Eurographics Symposium on Rendering 2020
C. Dachsbacher and M. Pharr
(Guest Editors)

Volume 39 (2020), Number 4

Adaptive Matrix Completion for Fast Visibility Computations with Many Lights Rendering Supplementary

S. Wang[1] and N. Holzschuch[1]

z‘ [1] Univ. Grenoble-Alpes, Inria, CNRS, Grenoble INP, LJK

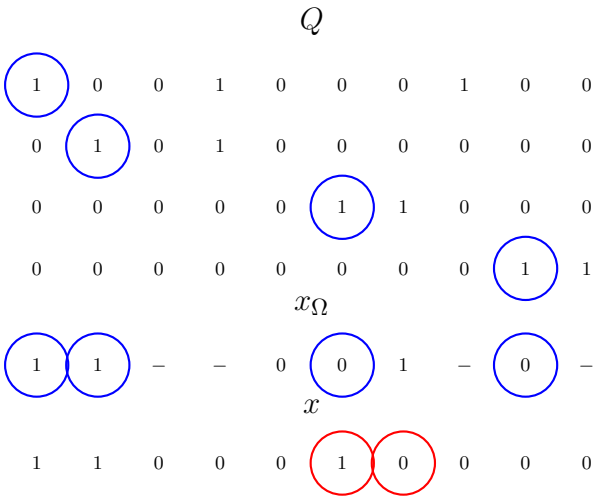

Figure 1: *A reduced matrix Q is used to reconstruct a sub-sampled row x_Ω to get x. The leading values of each row and the corresponding values in the sampled vector determine if a row should be considered. Here, the first and second rows are combined as the corresponding indices in the sub-sampled vector are 1. Not all values of x match x_Ω, indicating that x_Ω does not lie within the range space of Q.*

1. Gaussian Elimination

1.1. Reconstruction

Our aim is to reduce the basis vectors of Q so that we can easily determine whether or not a basis vector should be considered when reconstructing x. To achieve this, we define Q to be over GF(2), and perform Gauss-Jordan elimination on Q^T as the algorithm reduces rows. This puts Q^T in reduced row echelon form, where the leading 1 of each row to be the only 1 of its respective column, thus, determining whether or not to consider a row as a basis vector for x just requires checking if the corresponding index in x_Ω matches (Figure 1). As Q is defined over GF(2), the basis vectors are combined using the XOR operator.

1.2. Sampling

For sampling x_Ω, we ideally need the leading indices Υ to be in Ω in order to know which basis vectors to consider, as well as having a reasonable amount of other indices in Ω in order to check if x actually lies within the range space of Q. This is not always possible as there can be cases, in higher rank matrices, where Υ is similar in size, or even larger than Ω. To account for this, we set the number of leading indices sampled to be up to half the size of Ω, using uniform sampling to select the indices from Υ when its size exceeds that. The indices not in Υ are sampled using the same importance sampling method as described in the paper.

2. Scenes

We test our method across 12 scenes of varying geometric and material complexities. On the simple side, we have the Cornell box (Figure 4) and Sponza (Figure 10) scenes, both of which are fully diffuse and have relatively straightforward geometry. The living room scene (Figure 8) is also simple, but has a carpet that is complex geometrically, and has slightly glossy surfaces. The staircase scene (Figure 11) also has simple geometry, but likewise has glossy surfaces and the area underneath the staircase may be difficult to render accurately for clustering methods that don't handle visibility.

The San Miguel (Figure 9), classroom (Figure 5), and breakfast room (Figure 3) scenes all have complex geometry. The complexity in San Miguel is primarily caused by the large amounts of leaves on the trees. The breakfast room has complexity due to the blinds, resulting in a large number of high shadowing effects. The classroom is complex due to both the large number of windows, and the tables and chairs which cause many thin shadows.

The hairball scene (Figure 6) is the most complex geometrically, having a large amount of high frequency occluders. All of the above described scenes were rendered using 100k VPLs at 2x2 and were used to compare all methods.

We also have four highly glossy scenes, the bathroom (Figure 2), the kitchen (Figure 7), the Modern Hall (Figure 13),

and the Grey & White Room (Figure 12). These scenes also have different visibility challenges, with the bathroom having blinds and the kitchen having one of its light sources being occluded from most of the scene. Due to the glossy nature of these scenes, they were rendered using 100k VSLs at 2x2, and were used to compare our method to LightSlice. The other methods were not used here as they don't handle VSLs.

We also created tables showing our full statistics. Table 1 shows the time and RMSE for LightSlice across all scenes with varying numbers of columns. Table 2 shows the time, RMSE, and visibility sample percentage for Illumination-Cut across all scenes with varying error thresholds. Table 3 shows the time, RMSE, and visibility sample percentage for the Matrix Separation method across all scenes with varying numbers of maximum clusters. Table 4 shows the time, RMSE, and visibility sample percentage for both only clustering, and with AMC applied to visibility across all scenes with varying numbers of clusters per slice.

3. Other clustering methods

We experimented with both LightSlice and brute force (Figure 14). LightSlice was found to perform poorly due to its larger slice size and because it already partially handles visibility.

4. Importance sampling statistics

We compare our importance sampling methods to uniform sampling across three scenes. Table 5 shows these statistics. We found importance sampling to work poorly in the kitchen scene as one of the light sources are occluded from most of the scene, resulting in poor correlation with regards to direction.

5. Boolean AMC statistics

We compare our boolean AMC to the original algorithm by varying both rows and columns. Table 6 shows these statistics.

6. Adaptive Row Sample-Rate statistics

We compared our adaptive row sample-rate method to normal static row sample-rates in the classroom scene. Table 7 shows statistics of this.

7. A Comparison to Singular Value Thresholding

We performed some experiments with completing the visibility matrix with Singular Value Thresholding [CCS10], a popular nuclear-norm minimization method. For this implementation. We implemented this with both the normal version, as well as a version that uses the Randomized truncated SVD [Tul]. Figure 15 shows a couple of Cornell box scenes rendered with this method. Overall, we found this method to be far too slow, as it requires . It also requires far too many samples, and has slice-based artefacting and noise if not enough samples were assigned. The images generated in Figure 15 took 45.74 minutes for the non-truncated version, and 7.34 minutes for the truncated, 25% visibility samples were used to generate these images.

References

[CCS10] Cai, Jian-Feng, Candès, Emmanuel J, and Shen, Zuowei. "A singular value thresholding algorithm for matrix completion". *SIAM Journal on optimization* 20.4 (2010), 1956–1982 2.

[Tul] Tulloch, Andrew. *Fast Randomized SVD*. https://research.fb.com/blog/2014/09/fast-randomized-svd/. Accessed: 2020-03-22 2.

Columns		250	500	1000	2000	4000
Bathroom 100k VSLs 1024x1024 @ 4	Time	182.39s	234.11s	366.67s	662.31s	1261.27s
	Error (RMSE)	0.176	0.146	0.123	0.104	0.094
Breakfast room 100k VPLs 1280x720 @ 4	Time	100.92s	154.01s	273.36s	515.86s	998.68s
	Error (RMSE)	0.019	0.017	0.015	0.014	0.012
Cornell box 100k VPLs 800x600 @ 4	Time	26.09s	48.95s	96.36s	191.89s	378.43s
	Error (RMSE)	0.0038	0.0027	0.0019	0.0014	0.001
Classroom 100k VPLs 1280x720 @ 4	Time	122.82s	172.86s	289.17s	534.38s	1049.91s
	Error (RMSE)	0.097	0.078	0.063	0.06	0.056
Hairball 100k VPLs 800x600 @ 4	Time	80.94s	118.34s	194.89s	348.77s	658.16s
	Error (RMSE)	0.013	0.01	0.0087	0.0079	0.0065
Living room 100k VPLs 1280x720 @ 4	Time	124.67s	175.59s	297.67s	553.17s	1070.74s
	Error (RMSE)	0.019	0.015	0.012	0.0104	0.0102
Kitchen 100k VSLs 1280x720 @ 4	Time	67.53s	132.24s	264.66s	527.48s	1051.77s
	Error (RMSE)	0.125	0.115	0.098	0.082	0.069
San-Miguel 100k VPLs 1280x720 @ 4	Time	156.38s	220.4s	353.02s	624.1s	1175.71s
	Error (RMSE)	0.026	0.023	0.019	0.02	0.018
Sponza 100k VPLs 1280x720 @ 4	Time	79.45s	146.24s	273.11s	538.55s	1052.02s
	Error (RMSE)	0.017	0.013	0.011	0.009	0.0092
Staircase 100k VPLs 720x1280 @ 4	Time	130.03s	179.01s	297.74s	556.01s	1077.31s
	Error (RMSE)	0.027	0.022	0.018	0.019	0.015
Modern Hall 100k VPLs 1024x1024 @ 4	Time	169.95s	232.95s	374.96s	672.75s	1264.93s
	Error (RMSE)	0.0097	0.069	0.052	0.042	0.037
Grey & White Room 100k VPLs 1280x720 @ 4	Time	113.03s	166.71s	285.74s	531.94s	1038.54s
	Error (RMSE)	0.095	0.072	0.059	0.049	0.048

Table 1: *Statistics for our LightSlice across all scenes.*

Error threshold		10%	5%	2%	1%	0.5%
Breakfast room 100k VPLs 1280x720 @ 4	Time	101.95s	150.14s	262.91s	440.74s	745.58s
	Error (RMSE)	0.09	0.061	0.039	0.027	0.019
	Visibility samples	51.06%	47.36%	44.55%	47.55%	50.05%
Cornell box 100k VPLs 800x600 @ 4	Time	41.76s	59.21s	102.33s	161.01s	250.67s
	Error (RMSE)	0.014	0.0077	0.0047	0.0033	0.0026
	Visibility samples	1.87%	2.43%	2.58%	3.13%	3.76%
Classroom 100k VPLs 1280x720 @ 4	Time	164.12s	224.24s	343.13s	486.74s	775.85s
	Error (RMSE)	0.088	0.064	0.046	0.034	0.029
	Visibility samples	20.96%	26.44%	35.75%	44.06%	54.85%
Hairball 100k VPLs 800x600 @ 4	Time	53.68s	79.52s	122.83s	188.41s	304.46s
	Error (RMSE)	0.021	0.016	0.011	0.0083	0.0067
	Visibility samples	19.65%	20.72%	20.61%	20.61%	21.44%
Living room 100k VPLs 1280x720 @ 4	Time	122.65s	179.08s	303.48s	493.86s	772.62s
	Error (RMSE)	0.063	0.044	0.037	0.03	0.027
	Visibility samples	4.05%	3.9%	3.81%	3.9%	3.9%
San-Miguel 100k VPLs 1280x720 @ 4	Time	105.5s	146.51s	260.18s	410.68s	717.16s
	Error (RMSE)	0.074	0.059	0.054	0.042	0.034
	Visibility samples	40.3%	38.79%	38.75%	41.16%	46.48%
Sponza 100k VPLs 1280x720 @ 4	Time	153.81s	230.32s	404.11s	614.11s	975.07s
	Error (RMSE)	0.117	0.075	0.056	0.047	0.035
	Visibility samples	18.42%	15.53%	16.66%	18.94%	22.62%
Staircase 100k VPLs 720x1280 @ 4	Time	95.51s	138.49s	261.03s	436.58s	772.84s
	Error (RMSE)	0.056	0.049	0.036	0.029	0.026
	Visibility samples	3.3%	3.4%	3.56%	3.6%	3.61%

Table 2: *Statistics for IlluminationCut across all scenes.*

Max clusters		500	1000	2000	4000	8000
Breakfast room 100k VPLs 1280x720 @ 4	Time	30.46s	74.44s	173.18s	464.6s	1182.892s
	Error (RMSE)	0.035	0.034	0.034	0.034	0.033
	Visibility samples	50.48%	39.33%	30.17%	30.51%	29.67%
Cornell box 100k VPLs 800x600 @ 4	Time	7.83s	17.33s	36.16s	73.56s	135.44s
	Error (RMSE)	0.01	0.0094	0.0094	0.0094	0.0094
	Visibility samples	82.56%	82.46%	82.56%	81.77%	73.58%
Classroom 100k VPLs 1280x720 @ 4	Time	36.63s	76.03s	162.33s	353.43s	742.95s
	Error (RMSE)	0.032	0.029	0.027	0.026	0.026
	Visibility samples	38.41%	29.46%	23.37%	18.76%	14.65%
Hairball 100k VPLs 800x600 @ 4	Time	28.19s	65.52s	152.11s	344.92s	781.92s
	Error (RMSE)	0.02	0.02	0.02	0.02	0.019
	Visibility samples	45.2%	36.9%	31.81%	29.4%	51.39%
Living room 100k VPLs 1280x720 @ 4	Time	30.45s	68.69s	155.7s	350.38s	760.42s
	Error (RMSE)	0.014	0.013	0.012	0.014	0.014
	Visibility samples	65.09%	60.11%	60.39%	62.88%	65.12%
San-Miguel 100k VPLs 1280x720 @ 4	Time	28.23s	57.32s	125.23s	277.77s	618.08s
	Error (RMSE)	0.041	0.038	0.036	0.035	0.034
	Visibility samples	26.16%	16.9%	12.05%	9.42%	7.62%
Sponza 100k VPLs 1280x720 @ 4	Time	25.93s	60.65s	136.47s	299.73s	629.65s
	Error (RMSE)	0.029	0.026	0.026	0.026	0.025
	Visibility samples	50.7%	46.7%	45.43%	44.72%	43.82%
Staircase 100k VPLs 720x1280 @ 4	Time	38.26s	86.55s	182.42s	394.77s	835.47s
	Error (RMSE)	0.018	0.016	0.015	0.015	0.015
	Visibility samples	71.86%	67.47%	60.26%	57.73%	55.63%

Table 3: *Statistics for Matrix Separation across all scenes.*

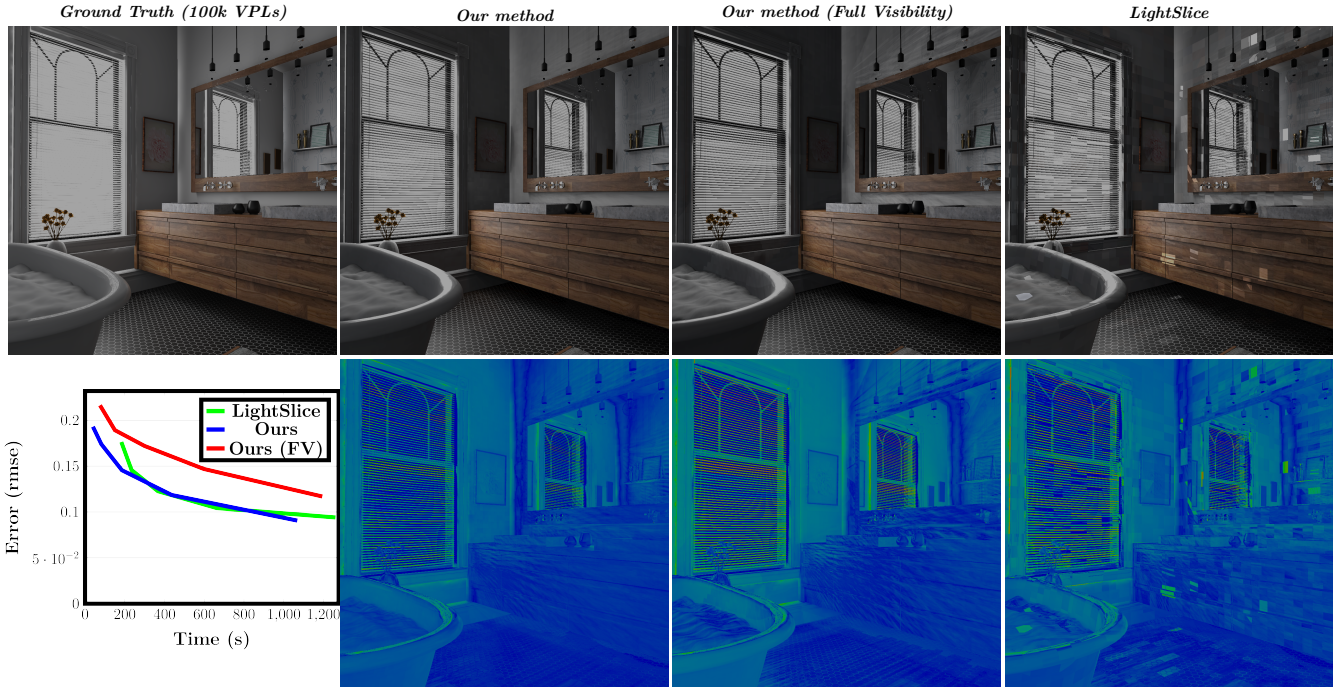

Figure 2: *Equal time renders (~185s) and time vs error plot for bathroom scene.*

Clusters per slice		250	500	1000	2000	4000	8000
Bathroom 100k VSLs 1024x1024 @ 4	Time		39.42s	81.98s	185.48s	436.16s	1068.98s
	Time (only clustering)	74.97s	150.26s	299.33s	597.96s	1194.73s	
	Error (RMSE)		0.193	0.173	0.146	0.119	0.091
	Error (RMSE, only clustering)	0.216	0.189	0.172	0.149	0.119	
	Visibility samples		19.92%	21.16%	23.52%	26.48%	29.4%
Breakfast room 100k VPLs 1280x720 @ 4	Time		30.89s	63.82s	145.78s	348.78s	845.09s
	Time (only clustering)	61.74s	123.12s	243.96s	489.2s	980.99s	
	Error (RMSE)		0.017	0.012	0.0083	0.0067	0.0067
	Error (RMSE, only clustering)	0.027	0.016	0.0105	0.008	0.0061	
	Visibility samples		18.9%	19.95%	22.76%	27.36%	31.99%
Cornell box 100k VPLs 800x600 @ 4	Time		7.51s	13.83s	28.19s	60.33s	129.12s
	Time (only clustering)	24.28s	46.95s	93.72s	189.95s	381.05s	
	Error (RMSE)		0.0037	0.0022	0.0016	0.0011	0.0008
	Error (RMSE, only clustering)	0.0065	0.0039	0.0023	0.0014	0.0011	
	Visibility samples		9.2%	9.26%	9.89%	10.99%	12.68%
Classroom 100k VPLs 1280x720 @ 4	Time		26.38s	54.99s	129.243s	310.51s	755.81s
	Time (only clustering)	65s	125.44s	253.82s	504.44s	1022.6s	
	Error (RMSE)		0.024	0.0196	0.0162	0.0152	0.0125
	Error (RMSE, only clustering)	0.033	0.024	0.02	0.016	0.015	
	Visibility samples		15.71%	16.83%	19.96%	23.77%	27.33%
Hairball 100k VPLs 800x600 @ 4	Time		40.56s	86.11s	194.23s	460.43s	1150.02s
	Time (only clustering)	40.53s	79.35s	157.71s	316.36s	639.95s	
	Error (RMSE)		0.006	0.0041	0.003	0.0023	0.0019
	Error (RMSE, only clustering)	0.0088	0.0059	0.004	0.0029	0.0022	
	Visibility samples		36.71%	37.98%	40.19%	43.11%	46.45%
Living room 100k VPLs 1280x720 @ 4	Time		27.09s	52.07s	109.36s	244.1s	574.04s
	Time (only clustering)	66.06s	132.77s	263.48s	526.1s	1043.42s	
	Error (RMSE)		0.019	0.012	0.0086	0.0076	0.0072
	Error (RMSE, only clustering)	0.024	0.018	0.012	0.0093	0.0074	
	Visibility samples		14.35%	13.46%	13.9%	15.53%	18.45%
Kitchen 100k VSLs 1280x720 @ 4	Time		25.2s	51.49s	114.15s	262.41s	612.66s
	Time (only clustering)	67.53s	132.24s	264.66s	527.48s	1051.77s	
	Error (RMSE)		0.112	0.097	0.083	0.068	0.054
	Error (RMSE, only clustering)	0.125	0.115	0.098	0.082	0.069	
	Visibility samples		15.56%	15.03%	16.89%	19.6%	22.57%
San-Miguel 100k VPLs 1280x720 @ 4	Time		30.77s	67.33s	158.75s	388.64s	966.41s
	Time (only clustering)	69.92s	139.83s	278.03s	549.3s	1103.81s	
	Error (RMSE)		0.023	0.0183	0.0181	0.0134	0.0135
	Error (RMSE, only clustering)	0.031	0.023	0.017	0.015	0.014	
	Visibility samples		12.97%	14.38%	16.64%	19.96%	23.15%
Sponza 100k VPLs 1280x720 @ 4	Time		16.3s	32.6s	71.76s	175.59s	442.49s
	Time (only clustering)	66.99s	131.17s	261.89s	521.69s	1043.18s	
	Error (RMSE)		0.026	0.015	0.0092	0.0067	0.0048
	Error (RMSE, only clustering)	0.043	0.03	0.015	0.0094	0.006	
	Visibility samples		7.16%	7.5%	8.93%	11.81%	16.16%
Staircase 100k VPLs 720x1280 @ 4	Time		25.58s	50.05s	106.82s	243.37s	583.3s
	Time (only clustering)	65s	132.14s	261.02s	522.5s	1048.93s	
	Error (RMSE)		0.017	0.012	0.0088	0.0064	0.0057
	Error (RMSE, only clustering)	0.023	0.017	0.012	0.0088	0.0085	
	Visibility samples		14.26%	14.01%	14.96%	17%	20.23%
Modern Hall 100k VPLs 1024x1024 @ 4	Time		41.52s	63.33s	105.38s	230.39s	551.69s
	Time (only clustering)	76.43s	151.05s	302.72s	600.54s	1204.17s	
	Error (RMSE)		0.102	0.072	0.052	0.041	0.034
	Error (RMSE, only clustering)	0.142	0.099	0.07	0.054	0.042	
	Visibility samples		10.16%	12.58%	13.65%	15.97%	19.3%
Grey & White Room 100k VPLs 1280x720 @ 4	Time		34.13s	52.1s	93.82s	206.88s	485.78s
	Time (only clustering)	64.9s	127.84s	251.43s	502.33s	998.86s	
	Error (RMSE)		0.139	0.089	0.063	0.047	0.04
	Error (RMSE, only clustering)	0.188	0.141	0.01	0.06	0.046	
	Visibility samples		13.89%	13.65%	14.92%	17.6%	21.68%

Table 4: *Statistics for our method across all scenes.*

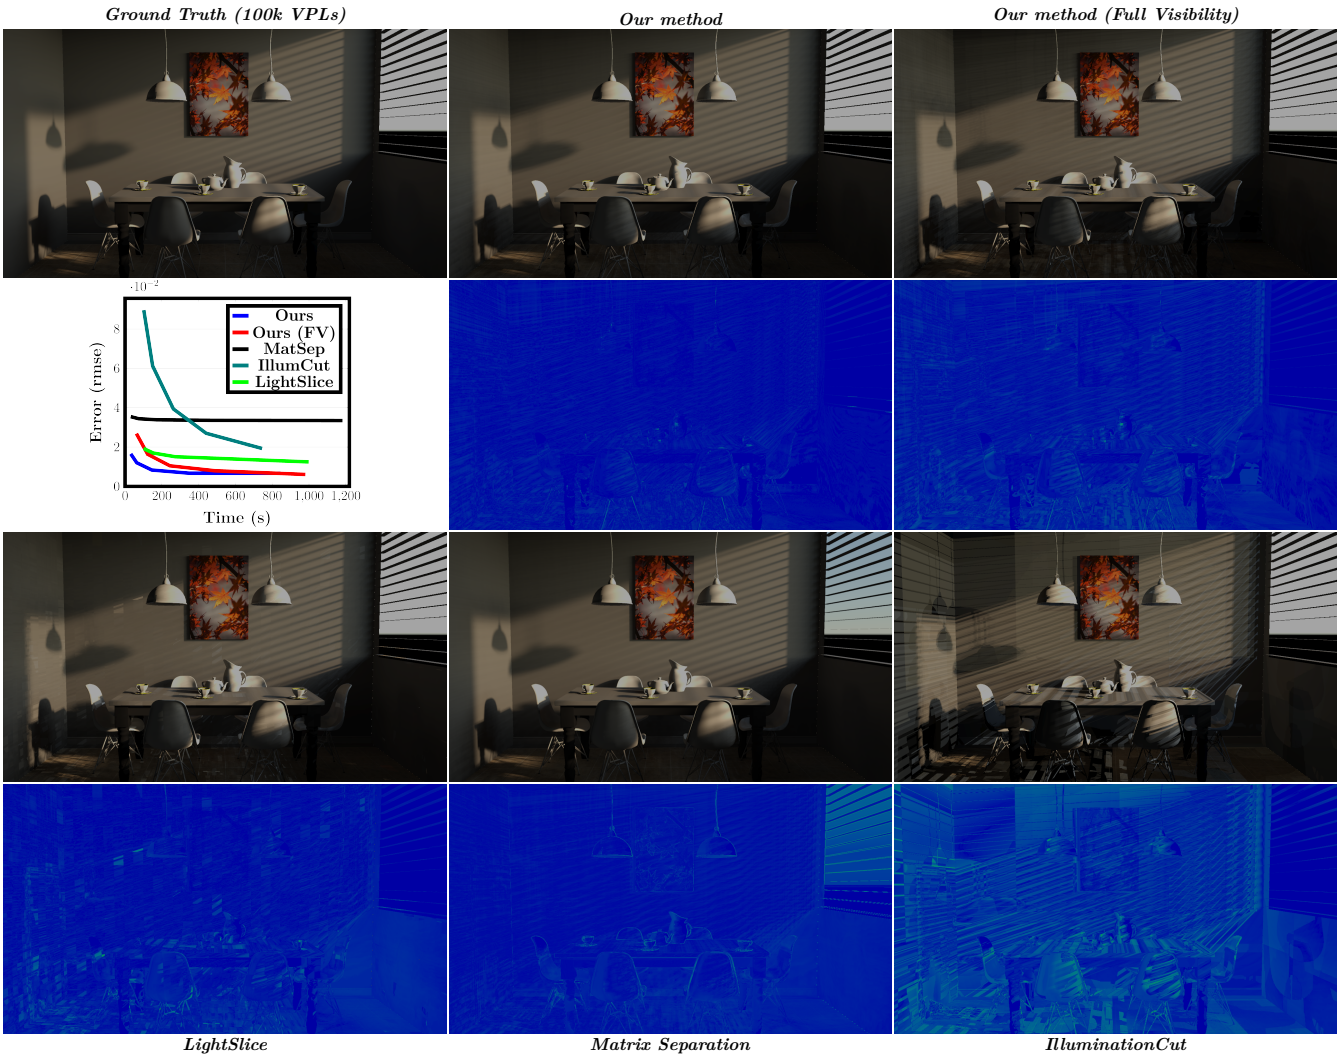

Figure 3: *Equal time renders (~80s) and time vs error plot for breakfast room scene*

		San Miguel 1280 x 720 @ 4, 100k VPLs			Classroom 1280 x 720 @ 4, 100k VPLs			Kitchen 1280 x 720 @ 4, 100k VPLs		
		Time	RMSE	Samples	Time	RMSE	Samples	Time	RMSE	Samples
Uniform Sampling	α=5%	51.7s	0.0437	9.5%	50.86s	0.0359	10.8%	52.67s	0.091	9.67%
	α=10%	85.82s	0.0319	15.14%	77.88s	0.0262	16.9%	81.26s	0.0889	15.57%
	α=15%	180.59s	0.0275	20.48%	210.76s	0.0238	22.36%	187.6s	0.0877	20.9%
	α=20%	607.81s	0.0247	25.69%	795.513s	0.021	27.51%	613.34s	0.088	26%
Importance Sampling	α=5%	48.01s	0.0348	9.7%	37.19s	0.034	10.9%	43.24s	0.09	10.09%
	α=10%	88.46s	0.0268	15.44%	71.92s	0.0249	17.02%	80.92s	0.0889	16.02%
	α=15%	190.24s	0.0218	20.77%	196.8s	0.0219	22.4%	198.12s	0.0888	21.3%
	α=20%	576.57s	0.0211	25.88%	801.92s	0.02	27.47%	640.53s	0.0894	26.35%

Table 5: *Importance sampling statistics when compared to no importance sampling in the Classroom, San Miguel, and Kitchen scenes.*

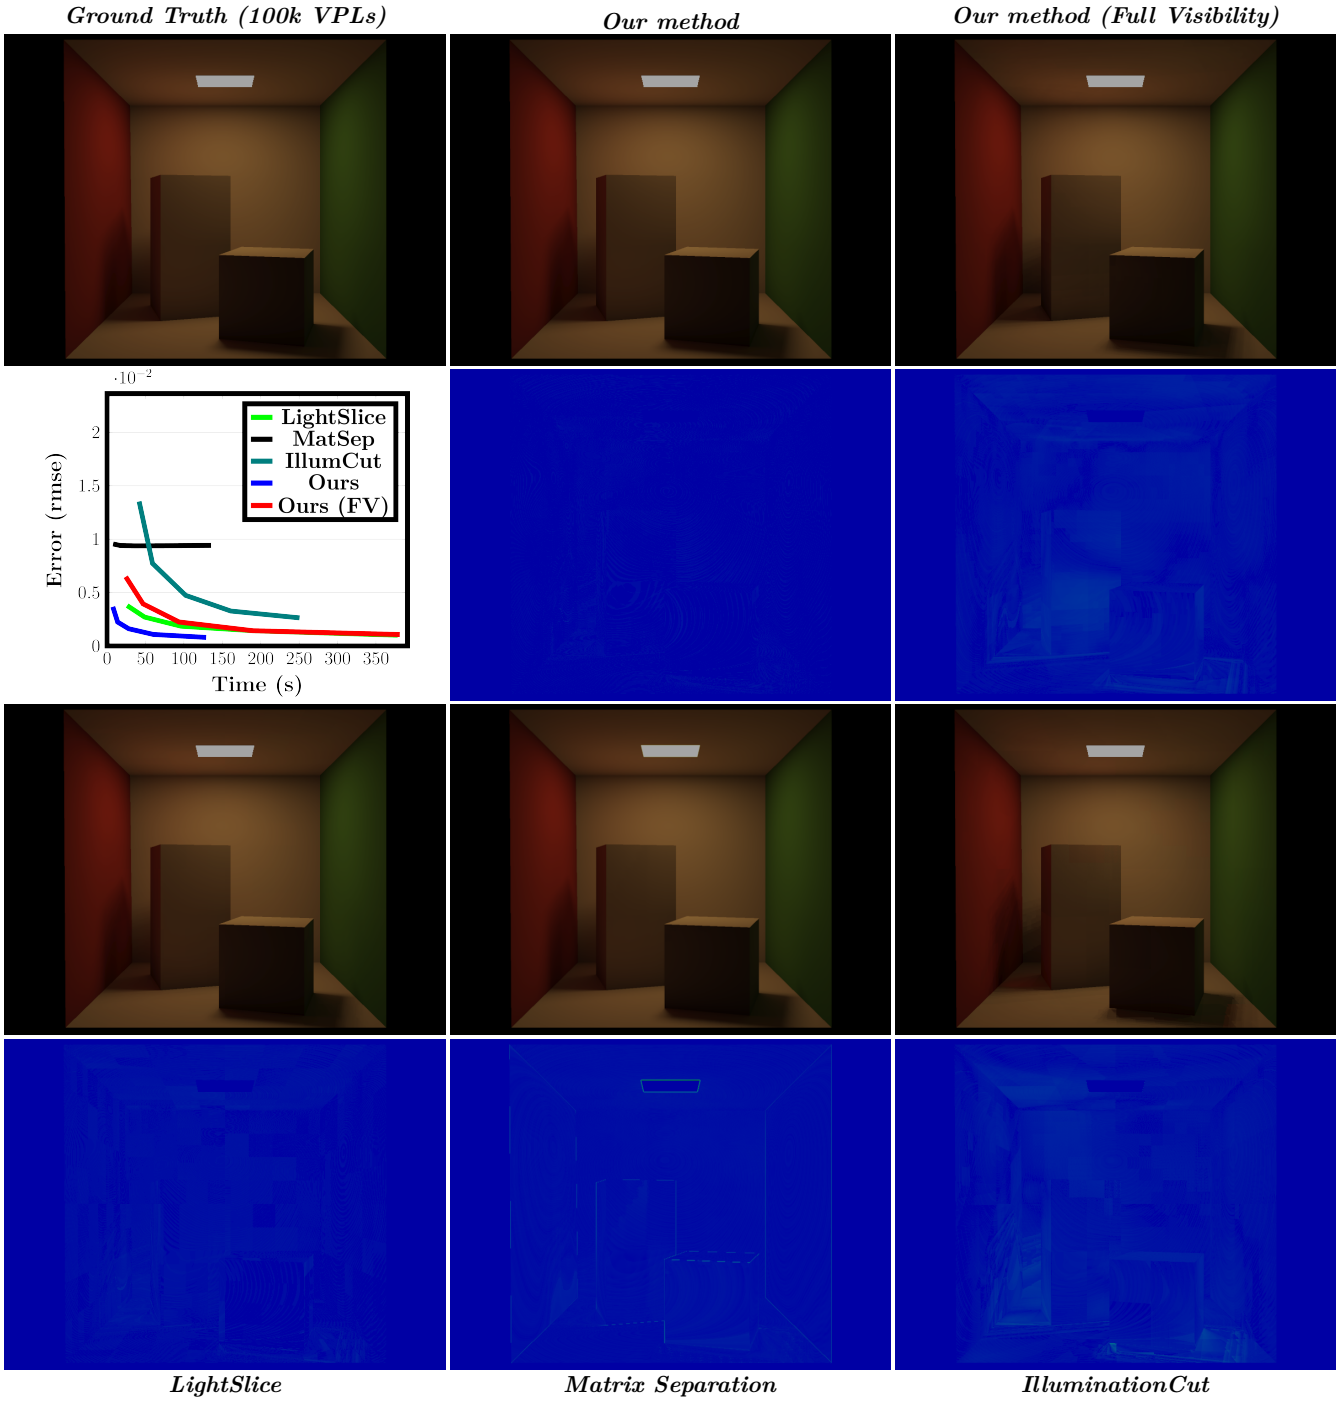

Figure 4: *Equal time renders (~30s) and time vs error plot for Cornell box scene*

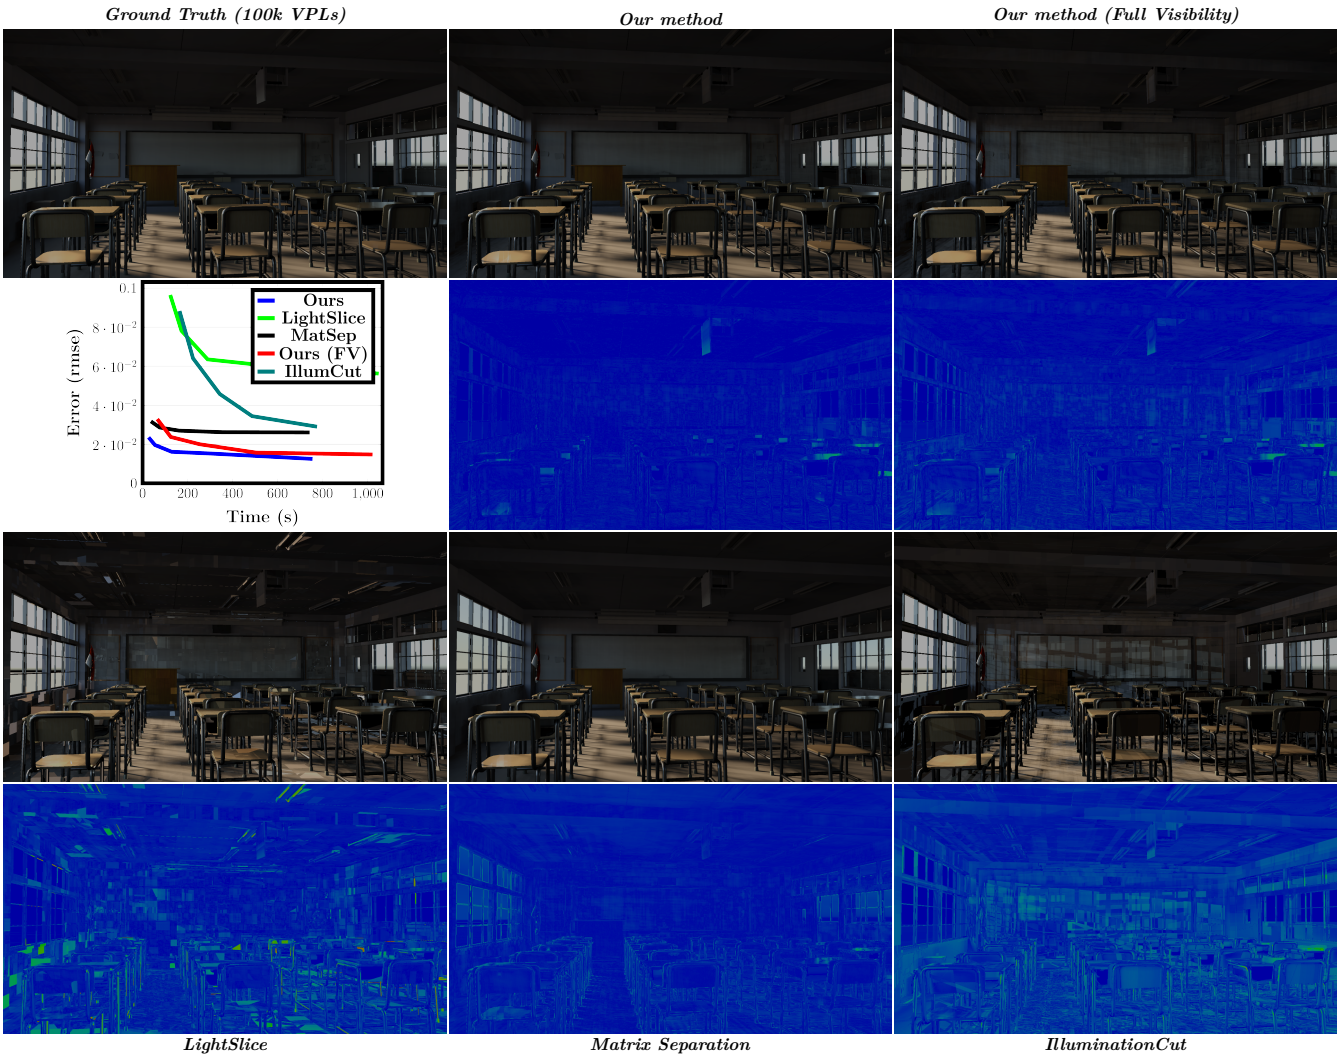

Figure 5: *Equal time renders (~55s) and time vs error plot for classroom scene*

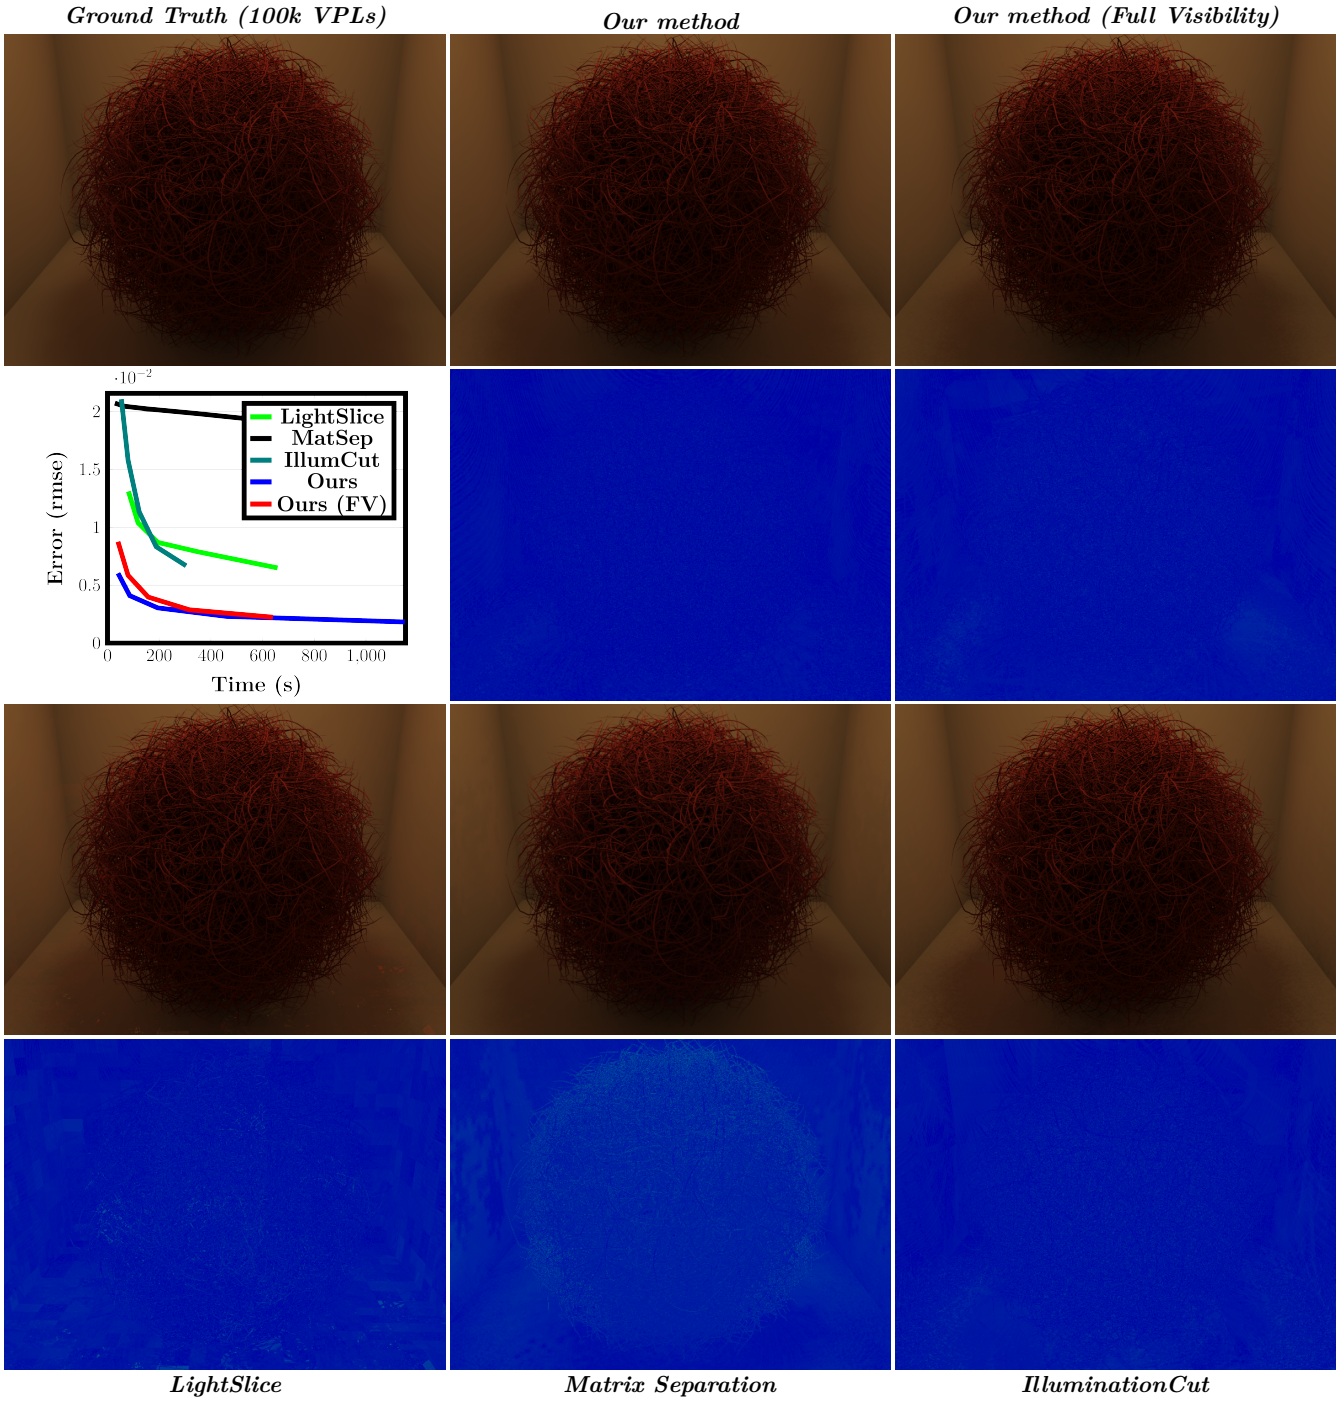

Figure 6: *Equal time renders (~80s) and time vs error plot for hairball scene*

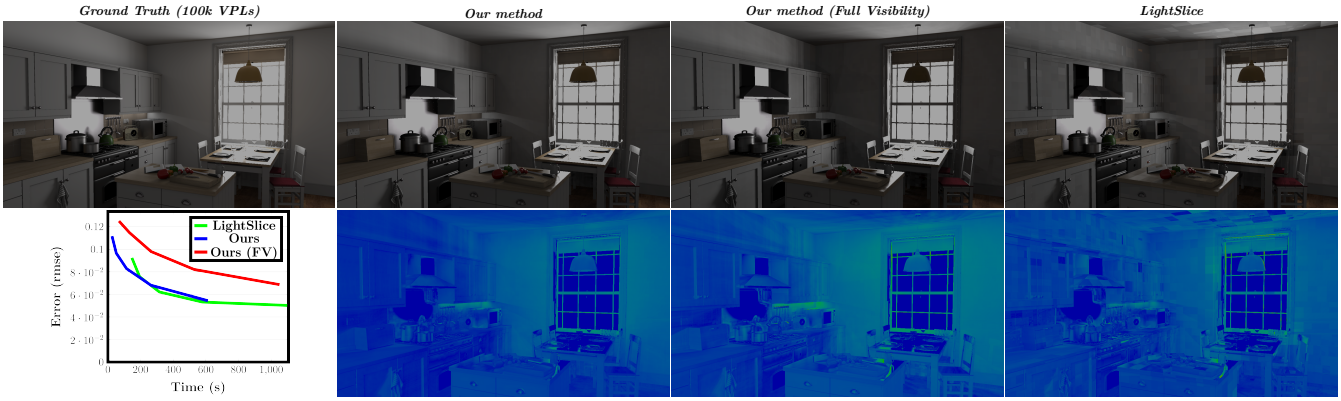

Figure 7: *Equal time renders (~60s) and time vs error for kitchen scene*

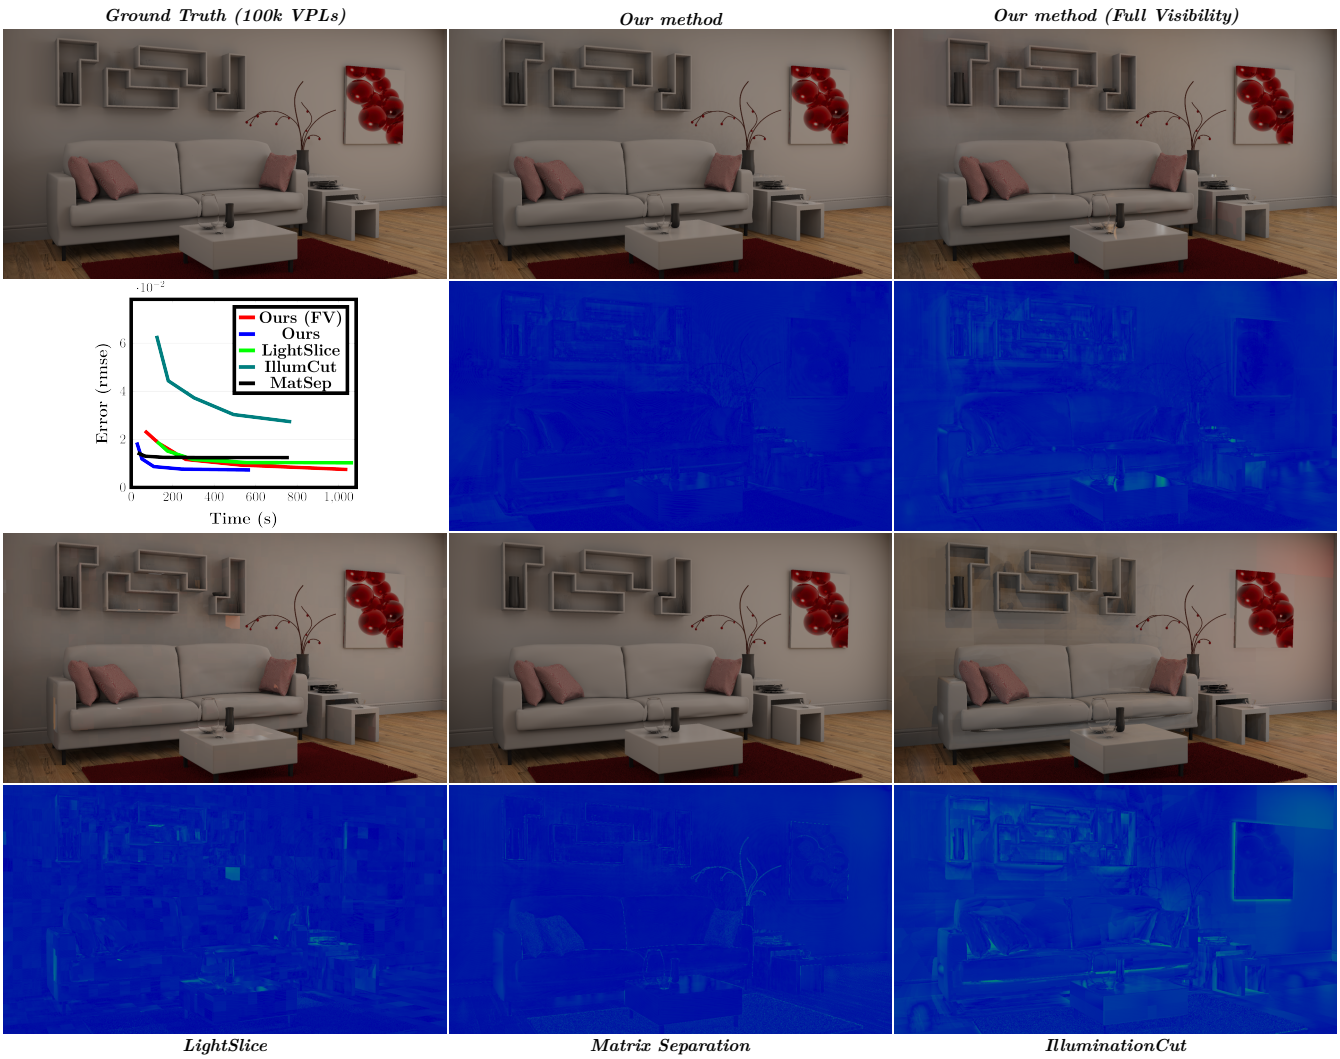

Figure 8: *Equal time renders (~110s) and time vs error for living room scene*

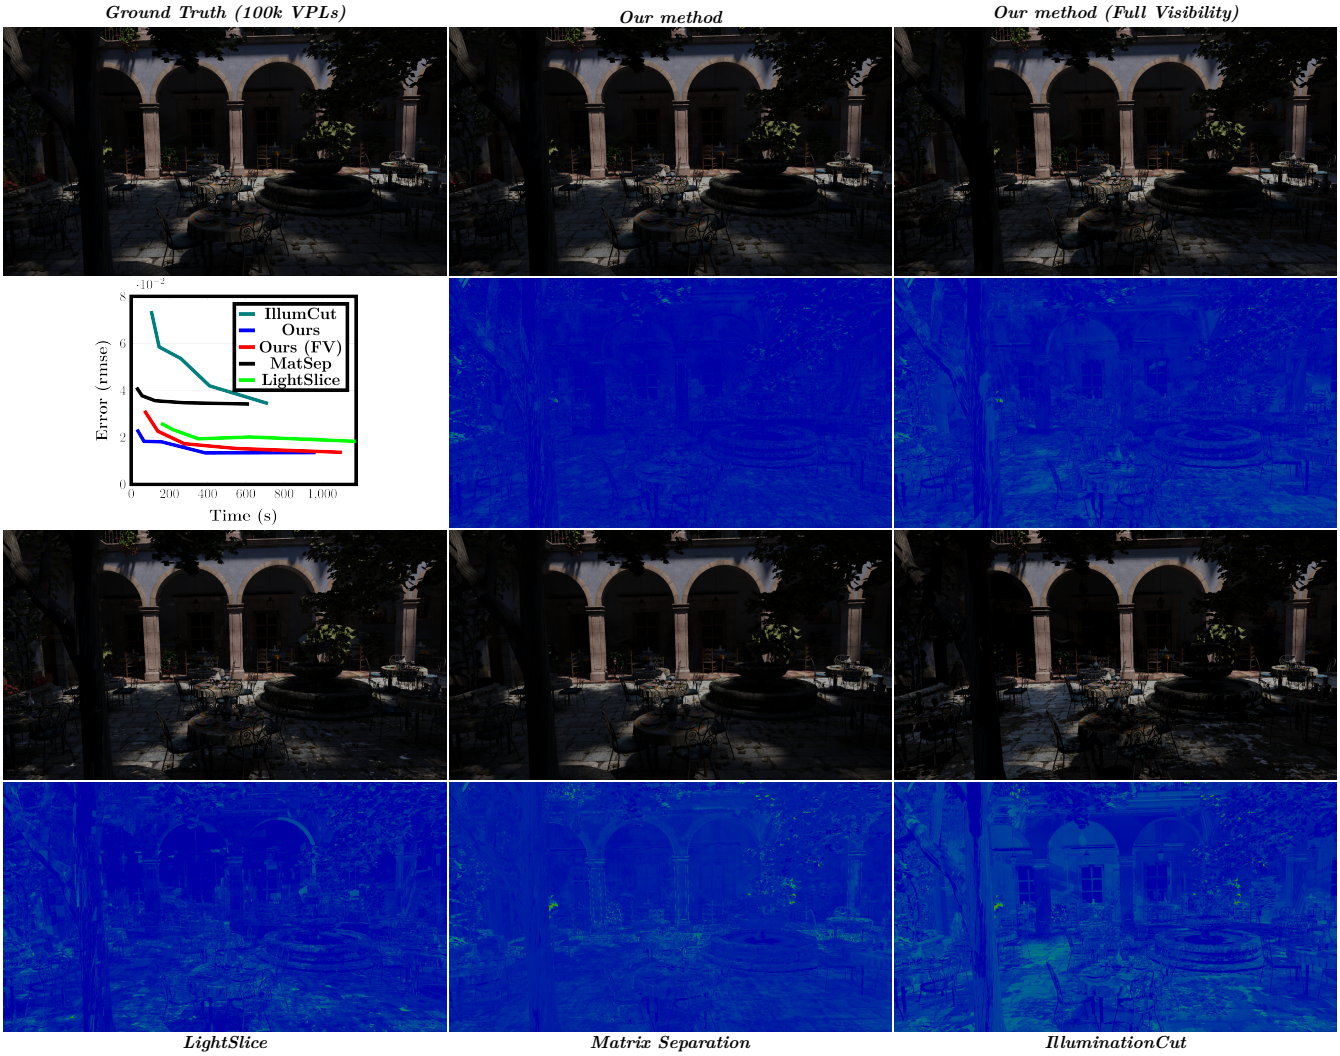

Figure 9: *Equal time renders (~125s) and time vs error plot for San Miguel scene*

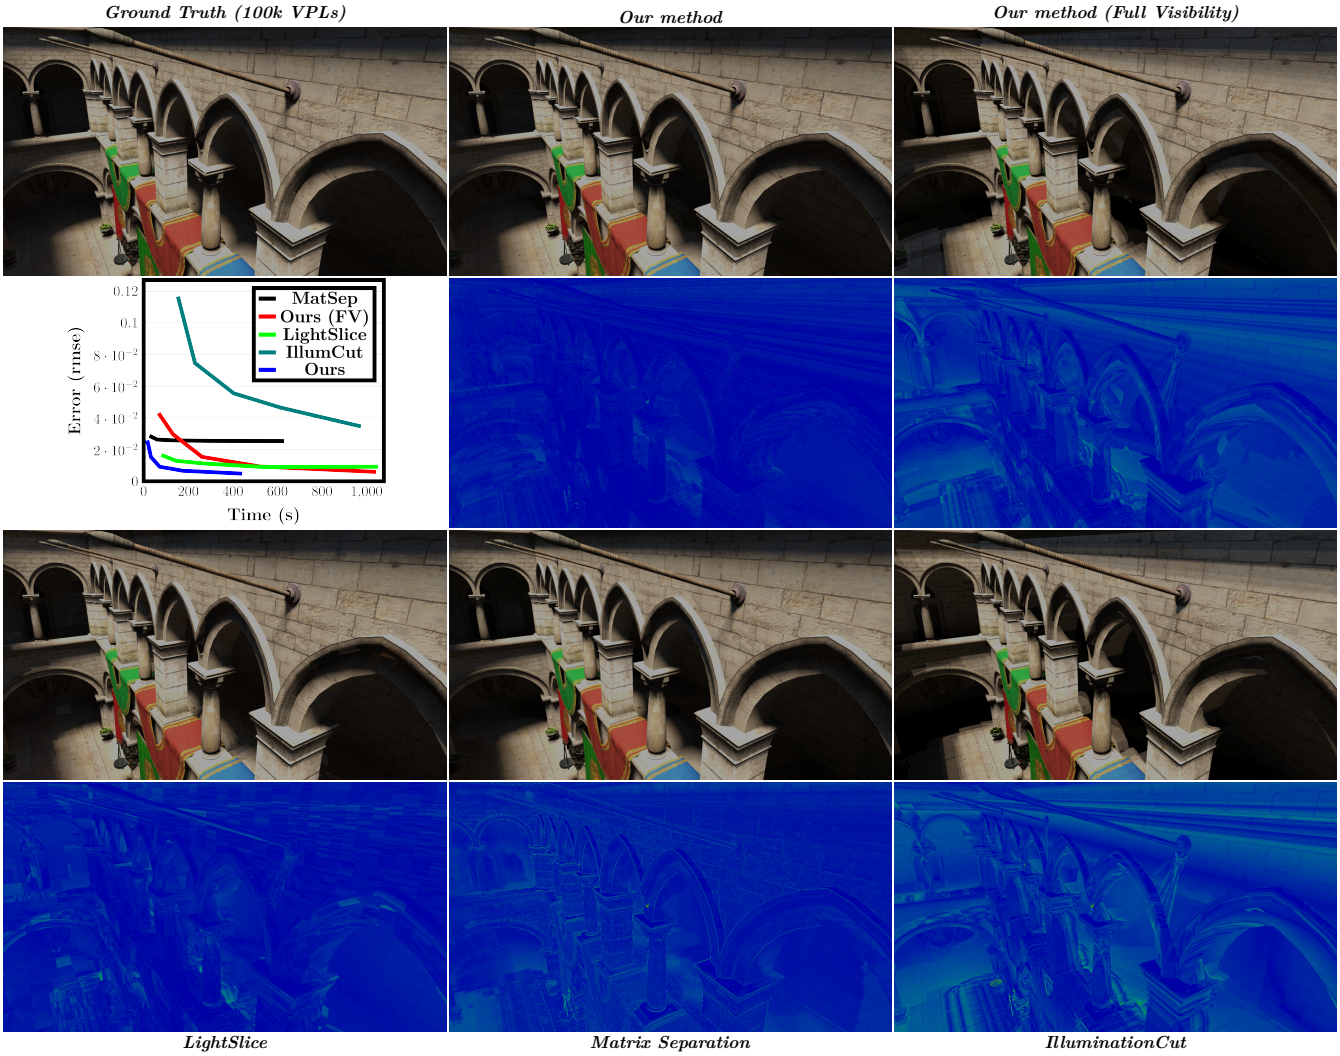

Figure 10: *Equal time renders (~140s) and time vs error plot for Sponza scene*

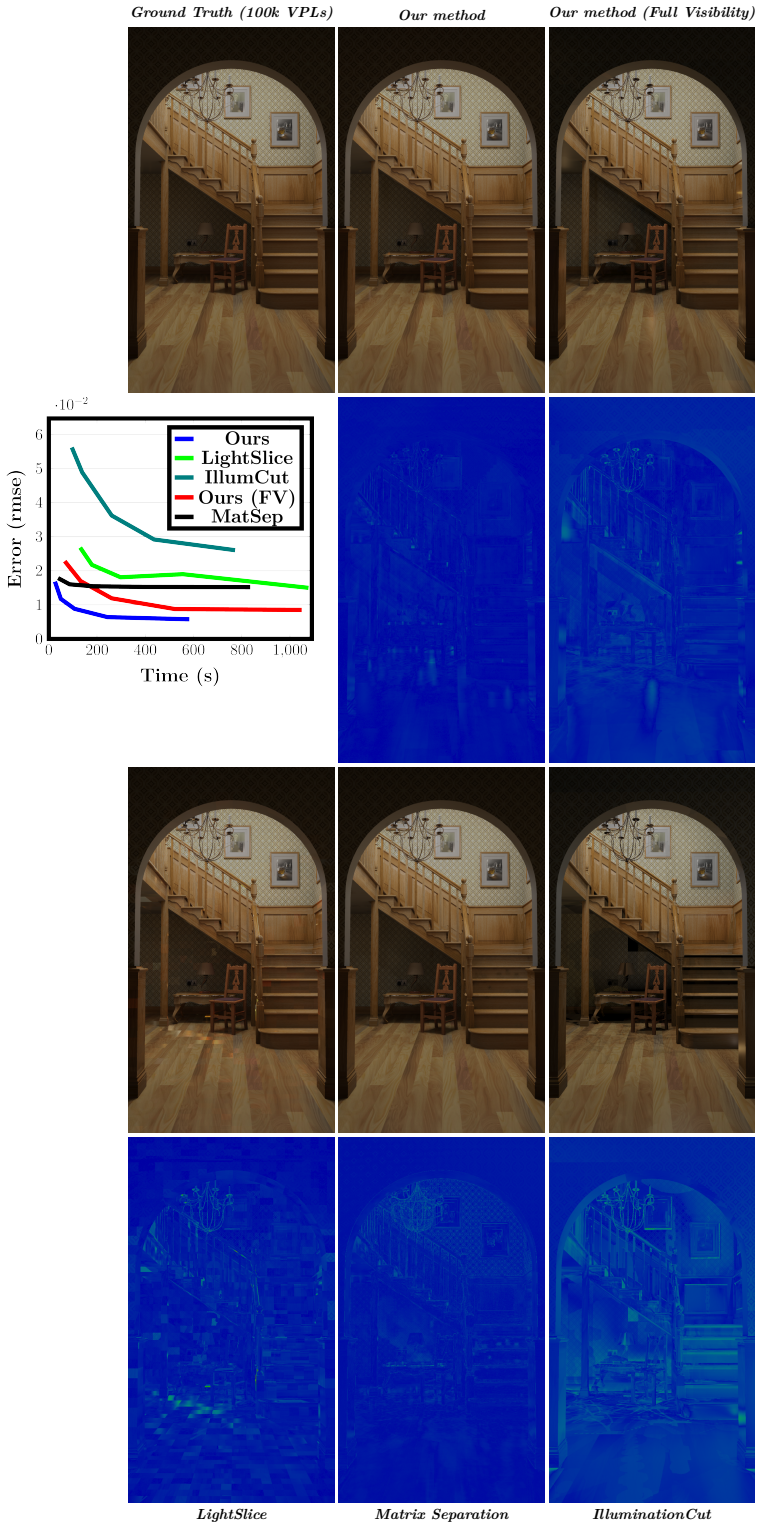

Figure 11: *Equal time renders (110s) and time vs error plot for staircase scene*

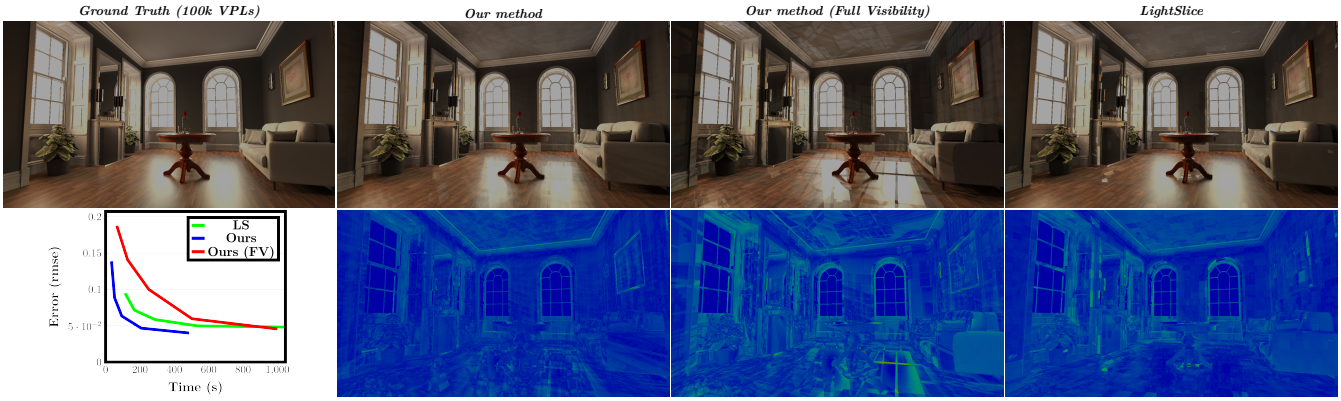

Figure 12: *Equal time renders (~220s) and time vs error for Grey & White Room scene*

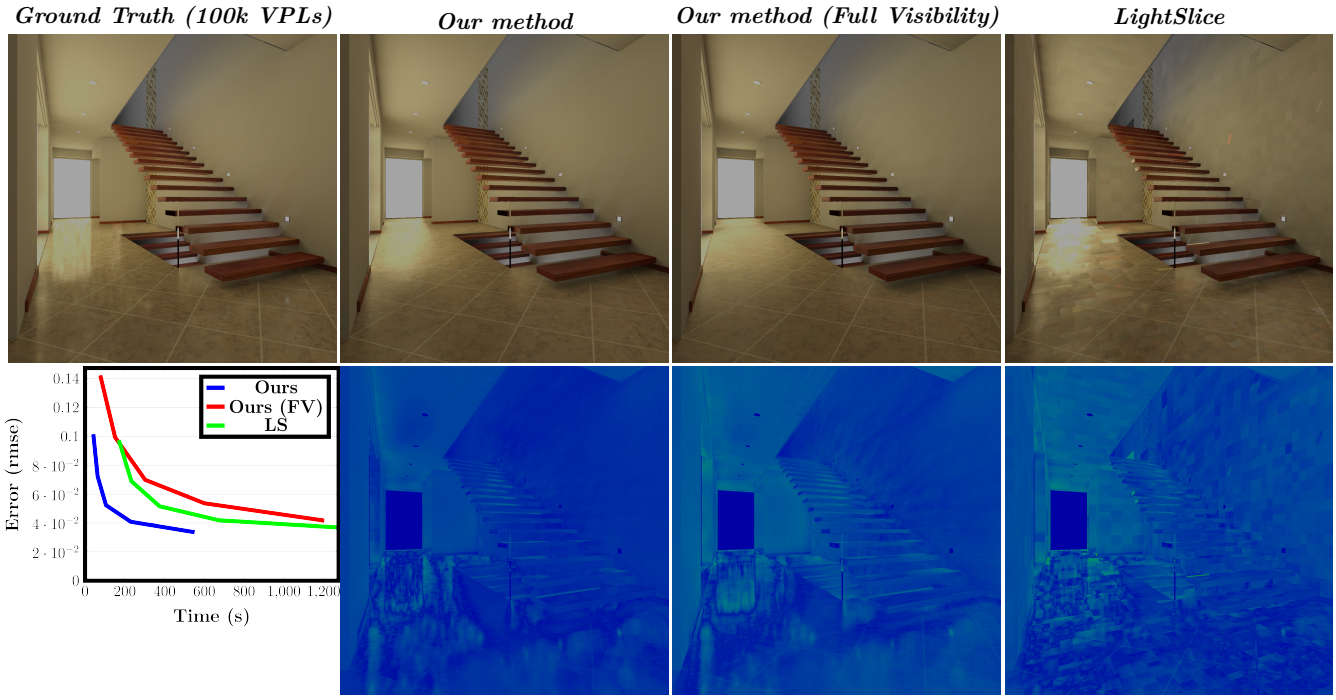

Figure 13: *Equal time renders (~230s) and time vs error for Modern Hall scene*

		San Miguel 1280 x 720 @ 4, 100k VPLs			Classroom 1280 x 720 @ 4, 100k VPLs			Kitchen 1280 x 720 @ 4, 100k VPLs		
		Time	RMSE	Samples	Time	RMSE	Samples	Time	RMSE	Samples
AMC	α=10% cols=1000	88.59s	0.0254	15.48%	71.60s	0.0242	16.98%	80.43s	0.0908	16.04%
	α=15% cols=1000	179.22s	0.0274	20.4%	201.55s	0.0226	22.57%	192.46s	0.0842	21.22%
	α=20% cols=1000	581.74s	0.0197	25.8%	805.46s	0.0209	27.49%	632.03s	0.0886	26.33%
	α=10% cols=2000	158.44	0.034	13.8%	146.36s	0.024	15.48%	153.79s	0.0762	14.26%
	α=10% cols=3000	230.08s	0.0297	13.18%	207.67s	0.0241	14.38%	223.68s	0.0602	13.79%
BAMC	α=10% cols=1000	64.28s	0.0252	17.25%	54.26s	0.0218	19.36%	53.05s	0.0865	17.22%
	α=15% cols=1000	81.77s	0.191	22.46%	69.3s	0.022	24.6%	68.17s	0.089	22.37%
	α=20% cols=1000	96.65s	0.0195	27.24%	81.24s	0.0197	29.03%	80.43s	0.0883	27.08%
	α=10% cols=2000	122.27s	0.0226	16.33%	105.74s	0.0211	18.96%	97.69s	0.079	15.92%
	α=10% cols=3000	178.53s	0.0253	15.7%	152.80s	0.0197	18.25%	142.86s	0.0597	15.44%
BAMC (GE)	α=10% cols=1000	156.88s	0.0199	18.63%	167.59s	0.02	21.72%	99.58s	0.091	19.55%
	α=15% cols=1000	171.21s	0.019	23.18%	175.37s	0.022	26%	106.12s	0.089	23.57%
	α=20% cols=1000	183.38s	0.018	27.69%	183.49s	0.0197	30.28%	112.81s	0.091	27.85%
	α=10% cols=2000	530.13s	0.023	18.2%	618.32s	0.03	21.8%	313.54s	0.075	20.41%
	α=10% cols=3000	1031.76s	0.032	17.75%	1039.85s	0.045	21.27%	670.02s	0.066	21.06%

Table 6: *Boolean Adaptive Matrix Completion statistics when compared to the original algorithm and the Gaussian elimination variant in the Classroom, San Miguel, and Kitchen scenes.*

Figure 14: *Our method used with LightSlice at 49.3% samples (left), and with brute force at 19.9% samples (right).*

Classroom				
	Time	RMSE	Samples	
Adaptive		412.93s	0.0142	25.02%
Static	α=10%	241.51s	0.0205	17.6%
	α=40%	659.26s	0.0139	46.29%

Table 7: *Statistics when comparing adaptive row-samplerates to static row sampling in the classroom scene.*

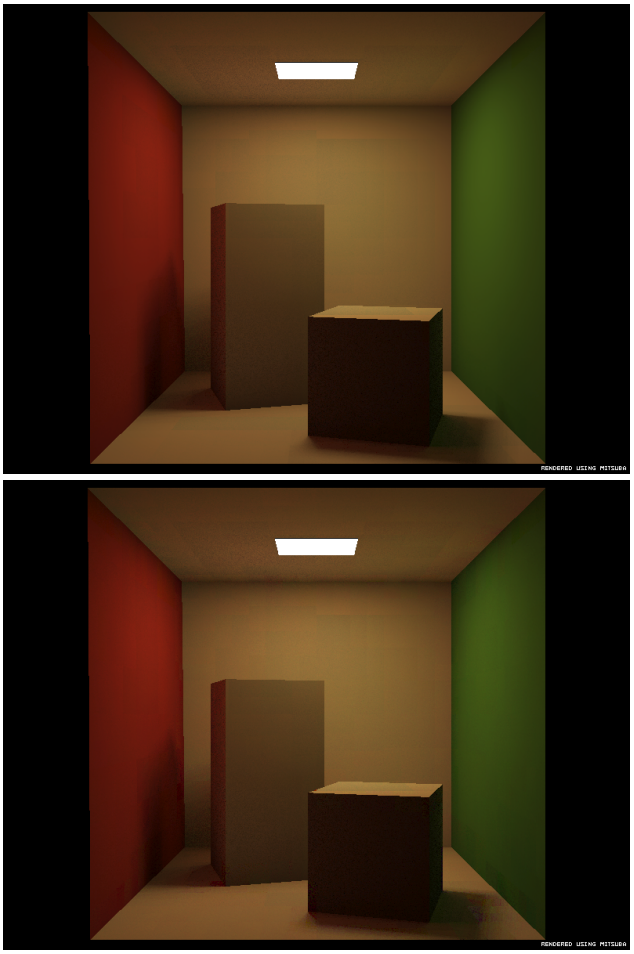

Figure 15: *Completing the visibility matrix with Singular Value Thresholding with (top) and without truncated svd (bottom) with 1000 clusters.*