





High Dynamic Range Point Clouds for Real-Time Relighting - Additional Material

Manuele Sabbadin¹ , Gianpaolo Palma¹ , Francesco Banterle¹ , Tamy Boubekeur², Paolo Cignoni¹ 

¹Visual Computing Lab - ISTI CNR, Pisa, Italy

²Telecom Paris, Institut Polytechnique de Paris & Adobe

1. Introduction

This document contains the supporting information of the paper "High Dynamic Range Point Cloud for Real-Time Relighting". In particular, Figures 1, 2 and 3 show the input data of tested scenes with an equirectangular view of the point cloud and the input HDR photo. Figures, from 4 to 15, show the rendering obtained with the expanded point cloud of scenes with HDR ground truth (SPONZA, SIBENIK, FIREROOM, ATRIUM, BUILDING and KITCHEN) with the relative the probability maps to detect differences from the ground truth rendering computed with HDR-VDP-2.2. Below each figure, there is a table with the error measures of each rendering from the ground truth. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity Index (SSIM). Figures 18 and 19 show the time and memory performance comparison of the three versions of the PBGI algorithm (Classic, MIP-PBGI, and X-PBGI) by changing the viewport size and the microbuffer size. For (SPONZA and TOYROOM), we used two viewpoints (see Figure 17): VIEW1 with a detail of the object that gets all the viewport; VIEW2 the second one where the entire object is visible in the viewport. Figure 20 shows the renderings of three different scenes, obtained with the three PBGI algorithms. For MIP-PBGI and X-PBGI, the PSNR value obtained from the comparison with the classic PBGI is reported. Figures, from 21 to 26, show the comparison of the X-PBGI rendering with the ground truth obtained with a path tracing, the classical environment mapping and two different versions of the Voxel Cone Tracing (VCT) algorithm (16 cones plus a specular cone for VCT16 and 1024 cones for VCT1024). Each rendering shows the relative error map from the ground truth using CIE LAB color space. Each figure shows the results for a pure diffuse BRDF and a GGX BRDF changing the roughness parameters.

References

- [BCD*13] BANTERLE F., CALLIERI M., DELLEPIANE M., CORSINI M., PELLACINI F., SCOPIGNO R.: Envydepth: An interface for recovering local natural illumination from environment maps. *Computer Graphics Forum* 32, 7 (October 2013), 411–420. 3

Panoramic view point cloud

HDR image



Figure 1: Input dataset with an equirectangular view of the point cloud and the input HDR image used for the expansion. These dataset have a ground truth HDR color for each point of the cloud. The point clouds are obtained by a Monte Carlo point sampling of three synthetic scenes after baking of the HDR diffuse color response using path tracing.



Figure 2: Input dataset with an equirectangular view of the point cloud and the input HDR image used for the expansion. These dataset have a ground truth HDR color for each point of the cloud. The point clouds were reconstructed from a single HDR panoramic image with a user-assisted method [BCD* 13].




Figure 3: Input dataset with an equirectangular view of the point cloud and the input HDR image used for the expansion. These dataset do not have a ground truth HDR color for each point of the cloud. The point clouds were acquired by photos (multi-view stereo) or with the Kinect. The scene TOYROOM is a synthetic dataset used only to evaluate the performance of the proposed PBGI algorithm.

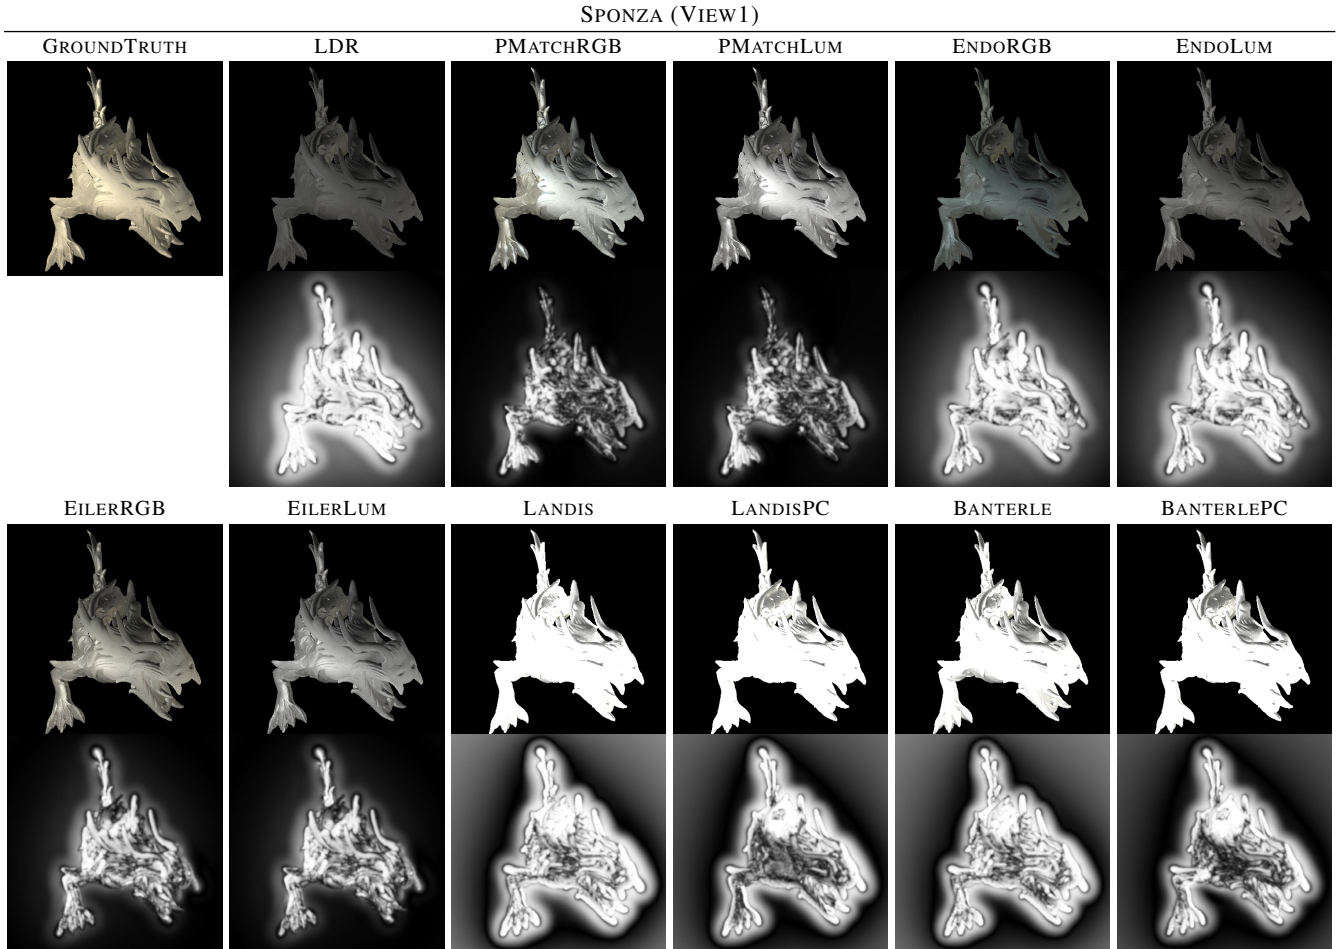


Figure 4: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Sponza computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	1.137	78.86	0.796	3.454
BANTERLE	10.58	78.03	0.843	8.314
LANDIS	16.97	77.60	0.833	12.97
BANTERLEPC	4.940	78.18	0.865	9.147
LANDISPC	6.844	78.11	0.857	17.15
EILERLUM	0.572	79.51	0.909	3.507
EILERRGB	0.567	79.50	0.909	3.507
ENDOLUM	1.045	79.00	0.811	3.454
ENDORGB	1.058	79.03	0.811	3.454
PMATCHLUM	0.442	79.73	0.941	3.555
PMATCHRGB	0.425	79.68	0.941	3.555

Table 1: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 4 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDR-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

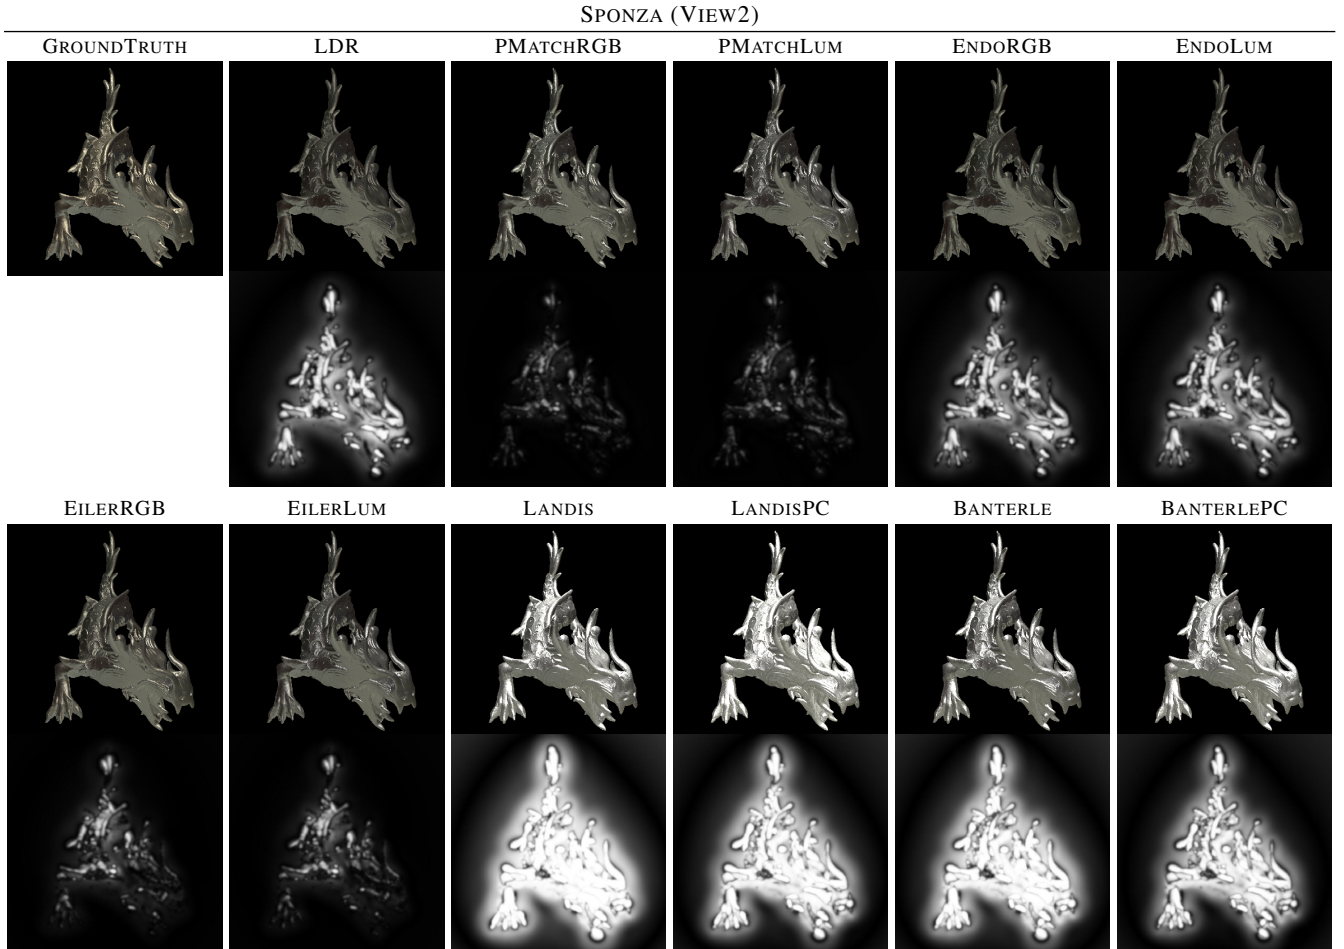


Figure 5: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Sponza computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.040	80.59	0.978	3.454
BANTERLE	0.469	78.30	0.896	8.314
LANDIS	0.820	77.70	0.871	12.97
BANTERLEPC	0.225	79.18	0.936	9.147
LANDISPC	0.351	78.69	0.909	17.15
EILERLUM	0.020	82.14	0.996	3.507
EILERRGB	0.019	82.10	0.996	3.507
ENDOLUM	0.038	80.80	0.982	3.454
ENDORGB	0.039	80.80	0.982	3.454
PMATCHLUM	0.013	83.30	0.998	3.555
PMATCHRGB	0.012	83.29	0.998	3.555

Table 2: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 5 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDR-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

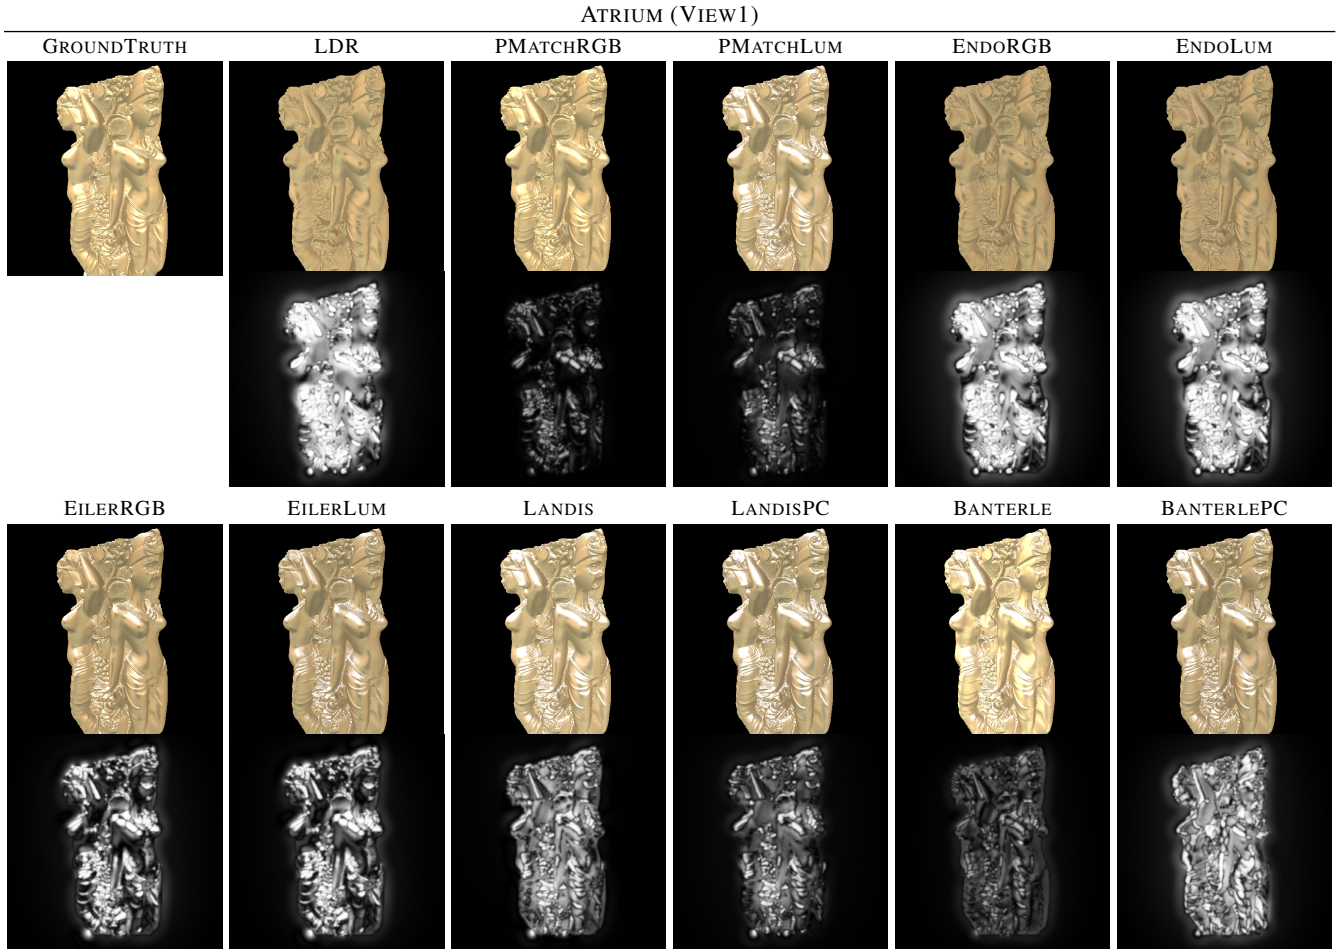


Figure 6: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Atrium computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.231	78.55	0.832	3.649
BANTERLE	0.114	81.82	0.938	2.383
LANDIS	0.175	80.03	0.932	3.349
BANTERLEPC	0.193	80.01	0.804	5.099
LANDISPC	0.150	80.52	0.948	2.834
EILERLUM	0.257	79.10	0.881	5.554
EILERRGB	0.240	79.11	0.881	5.554
ENDOLUM	0.245	78.87	0.824	3.549
ENDORGB	0.245	78.85	0.824	3.549
PMATCHLUM	0.109	81.33	0.969	2.145
PMATCHRGB	0.068	81.69	0.969	2.145

Table 3: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 6 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDR-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

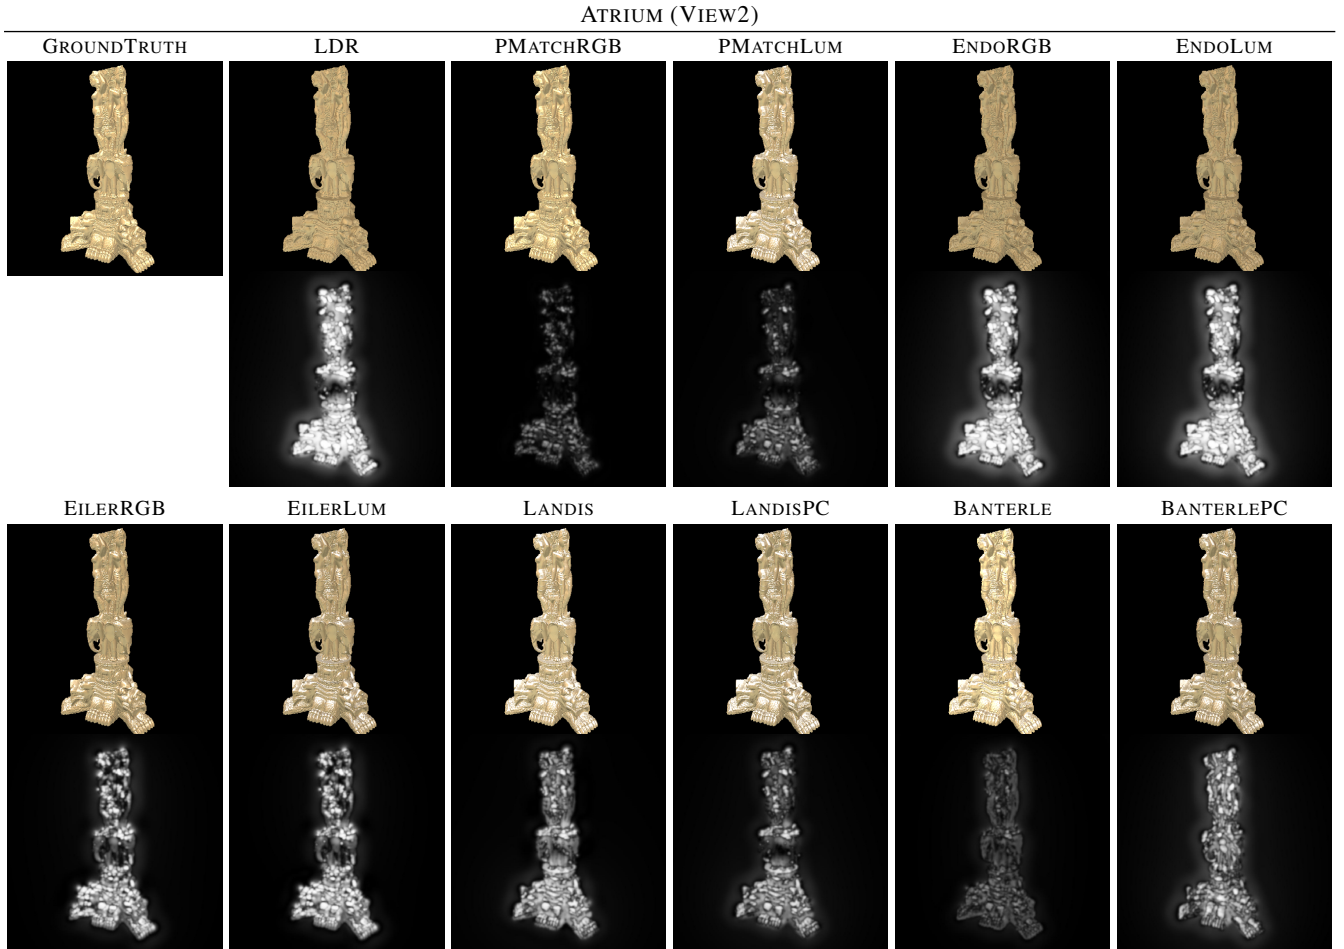


Figure 7: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Atrium computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.127	79.32	0.938	3.649
BANTERLE	0.061	82.53	0.982	2.383
LANDIS	0.101	80.45	0.973	3.349
BANTERLEPC	0.100	80.75	0.941	5.099
LANDISPC	0.088	80.93	0.974	2.834
EILERLUM	0.137	79.65	0.964	5.554
EILERRGB	0.126	79.72	0.964	5.554
ENDOLUM	0.136	79.55	0.937	3.549
ENDORGB	0.136	79.54	0.937	3.549
PMATCHLUM	0.060	81.86	0.990	2.145
PMATCHRGB	0.035	82.35	0.990	2.145

Table 4: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 7 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDR-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

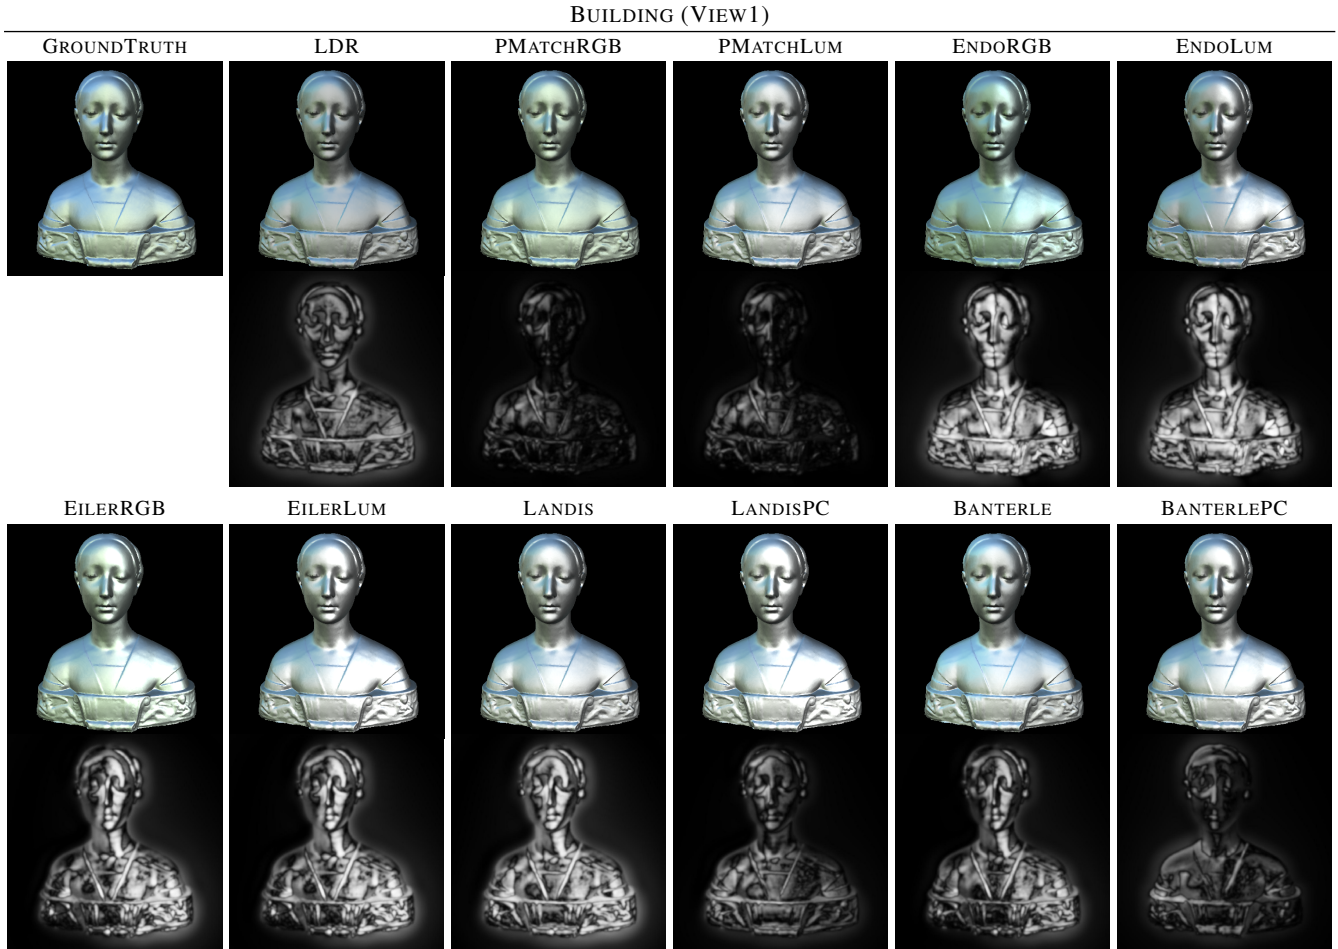


Figure 8: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Building computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.129	82.13	0.833	0.064
BANTERLE	0.182	81.63	0.913	0.087
LANDIS	0.352	80.87	0.845	0.137
BANTERLEPC	0.107	83.08	0.911	0.493
LANDISPC	0.128	82.35	0.930	0.049
EILERLUM	0.310	80.79	0.880	0.130
EILERRGB	0.278	80.78	0.880	0.130
ENDOLUM	0.154	80.77	0.835	0.100
ENDORGB	0.157	80.69	0.835	0.100
PMATCHLUM	0.059	83.38	0.985	0.046
PMATCHRGB	0.043	83.54	0.985	0.046

Table 5: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 8 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

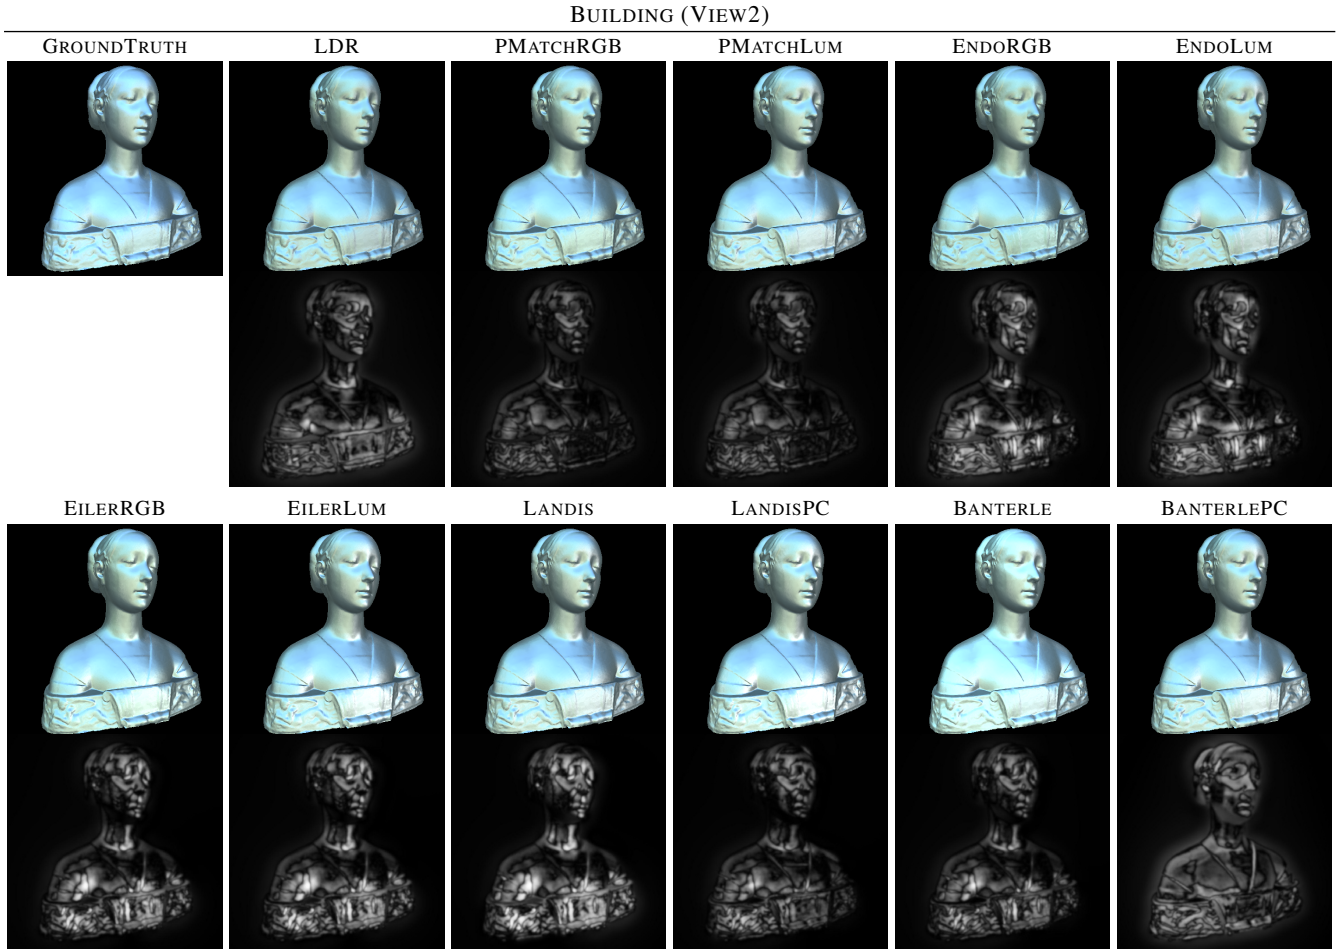


Figure 9: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Building computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.080	82.93	0.890	0.064
BANTERLE	0.065	82.78	0.933	0.087
LANDIS	0.080	81.80	0.909	0.137
BANTERLEPC	0.090	82.83	0.815	0.493
LANDISPC	0.066	82.99	0.938	0.049
EILERLUM	0.073	82.10	0.917	0.130
EILERRGB	0.071	82.16	0.917	0.130
ENDOLUM	0.076	82.52	0.893	0.100
ENDORGB	0.078	82.49	0.893	0.100
PMATCHLUM	0.065	83.61	0.938	0.046
PMATCHRGB	0.067	83.68	0.938	0.046

Table 6: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 9 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

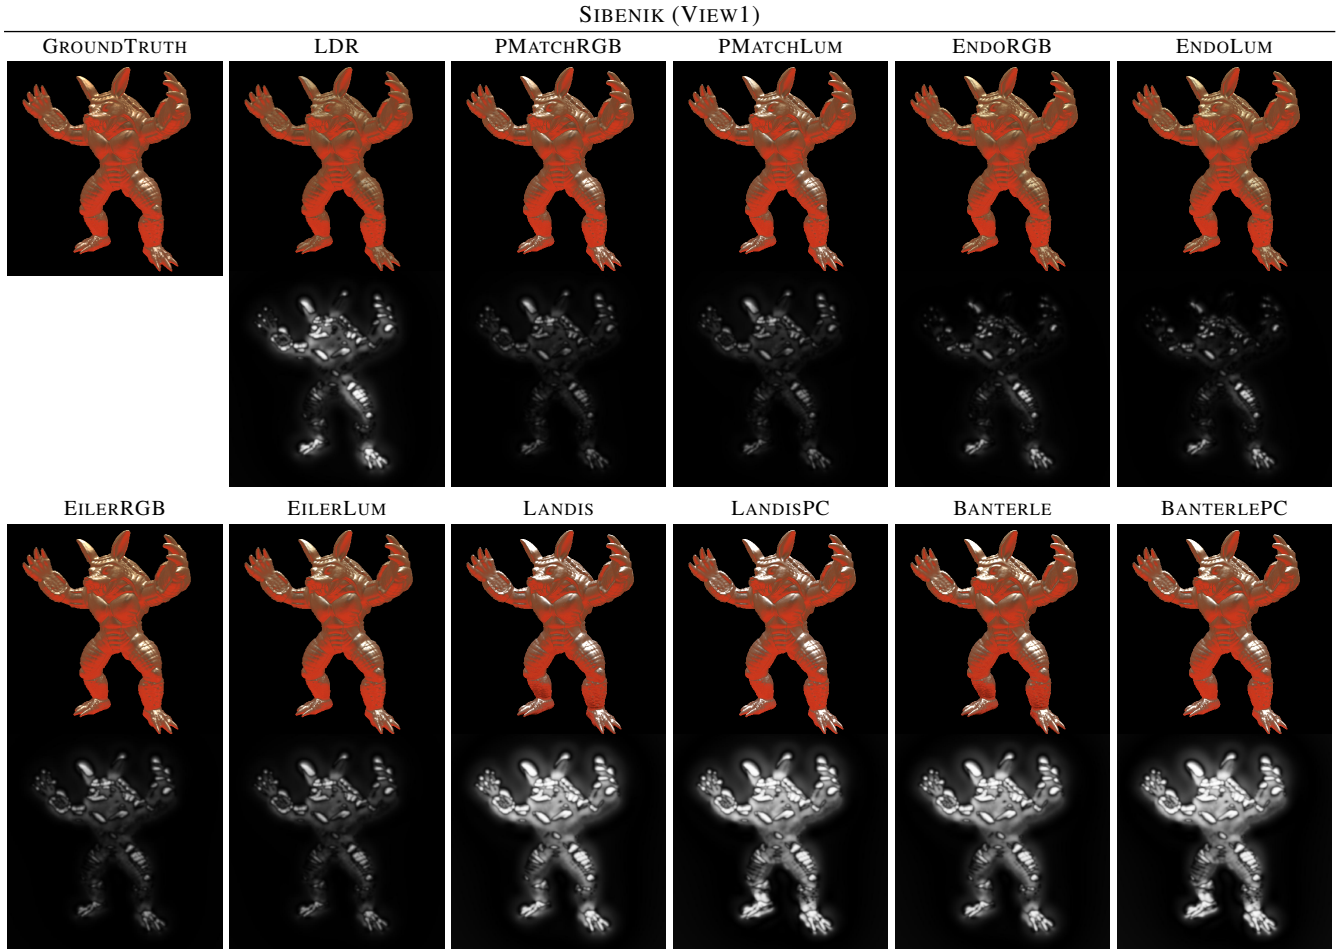


Figure 10: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Sibenik computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.042	80.99	0.986	5.090
BANTERLE	0.093	80.45	0.974	17.10
LANDIS	0.102	80.34	0.972	59.09
BANTERLEPC	0.117	80.11	0.946	15.51
LANDISPC	0.106	80.25	0.951	19.69
EILERLUM	0.028	82.17	0.994	4.878
EILERRGB	0.024	82.46	0.994	4.878
ENDOLUM	0.026	82.37	0.996	5.088
ENDORGB	0.025	82.31	0.996	5.088
PMATCHLUM	0.020	82.91	0.998	5.026
PMATCHRGB	0.017	83.04	0.998	5.026

Table 7: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 10 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

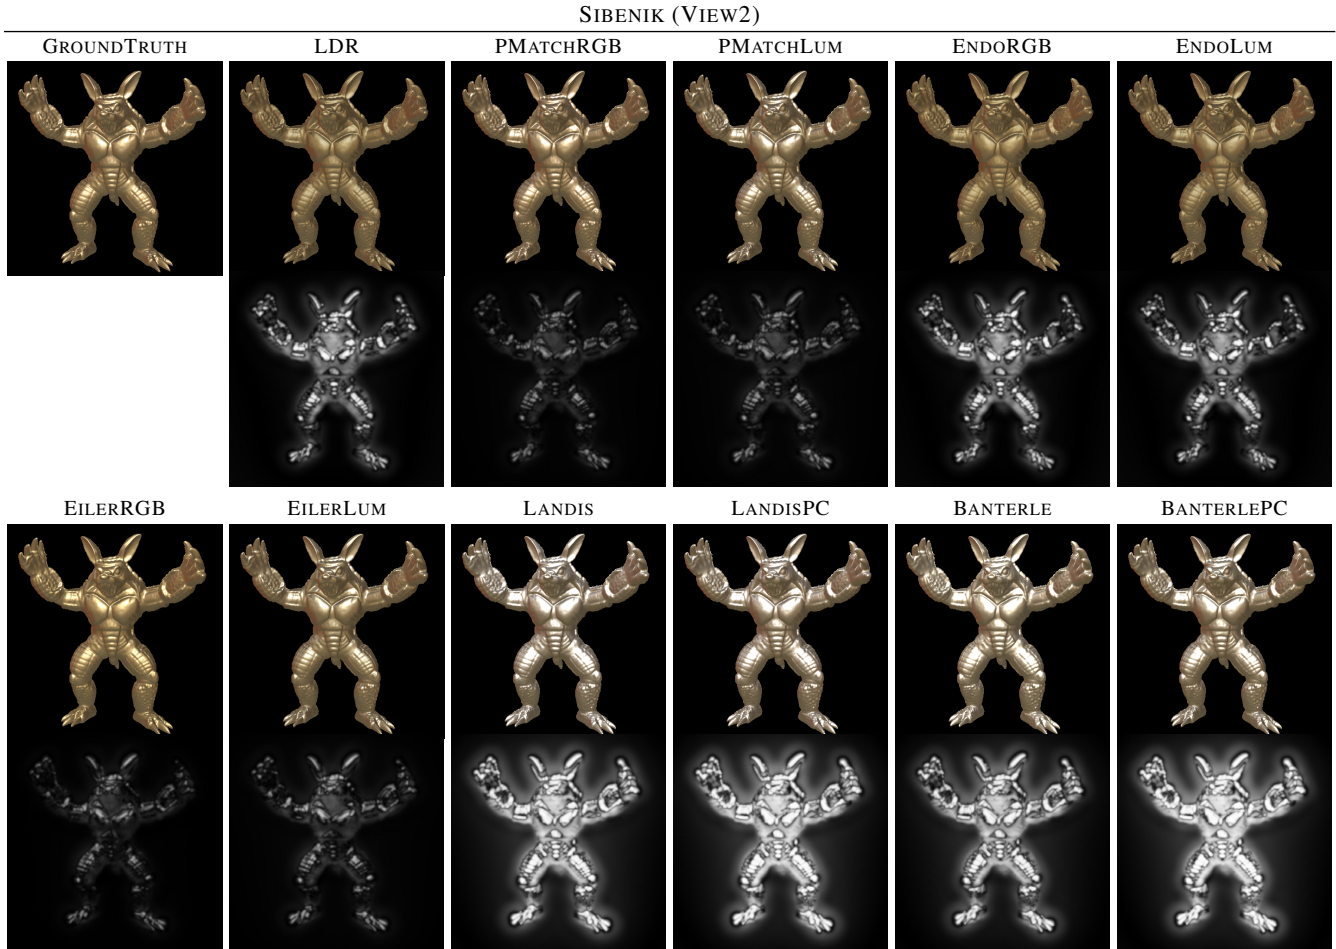


Figure 11: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Sibenik computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.073	80.83	0.976	5.090
BANTERLE	0.186	80.27	0.941	17.10
LANDIS	0.207	80.14	0.935	59.09
BANTERLEPC	0.281	79.76	0.905	15.51
LANDISPC	0.250	79.92	0.917	19.69
EILERLUM	0.034	82.38	0.991	4.878
EILERRGB	0.030	82.75	0.991	4.878
ENDOLUM	0.076	80.64	0.975	5.088
ENDORGB	0.076	80.60	0.975	5.088
PMATCHLUM	0.032	82.72	0.993	5.026
PMATCHRGB	0.027	82.85	0.993	5.026

Table 8: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 11 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

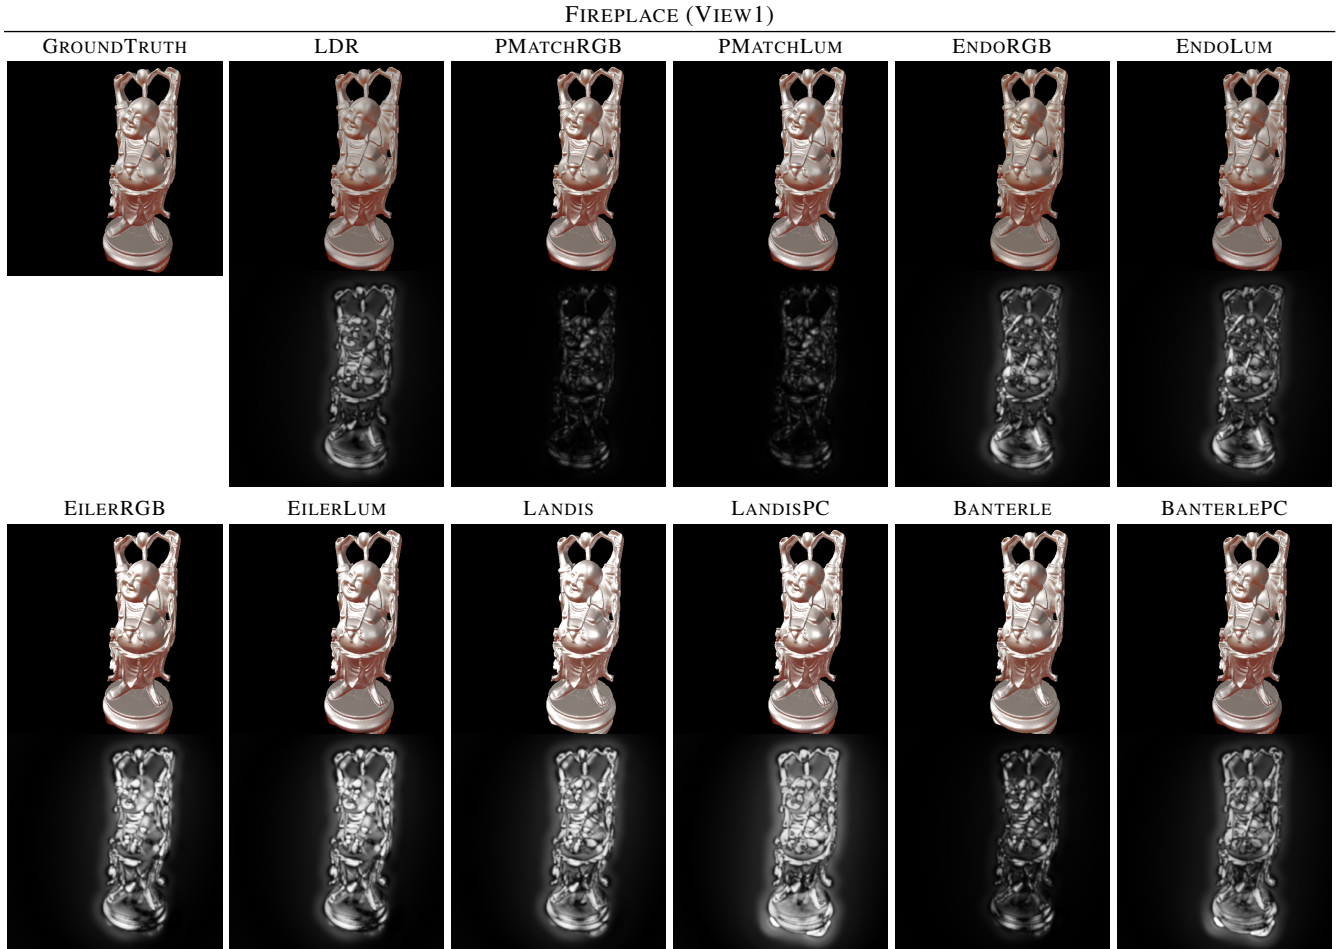


Figure 12: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud FirePlace computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.052	81.97	0.972	0.551
BANTERLE	0.050	81.87	0.962	0.611
LANDIS	0.156	80.52	0.901	0.994
BANTERLEPC	0.093	80.40	0.912	1.223
LANDISPC	0.176	79.97	0.862	0.939
EILERLUM	0.160	80.30	0.914	1.690
EILERRGB	0.160	80.31	0.914	1.690
ENDOLUM	0.048	81.57	0.963	0.638
ENDORGB	0.050	81.52	0.963	0.638
PMATCHLUM	0.012	83.53	0.990	0.544
PMATCHRGB	0.012	83.54	0.991	0.544

Table 9: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 12 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).




Figure 13: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud FirePlace computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.023	82.31	0.992	0.551
BANTERLE	0.015	82.73	0.993	0.611
LANDIS	0.051	81.21	0.978	0.994
BANTERLEPC	0.050	81.27	0.970	1.223
LANDISPC	0.085	80.55	0.955	0.939
EILERLUM	0.049	81.17	0.983	1.690
EILERRGB	0.049	81.19	0.983	1.690
ENDOLUM	0.032	81.04	0.975	0.638
ENDORGB	0.032	81.01	0.975	0.638
PMATCHLUM	0.010	83.30	0.995	0.544
PMATCHRGB	0.010	83.24	0.995	0.544

Table 10: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 13 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

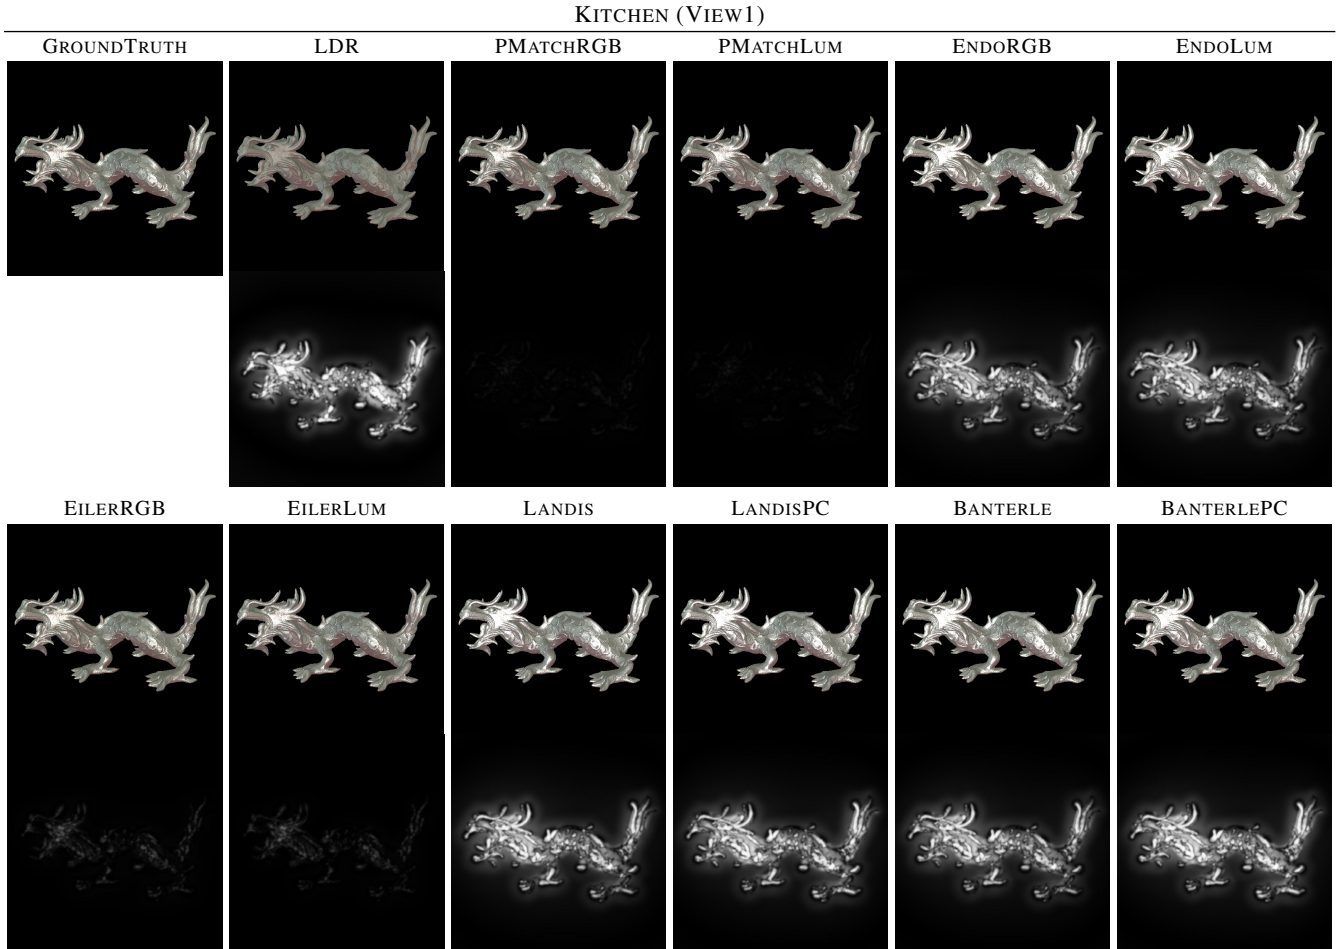


Figure 14: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Kitchen computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.087	80.14	0.955	1.806
BANTERLE	0.124	80.47	0.957	1.696
LANDIS	0.130	80.40	0.955	3.202
BANTERLEPC	0.141	80.16	0.949	3.857
LANDISPC	0.129	80.28	0.954	3.064
EILERLUM	0.009	84.31	0.998	1.248
EILERRGB	0.010	84.33	0.998	1.248
ENDOLUM	0.127	80.43	0.956	1.724
ENDORGB	0.122	80.47	0.956	1.724
PMATCHLUM	0.004	86.60	0.999	1.067
PMATCHRGB	0.004	86.56	0.999	1.067

Table 11: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 14 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).

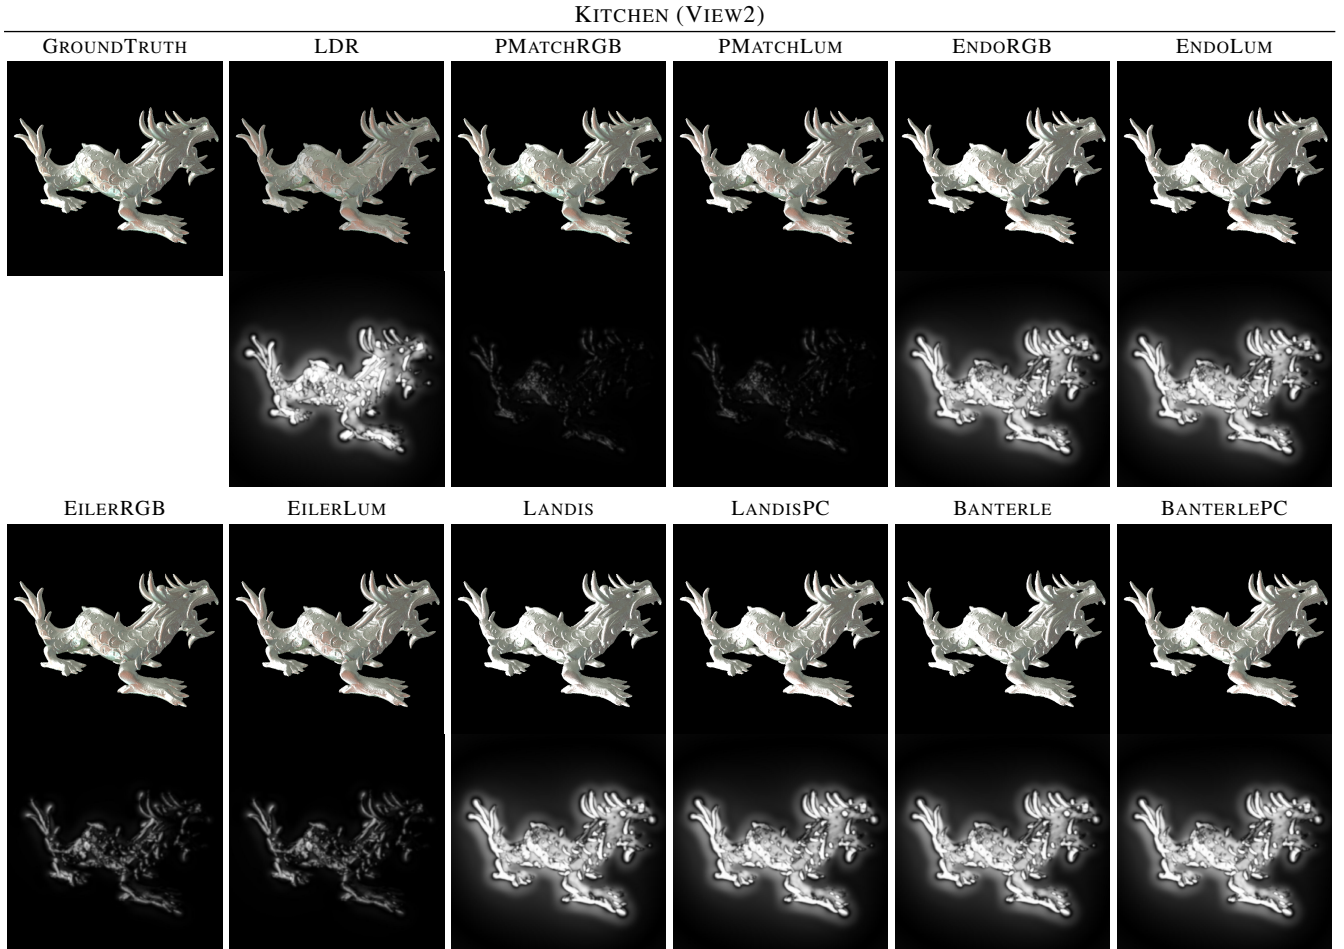


Figure 15: Ground truth comparison of the renderings obtained with the different HDR expanded versions of the expanded point cloud Kitchen computed with the methods in Section 6.2. The bottom images show the probability map to detect differences from the rendering obtained with the ground truth HDR cloud. The probability map is computed with HDR-VDR-2.2.

	RENDERINGS			POINT CLOUD
	RMS	HDR-VDP	SSIM	RMS
LDR	0.225	79.06	0.905	1.806
BANTERLE	0.493	79.28	0.897	1.696
LANDIS	0.510	79.23	0.894	3.202
BANTERLEPC	0.511	79.16	0.896	3.857
LANDISPC	0.483	79.24	0.901	3.064
EILERLUM	0.065	81.82	0.980	1.248
EILERRGB	0.062	81.87	0.980	1.248
ENDOLUM	0.501	79.25	0.895	1.724
ENDORGB	0.485	79.28	0.896	1.724
PMATCHLUM	0.028	83.52	0.993	1.067
PMATCHRGB	0.024	83.55	0.993	1.067

Table 12: The table contains the error measures of the expanded point clouds with respect to the ground truth HDR point cloud. The columns contain the error measures of the renderings in Figure 15 obtained with the X-PBGI. The used error metrics are the RMS error, the quality of HDV-VDP-2.2, and the Structure Similarity (SSIM). The green text highlights the best result for each test (for HDR-VDP and SSIM higher values are better).




Figure 16: X-PBGI rendering of a sphere inside the scene TOYROOM by varying the parameters of a Disney Principled BRDF. (Top) Rendering with increasing roughness (from 0.1 to 1.0) with fixed metalness (1.0). (Bottom) Rendering with increasing metalness (from 0.1 to 1.0) with fixed roughness (0.2).




Figure 17: Viewpoints used for the performance comparison in Fig. 18 and 19 of the three version of the PBGI algorithms.

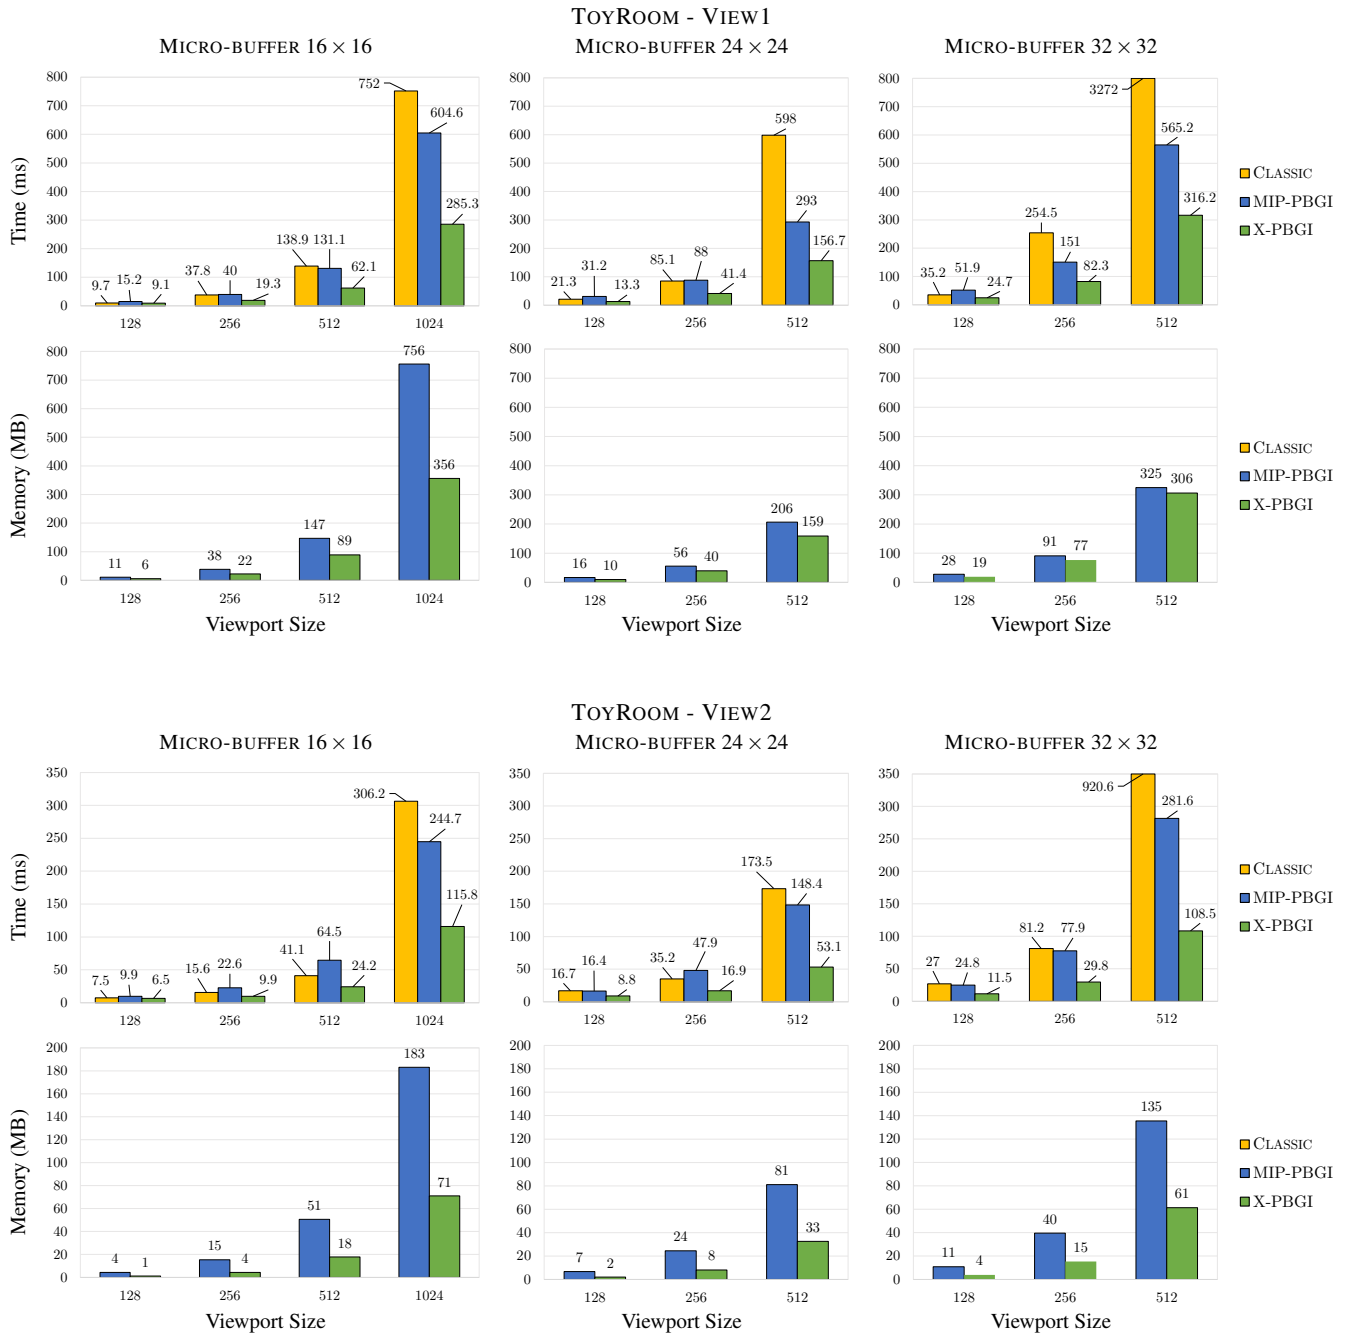


Figure 18: Performance comparison (time and memory occupancy) of the three PBGI algorithms – Classic PBGI, and the proposed MIP-PBGI and X-PBGI – varying the viewport and the micro-buffer size. These tests are performed on the point clouds TOYROOM for the two viewpoints in Fig 17. For the memory occupancy, we report the additional memory required to store the output primitives of the Geometry Shader in the Transform Feedback buffers.

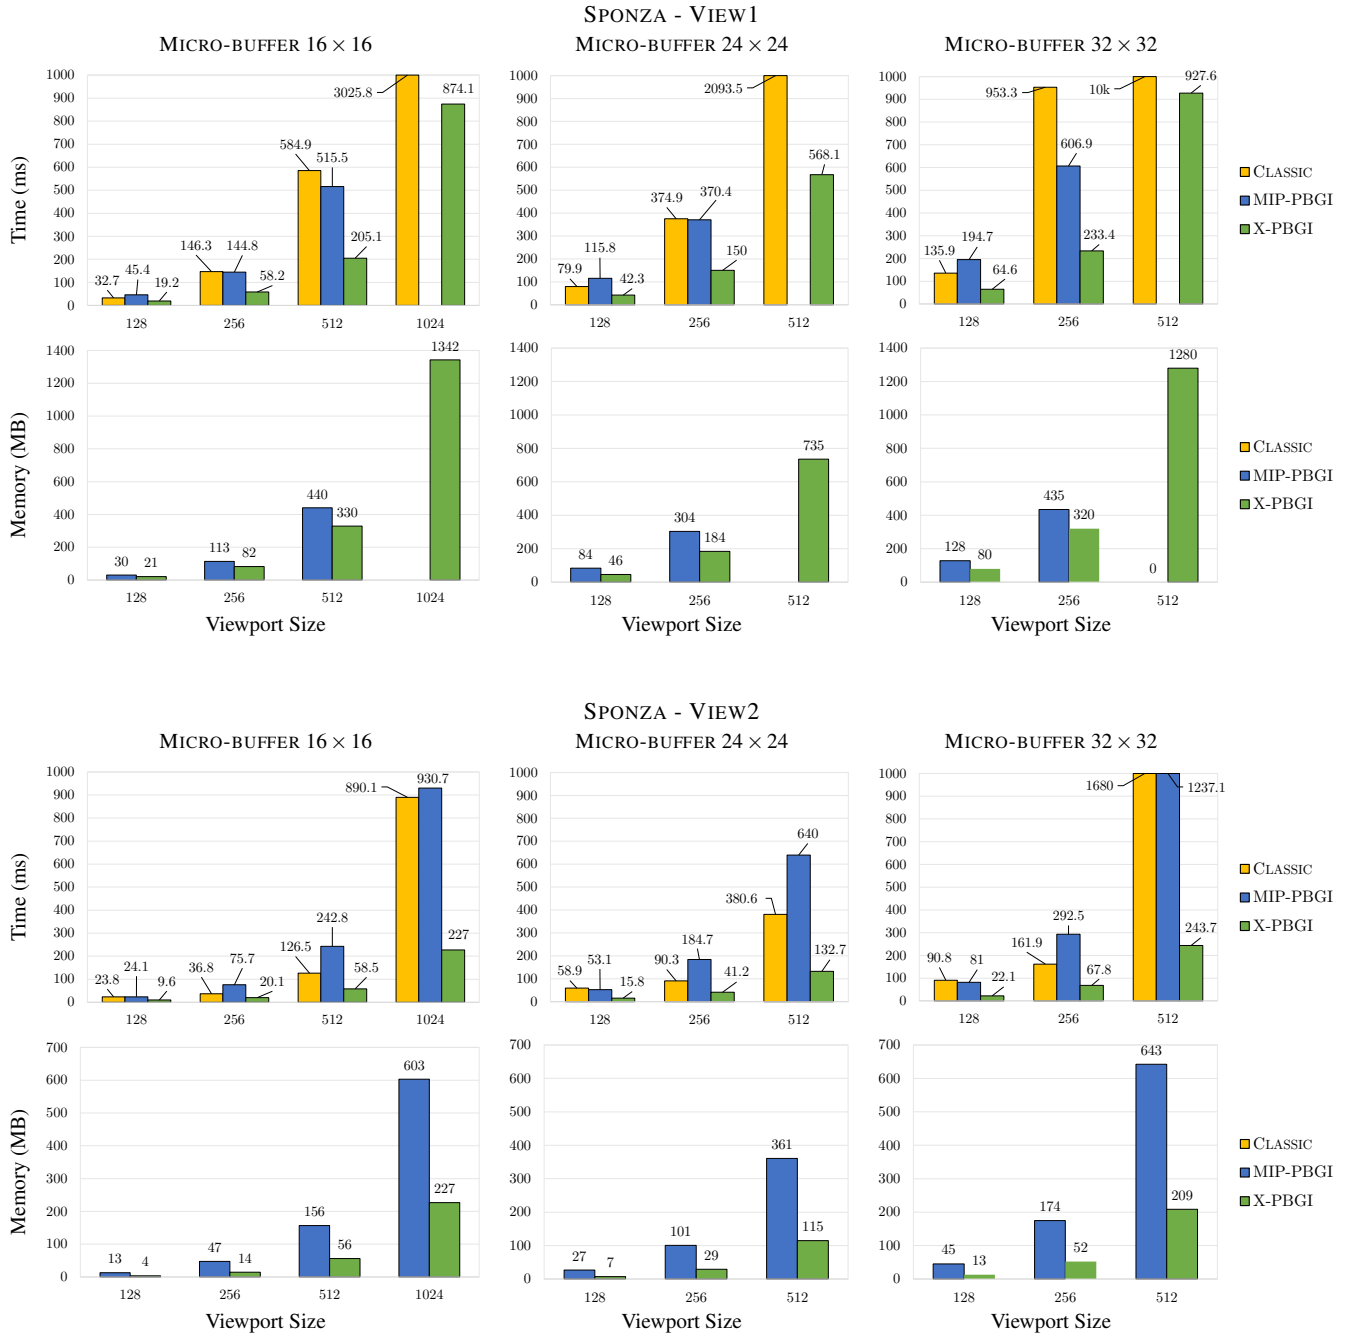


Figure 19: Performance comparison (time and memory occupancy) of the three PBGI algorithms – Classic PBGI, and the proposed MIP-PBGI and X-PBGI – varying the viewport and the micro-buffer size. These tests are performed on the point clouds SPONZA for the two viewpoints in Fig 17. For the memory occupancy, we report the additional memory required to store the output primitives of the Geometry Shader in the Transform Feedback buffers.




Figure 20: Comparison of the three different algorithms (columns) on different scenes (rows). The reported PSNR values show that changing the classic PBGI algorithm with the proposed ones (MIP-PBGI and X-PBGI) does not affect the quality of the final result.

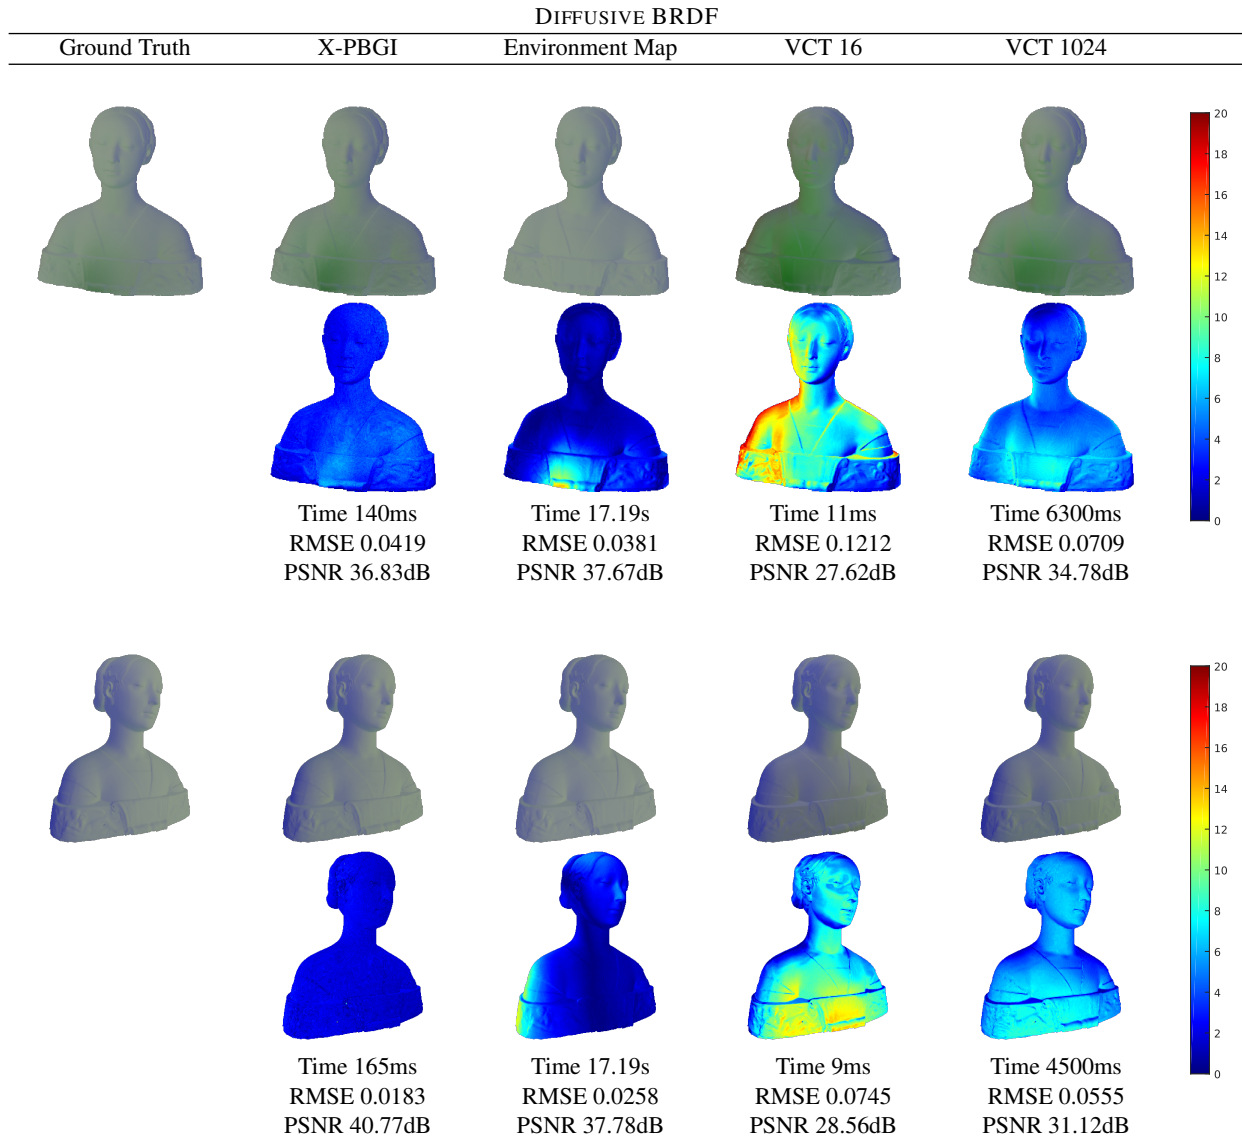


Figure 21: Comparison of the X-PBGI rendering with the ground truth obtained with a path tracing, the classical environment mapping and two different versions of the Voxel Cone Tracing (VCT) algorithm (16 cones plus a specular cone for VCT16 and 1024 cones for VCT1024). Each rendering shows the relative error map from the ground truth. The rendered object presents a pure diffuse BRDF.

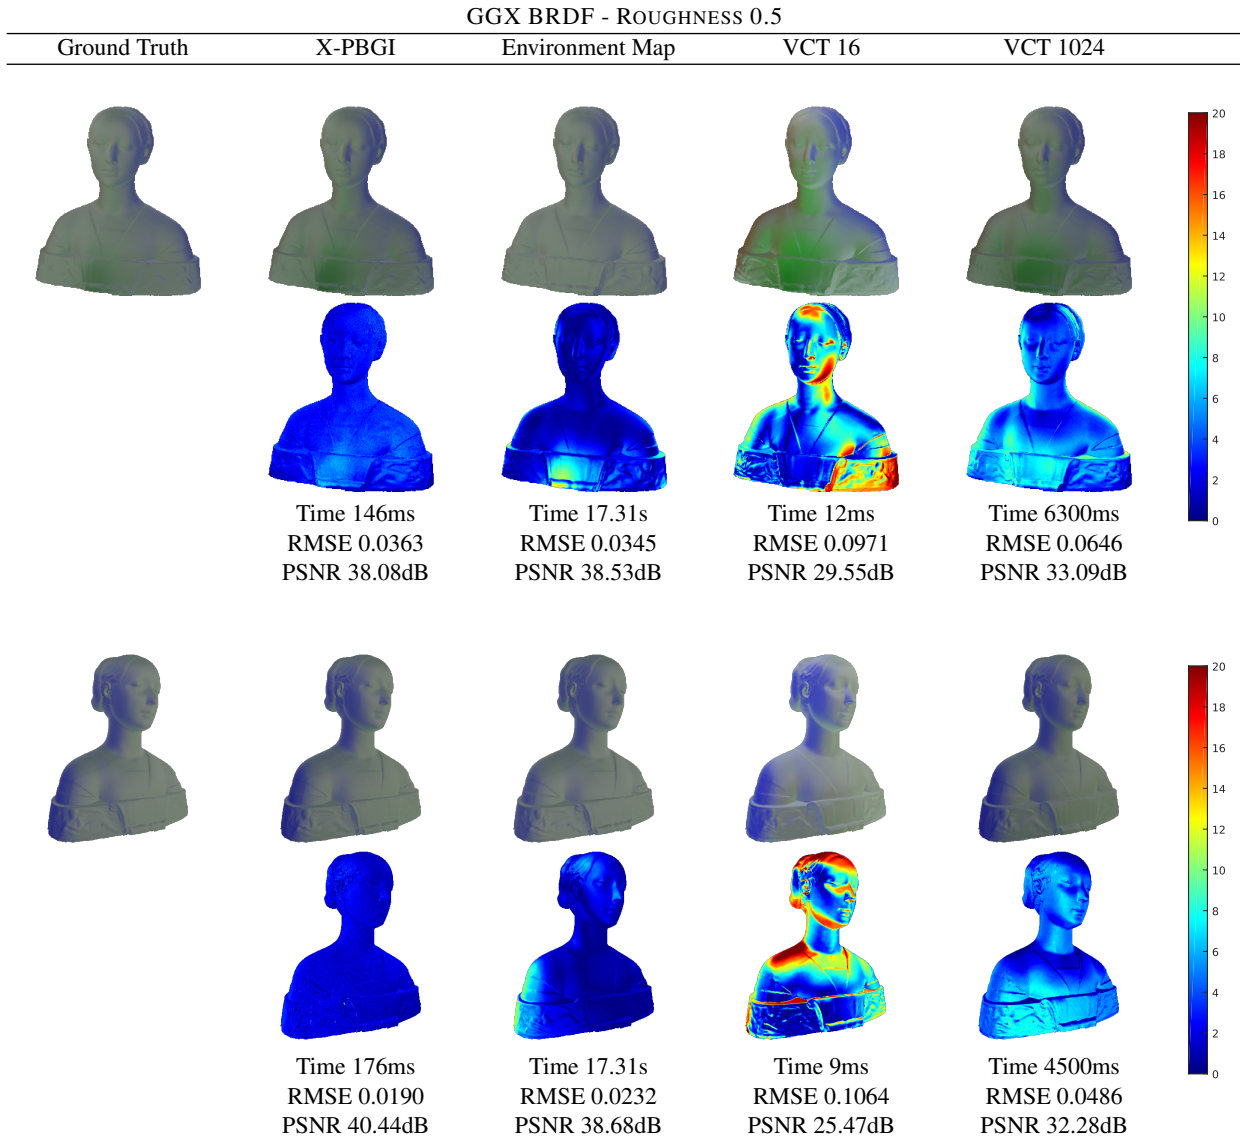


Figure 22: Comparison of the X-PBGI rendering with the ground truth obtained with a path tracing, the classical environment mapping and two different versions of the Voxel Cone Tracing (VCT) algorithm (16 cones plus a specular cone for VCT16 and 1024 cones for VCT1024). Each rendering shows the relative error map from the ground truth. The rendered object presents a GGX BRDF with roughness 0.5.




Figure 23: Comparison of the X-PBGI rendering with the ground truth obtained with a path tracing, the classical environment mapping and two different versions of the Voxel Cone Tracing (VCT) algorithm (16 cones plus a specular cone for VCT16 and 1024 cones for VCT1024). Each rendering shows the relative error map from the ground truth. The rendered object presents a GGX BRDF with roughness 0.4.

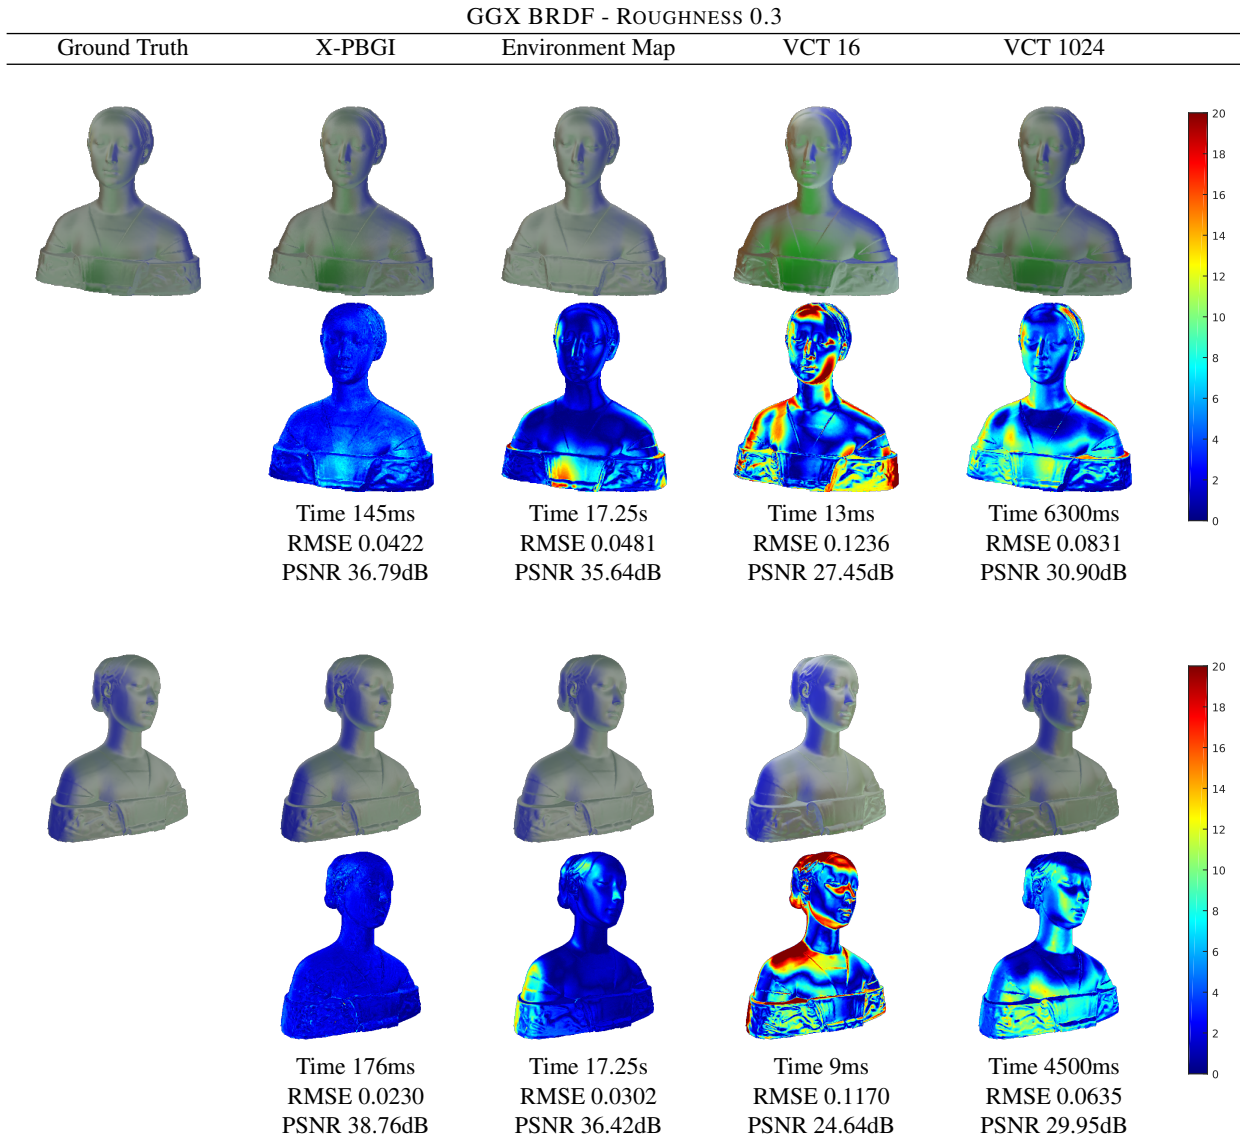


Figure 24: Comparison of the X-PBGI rendering with the ground truth obtained with a path tracing, the classical environment mapping and two different versions of the Voxel Cone Tracing (VCT) algorithm (16 cones plus a specular cone for VCT16 and 1024 cones for VCT1024). Each rendering shows the relative error map from the ground truth. The rendered object presents a GGX BRDF with roughness 0.3.

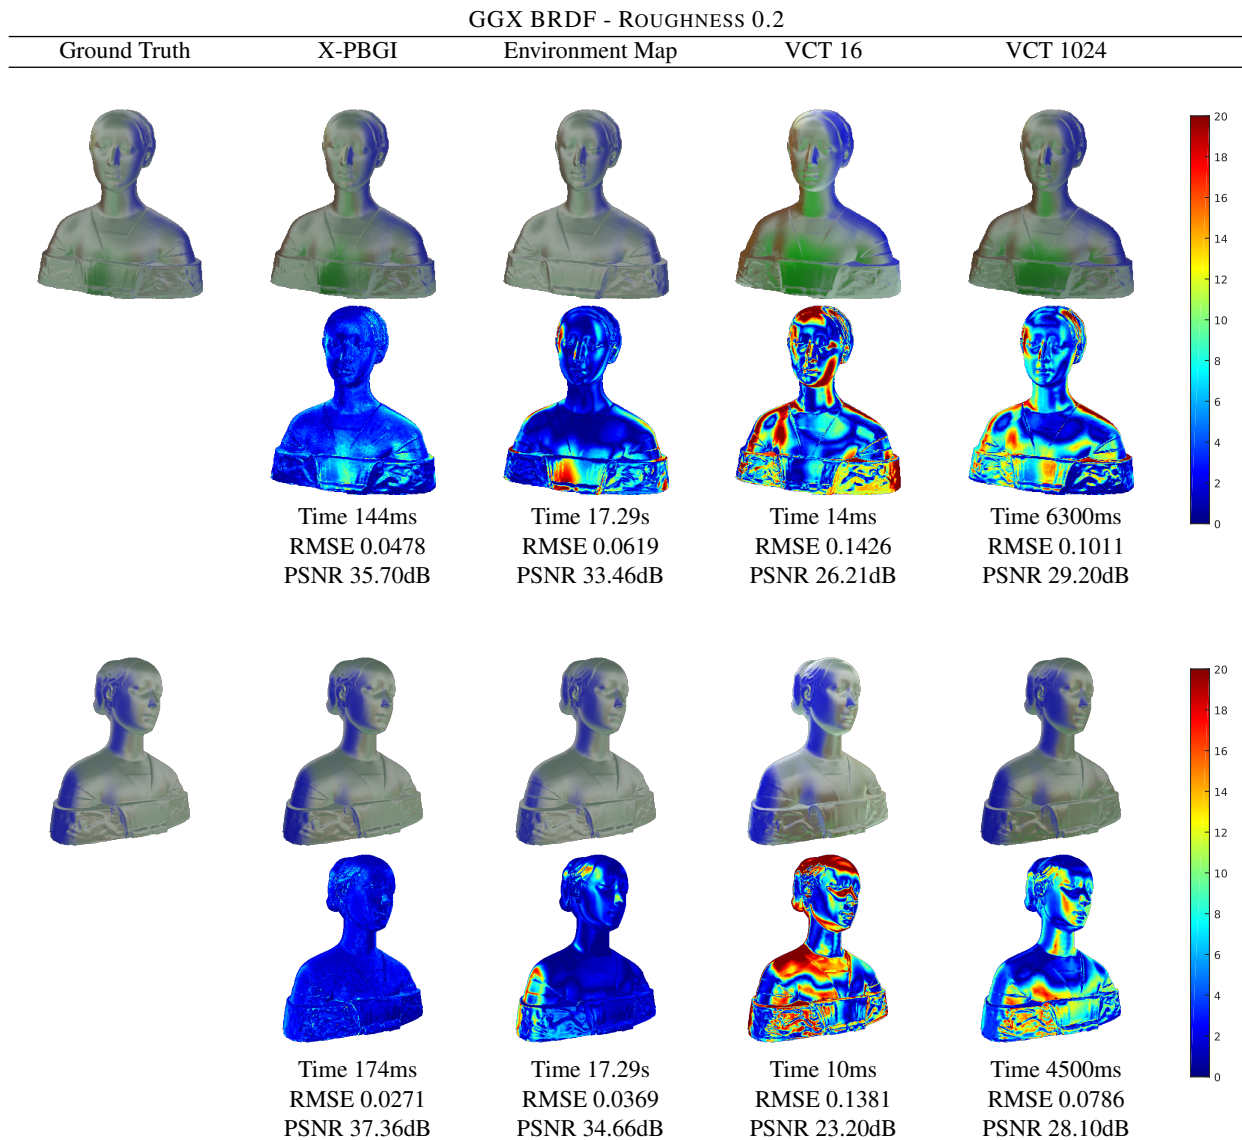


Figure 25: Comparison of the X-PBGI rendering with the ground truth obtained with a path tracing, the classical environment mapping and two different versions of the Voxel Cone Tracing (VCT) algorithm (16 cones plus a specular cone for VCT16 and 1024 cones for VCT1024). Each rendering shows the relative error map from the ground truth. The rendered object presents a GGX BRDF with roughness 0.2.

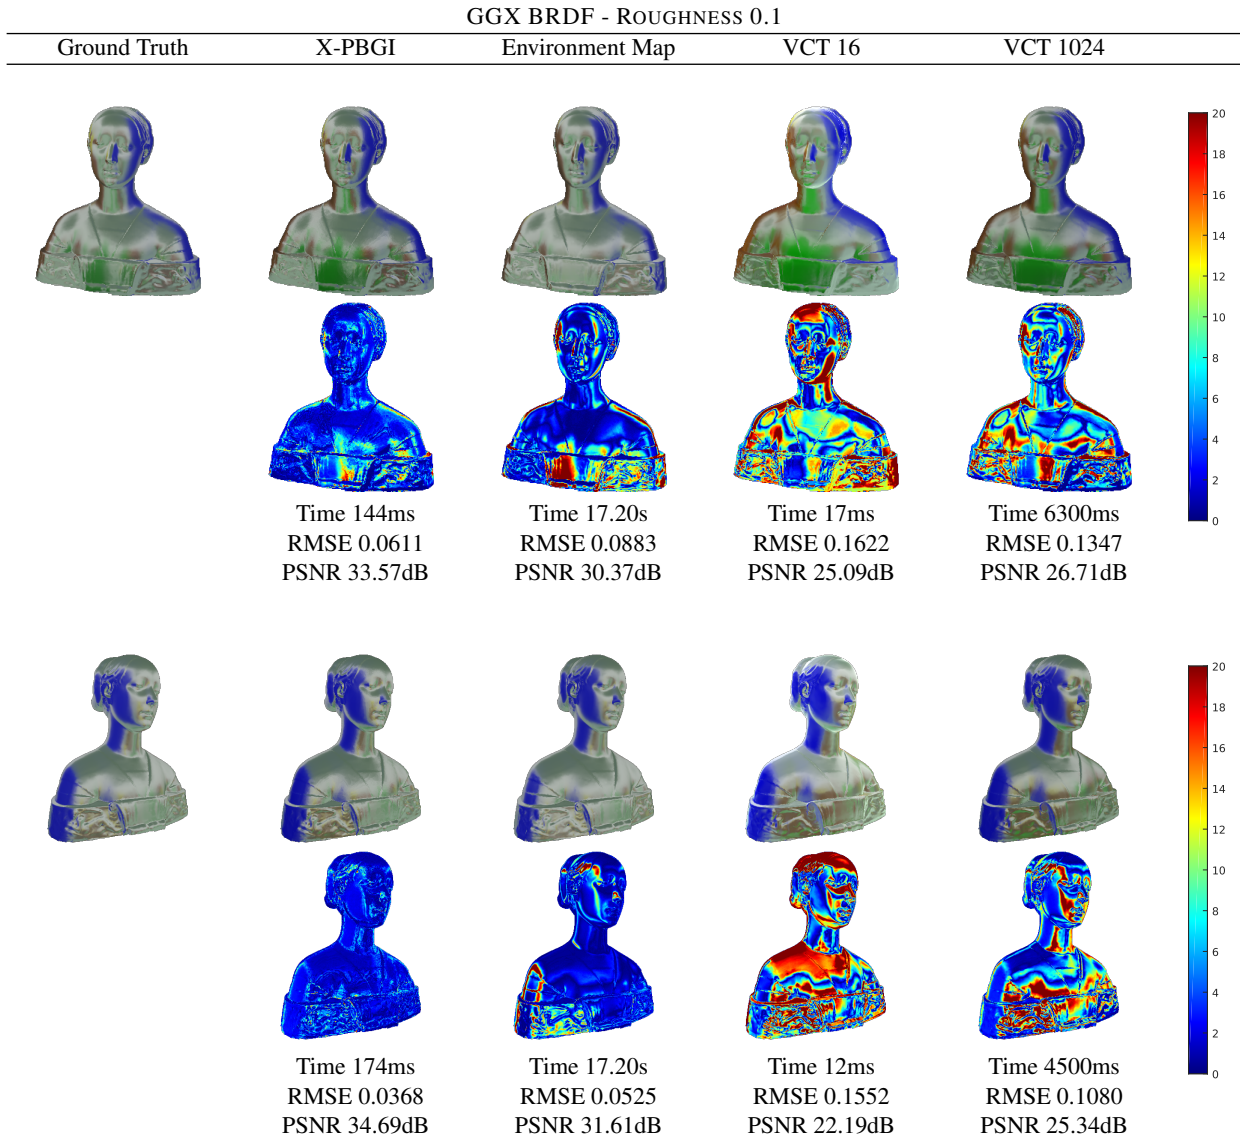


Figure 26: Comparison of the X-PBGI rendering with the ground truth obtained with a path tracing, the classical environment mapping and two different versions of the Voxel Cone Tracing (VCT) algorithm (16 cones plus a specular cone for VCT16 and 1024 cones for VCT1024). Each rendering shows the relative error map from the ground truth. The rendered object presents a GGX BRDF with roughness 0.1.