et al. [1]

ORIGINALLY Viola et al. [1] neither handled surfaces nor shadows from a point light. Thus, we extended their method to enable a fair comparison. In Fig. 1, We show four data sets rendered with maximum importance projection (MImP) including surfaces with and without an additional MImP pass from the light view. MImP, as proposed in [1], is based on the idea of rendering only the most important structure on a ray fully opaque and removing everything in front and behind it. To preserve the volumetric look of the renderings, we adapt the algorithm to include structures behind the most important structure, as well. Similarly to the original algorithm, we start by raymarching through the volume collecting importance values until we find the maximum. Instead of stopping and rendering that value, we start a classical transmittance based direct volume rendering from that maximum point to the end of the ray. Additionally, we employ a threshold, allowing us to decide how important a structure has to be to be considered in the query for the maximum value.

While the algorithm in the original paper does not consider indirect volume rendering, e.g., surface rendering, the method by default shows similarities to cutout views. Using our threshold, it is straightforward to extend the method to surfaces, by cutting parts of the surface where our threshold for the maximum importance is exceeded. Results on the left side of Fig. 1 demonstrate the combination of direct and indirect volume rendering with maximum importance projection.

A disadvantage of the MImP approach in comparison to more recent methods is the lack of lighting. Enabling our point light by default resulted in dark images, since all important objects were usually in shadow. For better comparison, we extend Viola et al. [1] to lighting and shadows. Therefore, to account for binary visibility from a point light, we compute the MImP from the light view and query for every point $\mathbf{x}$ on our view ray, if the most important structure on the light ray is blocking the light from reaching $\mathbf{x}$. In Fig. 1 on the right, we show the final result, where MImP is combined with surfaces and shadows from a point light source.

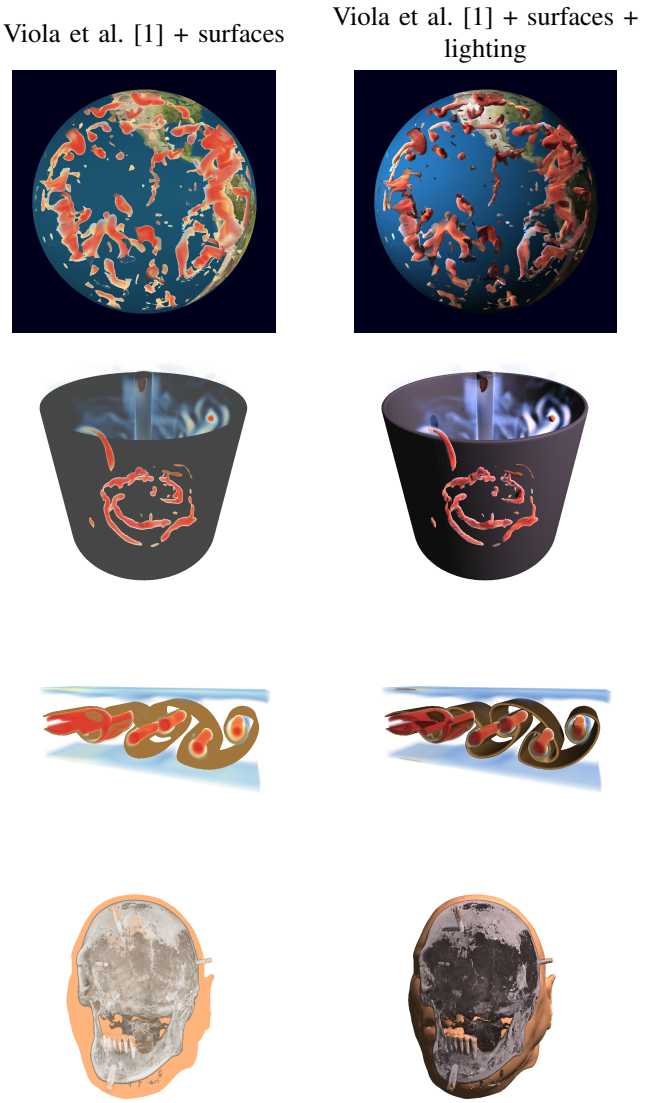


Fig. 1. Here, we show extensions of the maximum importance projection, originally introduced by Viola et al. [1]. On the left, we include surfaces, and on the right, we add shadows from a point light source.

Paul Himmler and Tobias Günther are with the Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany.

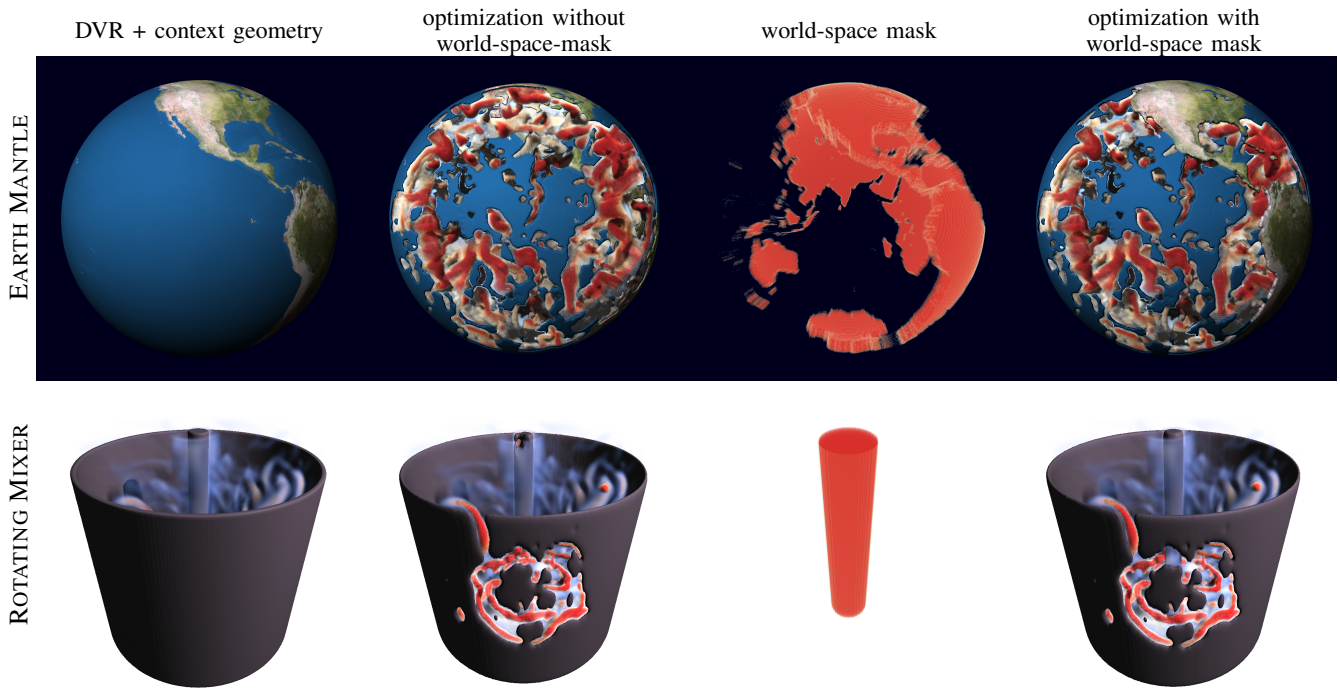


Fig. 2. We show results with world-space masks enabled and disabled for both the EARTH MANTLE and ROTATING MIXER data set. The used world-space mask is shown in the third column.

## II. WORLD-SPACE MASKS

In Fig. 2, we visualize the world-space masks for the two data sets, in which they have been utilized. From left to right, the columns show the unoptimized context geometry, the result without the world-space mask, the world-space mask itself and lastly the result of the optimization with world-space mask.

## III. INDIVIDUAL COMPONENTS

The main paper contained visibility optimizations for our four data sets. Here, in Fig. 3, we show the individual components that were used in the optimization, i.e., a standard direct volume rendering (DVR) of the scalar field, a depiction of the importance field, the DVR together with the complete context geometry, the result of our method, and for explanatory purposes our result from a different viewpoint. In the EARTH MANTLE, the context geometry, i.e., the Earth's surface, occludes the entire volume. The interesting structures are the cold slabs, sinking into the crust, which become visible once the occluding surface is removed. The ROTATING MIXER contains a numerical simulation of a liquid in a cylindrical container that is stirred into motion by three rotating paddles. By preserving the central rod as context, the vortices can be seen to be placed around the rod. The VORTEX STREET contains a von-Kármán vortex street, with vortices being separated by streak surfaces. The streak surfaces would hide large parts of the vortices, as they tend to fold around them. Lastly, the VISIBLE HUMAN shows the skin as context around the skull of a male. After rotating the view, the placement of the skull inside the context geometry can be seen well, including for example, the distance between the cheek skin and the jaw bone.

## IV. FLOATING POINT PRECISION

In Fig. 4, we assess the impact on the visual quality for 16-bit and 32-bit floating point precision. We employ the ꟻLIP metric [2], which shows a negligible difference across all datasets. We compare both setups for three of the four data sets of the main paper; the VISIBLE HUMAN data set did not fit into memory when using 32-bit floating point precision. Since both the energy we optimize for (e.g. for the transmittance), as well as the AdEMAMix [3] optimizer depend on floating point precision, there are minor numerical differences in the resulting energy residual. Nevertheless, in our experiments the results are visually indistinguishable.

## V. USER STUDY DATA

Figs. 5 and 6 display all images that have been shown one-by-one in randomized order to the users during the user study. Each image contains two marked points and participants were asked to state whether the shown image conveys whether one specific marked point is closer to the camera than the other. The responses were collected on a Likert scale from 1 (strongly disagree) to 5 (strongly agree). The mean and standard deviation are listed below for each image.

Further, we conducted a statistical analysis of the study results. Table I reports the mean Likert rating of each method as well as the corresponding rank for each user. The latter is used to demonstrate a significant difference between the methods. For this, we employ a Friedman test with a significance level of 5% and the null hypotheses $H_0$, which states that there is no significant difference between the four methods. For our data, the resulting p-value is $\approx 2.6 \times 10^{-5}$, which is $\leq 0.05$, and thus we reject the null hypothesis and conclude that there is a significant difference between at least one of the methods compared to the others. To further show the

TABLE I
MEAN RATING OF EACH METHOD AS WELL AS THE CORRESPONDING RANKS IN OUR USER STUDY LISTED FOR EACH PARTICIPANT. THE RANKS ARE USED FOR THE FRIEDMAN TEST AND FOR THE WILCOXON SIGNED-RANK TEST IN TABLE II.

| User | Maximum importance projection [1] | | Extinction opt. [4] + cut away | | Extinction opt. [4] + transp. modulation | | Our method | |
|---|---|---|---|---|---|---|---|---|
| | Mean | Rank | Mean | Rank | Mean | Rank | Mean | Rank |
| 1 | 1.875 | 4 | 2.25 | 2 | 2.125 | 3 | **3.875** | **1** |
| 2 | 2.5 | 3 | 2.25 | 4 | 2.875 | 2 | **4.125** | **1** |
| 3 | 2.5 | 3 | 2 | 4 | 2.75 | 2 | **3.875** | **1** |
| 4 | 2.375 | 3 | 2 | 4 | 2.5 | 2 | **3.75** | **1** |
| 5 | 2.125 | 4 | 2.375 | 2 | 2.25 | 3 | **4.375** | **1** |
| 6 | 2.375 | 2.5 | 2 | 4 | 2.375 | 2.5 | **3.375** | **1** |
| 7 | 2.25 | 4 | 2.375 | 3 | 2.75 | 2 | **3.875** | **1** |
| 8 | 2.25 | 3 | 2 | 4 | 3 | 2 | **4.5** | **1** |
| 9 | 1.875 | 4 | 2.375 | 2.5 | 2.375 | 2.5 | **4.375** | **1** |
| 10 | 2.75 | 4 | 3.25 | 2 | 3 | 3 | **4.375** | **1** |
| 11 | 3.125 | 3 | 2.75 | 4 | 3.25 | 2 | **3.375** | **1** |

TABLE II
THE TEST STATISTIC VALUES FOR THE TWO-TAILED WILCOXON SIGNED-RANK TEST FOR EACH PAIR OF METHODS. FOR A SIGNIFICANCE LEVEL OF 5% AND $n = 11$ USERS THE CRITICAL VALUE IS 13. FOR TEST STATISTIC VALUES SMALLER THAN THIS CRITICAL VALUE WE HAVE TO REJECT $H_0$ AND ASSUME A SIGNIFICANT DIFFERENCE BETWEEN THE METHODS.

| | Maximum importance projection [1] | Extinction opt. [4] + cut away | Extinction opt. [4] + transp. modulation | Our method |
|---|---|---|---|---|
| Maximum importance projection [1] | - | 14.5 | 0 | 0 |
| Extinction opt. [4] + cut away | 14.5 | - | 5.5 | 0 |
| Extinction opt. [4] + transp. modulation | 0 | 5.5 | - | 0 |
| Our method | 0 | 0 | 0 | - |

statistical significance of that difference and get details about which method differs from which other, we also perform a Wilcoxon signed-rank test for each pair. Again, we place a significance level of 5% for the null hypotheses $H_0$, stating that there is no significant difference between the methods. For this significance level and our $n = 11$ users, the critical value to reject this hypothesis for a two-tailed test is 13, which can be looked up from the commonly-used Wilcoxon signed rank tables. Table II shows the value of the test statistics for each pair. We can again conclude that there is a statistically significant difference between our method and the others. As can be seen, $H_0$ cannot be rejected for the combination of maximum importance projection and the cut-away strategy, which is unsurprising given their similar operating principles.

## REFERENCES

[1] I. Viola, A. Kanitsar, and M. E. Gröller, "Importance-driven feature enhancement in volume visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 4, p. 408–418, 7 2005.

[2] P. Andersson, J. Nilsson, T. Akenine-Möller, M. Oskarsson, K. Åström, and M. D. Fairchild, "FLIP: A difference evaluator for alternating images," *Proc. ACM Computer Graphics and Interactive Techiques*, vol. 3, no. 2, Aug. 2020.

[3] M. Pagliardini, P. Ablin, and D. Grangier, "The AdEMAMix optimizer: Better, faster, older," 2024, https://arxiv.org/abs/2409.03137.

[4] P. Himmler and T. Günther, "Transmittance-based extinction and viewpoint optimization," *Computer Graphics Forum (EuroVis)*, vol. 43, no. 3, p. e15096, 2024.

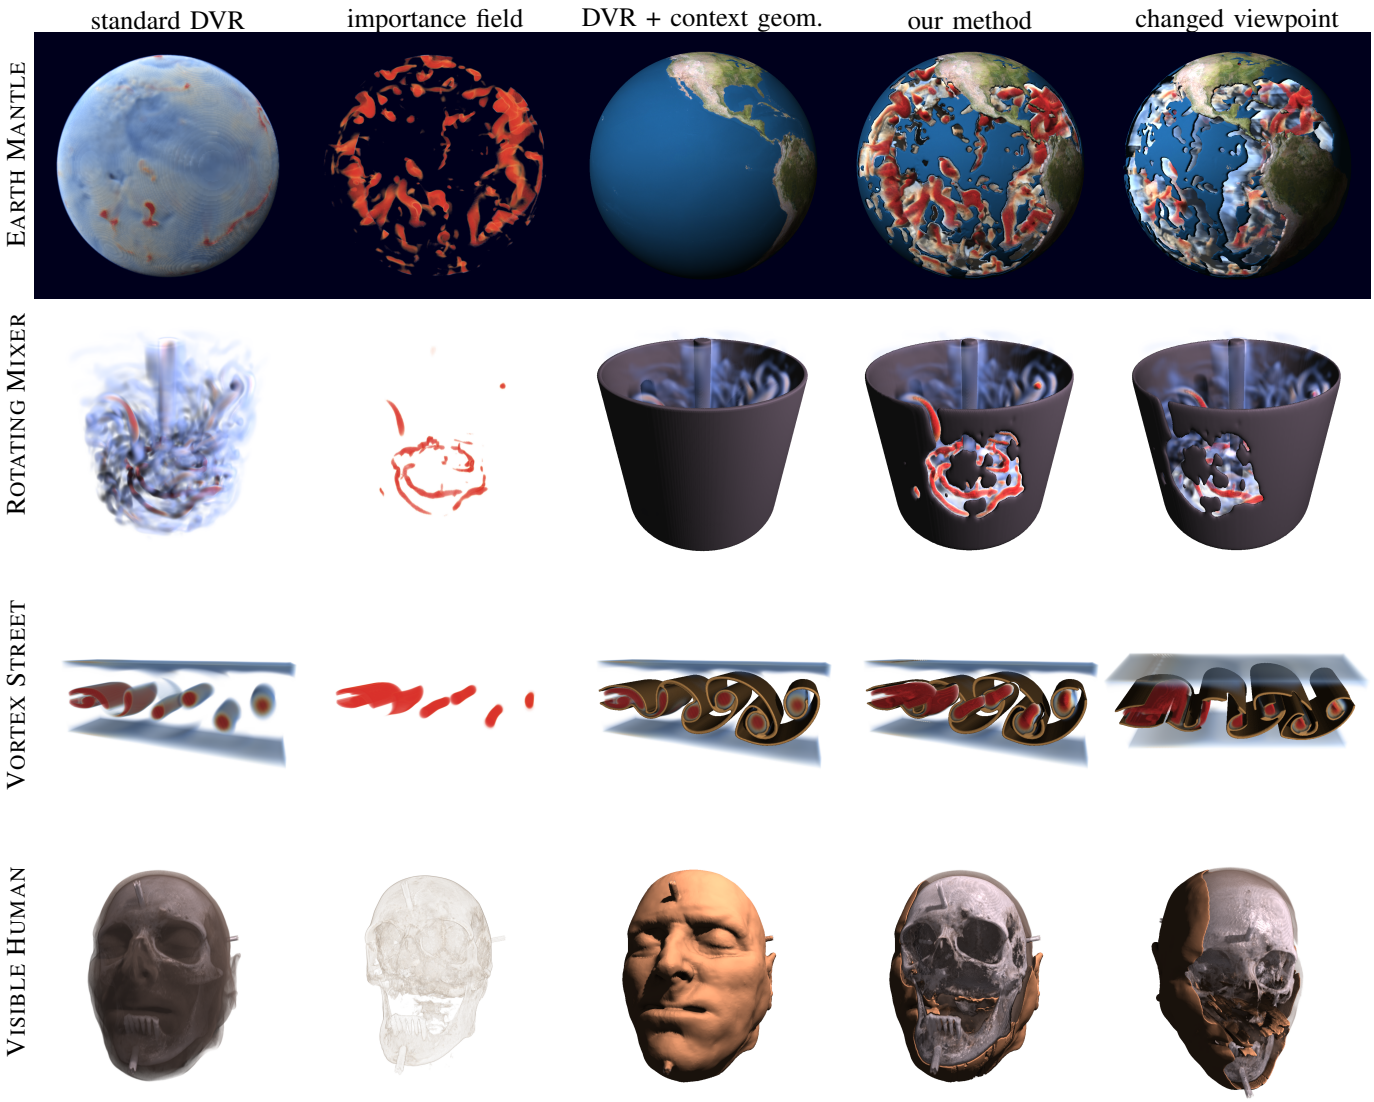


Fig. 3. Here, we show results of our visibility optimization in four scientific data sets. The first two columns depict the input to our optimization, i.e., a standard direct volume rendering, and an importance scalar field that identifies relevant structures. Our approach augments the direct volume rendering by adding context geometry, which is shown in the third column. Since the context geometry hides the relevant structures, we perform a visibility optimization that displaces the geometry using level set propagation in a signed distance field. The result is shown in the fourth column. The final column shows the same optimized scene from a different viewpoint.

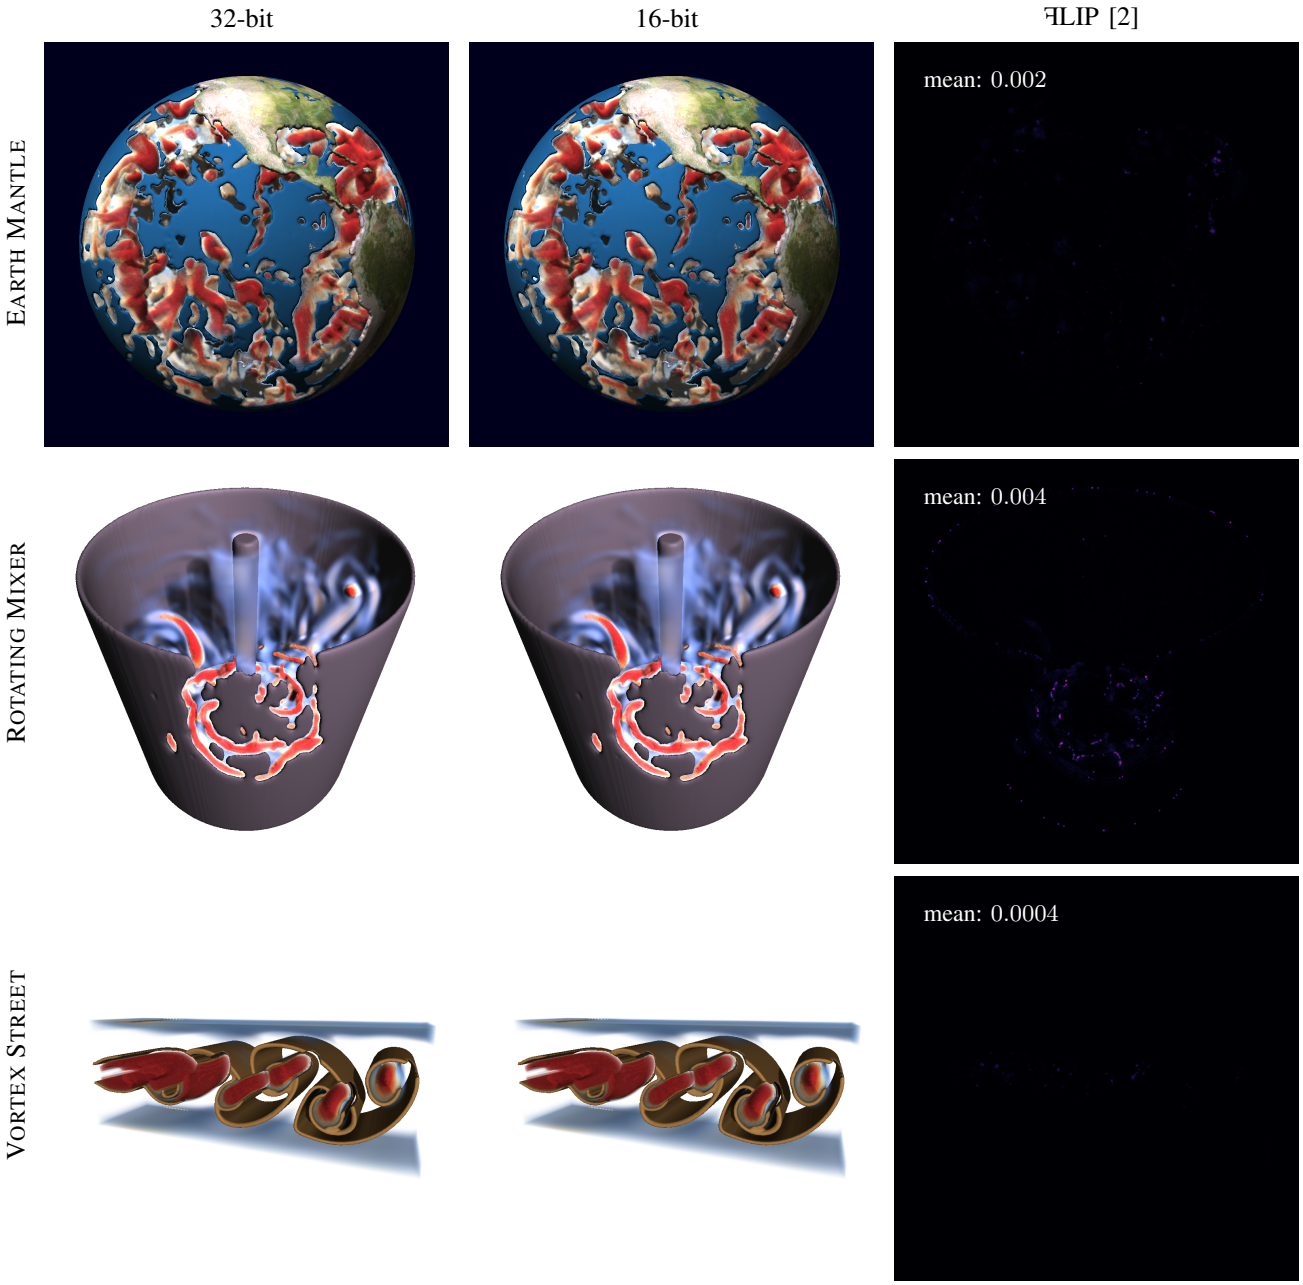


Fig. 4. Comparison of the visual quality for 32-bit and 16-bit floating point precision. Although there are minor numerical differences in the resulting energy residual due to the floating point precision, our experiments show that the visual results are visually indistinguishable. We use the ꟻLIP metric [2] to assess the differences across multiple datasets.

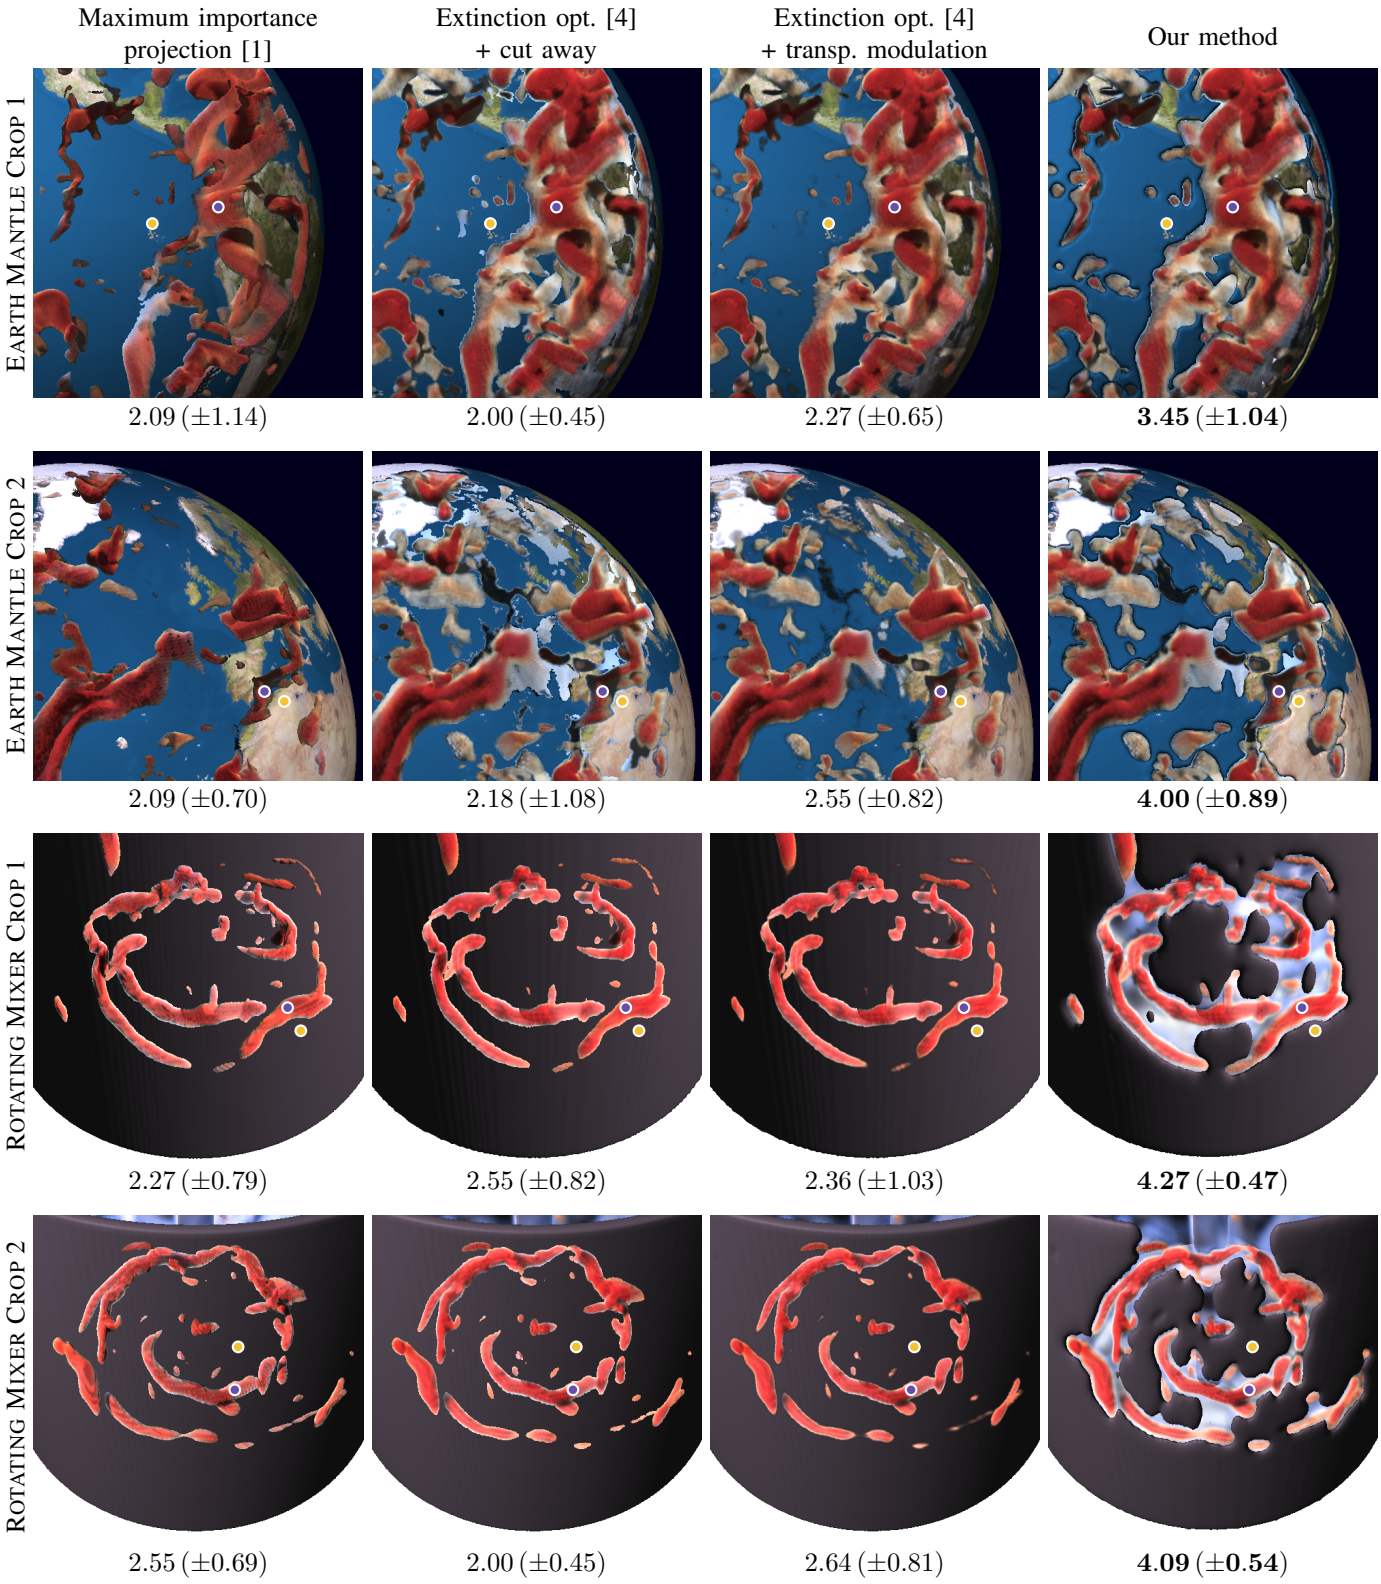


Fig. 5. Images (1/2) that have been used in the user study. For each image, the mean and standard deviation of the user responses are listed below.

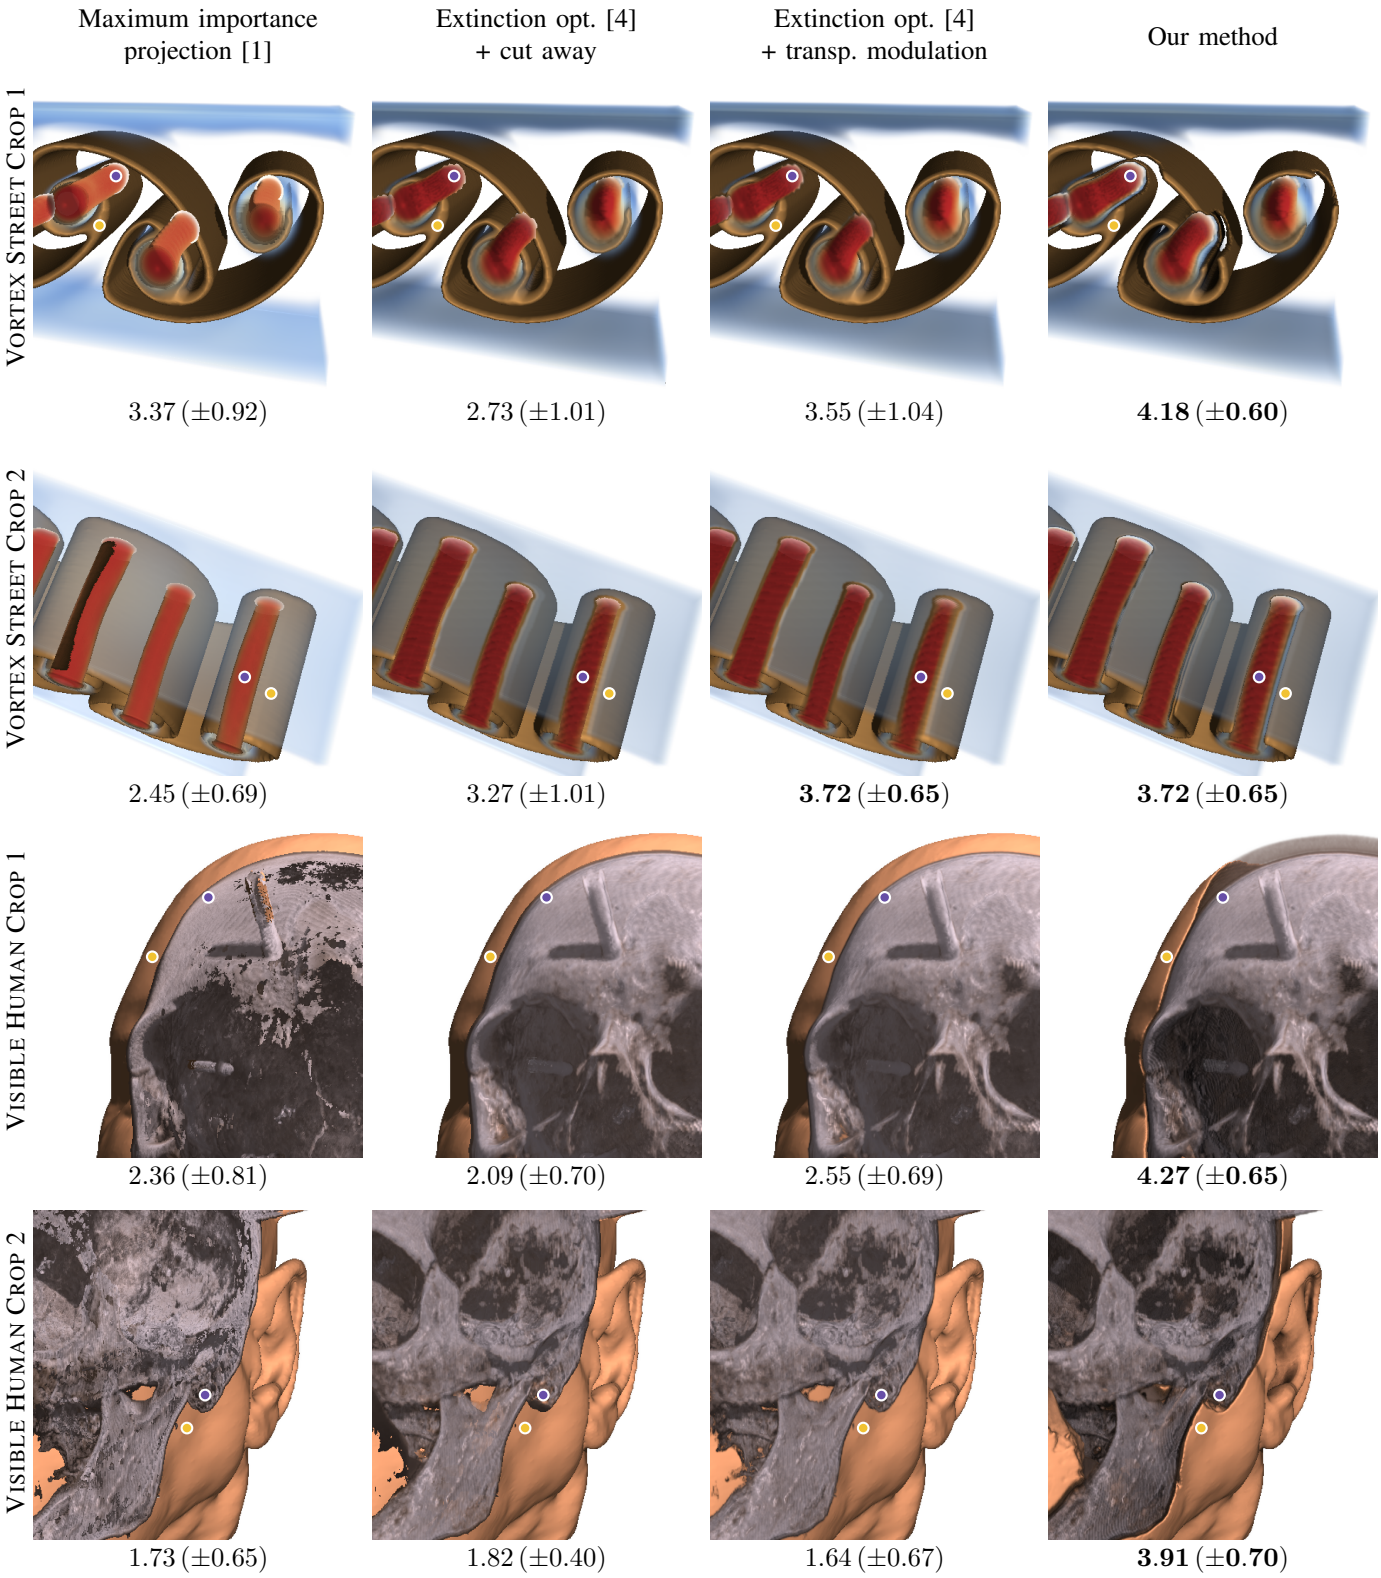


Fig. 6. Images (2/2) that have been used in the user study. For each image, the mean and standard deviation of the user responses are listed below.