

Understanding the Role of Phase Function in Translucent Appearance: Supplementary Material

IOANNIS GKIOULEKAS

Harvard School of Engineering and Applied Sciences
BEI XIAO

Massachusetts Institute of Technology

SHUANG ZHAO

Cornell University

EDWARD H. ADELSON

Massachusetts Institute of Technology

TODD ZICKLER

Harvard School of Engineering and Applied Sciences
and

KAVITA BALA

Cornell University

In this supplementary material we discuss details related to the main paper. We group these results into sections corresponding to the sections of the same title in the main paper.

1. CHARACTERIZING TRANSLUCENT APPEARANCE

In Figure 1, we show an example image of the scene “Lucy + Campus” rendered with phase function of average cosine value $\bar{C} > 0.8$ that we removed from the set of phase functions we use in all our experiments.

Authors’ addresses: I. Gkioulekas (corresponding author), T. Zickler, Harvard School of Engineering and Applied Sciences, Cambridge, MA 02138; email: {igkiou,zickler}@seas.harvard.edu. B. Xiao, Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139; email: beixiao@mit.edu. E.H. Adelson, CSAIL, Massachusetts Institute of Technology, Cambridge, MA 02139; email: adelson@csail.mit.edu. S. Zhao, K. Bala, Computer Science Department, Cornell University, Ithaca, NY 14853, {szhao,kb}@cs.cornell.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0730-0301/YYYY/13-ARTXXX \$10.00

DOI 10.1145/XXXXXXX.YYYYYY

<http://doi.acm.org/10.1145/XXXXXXX.YYYYYY>

2. COMPUTATIONAL EXPERIMENT

2.1 Embedding for Laboratory Setup

In Figure 2, we show the embedding produced using the various image difference metrics, for the “laboratory” setup with laser illumination (Figure 3 of the main paper). We observe that, for this scene, the cubic root metric produces the embedding most consistent with the other scenes, whereas the embedding produced by the L_2 -norm is significantly different from the rest.

2.2 Embeddings Using the L_1 -norm Metric

In Figure 3, we show the two dimensional embeddings produced using the L_1 -norm metric for all the nine scenes. In Tables I-III, we report numerical values for each of the three embedding similarity measures and all the scene pairs, for embeddings produced with the L_1 -norm metric.

2.3 Embeddings Using the L_2 -norm Metric

In Figure 4, we show the two dimensional embeddings produced using the L_2 -norm metric for all of the nine scenes. In Tables IV-VI, we report numerical values for each of the three embedding similarity measures and all the scene pairs, for embeddings produced with the L_2 -norm metric

2.4 Embedding Using Mean-Opinion-Score Metric

In Figure 5, we show the two-dimensional embeddings produced when using the mean-opinion-score metric by Mantiuk et al. [2011] for all of the nine scenes, using only the 40 phase functions selected by the clustering approach we describe in Section 6 of the main paper. We use the implementation that the authors have made publicly available¹. In Figure 6, we compare with the corresponding embeddings produced using the cubic root metric. We observe that the geometric structure of the embeddings produced by the two metrics is very similar. For corresponding embeddings produced by the two different metrics, using a linear regression hypothesis test, we found that, for both coordinates, the hypothesis of linear relation between them is statistically significant at the 99% confidence level.

¹<http://hdrvdp.sourceforge.net/>




Fig. 1. Image of the scene “Lucy + Campus” rendered with phase functions of average cosine value $\bar{C} > 0.8$.

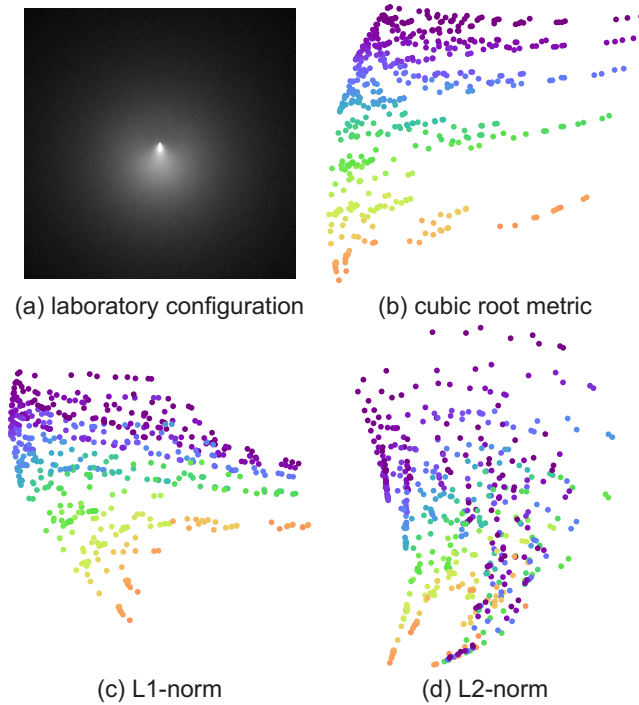


Fig. 2. Two-dimensional embeddings of images rendered with a representative set of phase functions sampled from the physical parameter space, for the “laboratory” setup. The embedded dots, one per rendered image, are colored according to the square of the average cosine of the phase function used to render them.

3. PSYCHOPHYSICAL EXPERIMENTS

3.1 Instructions for Psychophysical Experiment

Before taking the study, every observer was provided with a set of instructions, which were also explained to them by the author supervising the study. A flyer was used as a visual aid during this stage, which is provided separately as supplementary material as file “instructions.pdf”.

3.2 Clustering for Stimuli Selection

In Figures 7-9, we show all members of three representative clusters, out of the forty clusters produced by the clustering algorithm used for the stimulus image selection. We show these examples to demonstrate that the clusters includes materials that appear very similar, and can be represented well by an exemplar in a psychophysical experiment.

3.3 Visualization of Embeddings

In Figures 10-13, we show cropped versions of all 40 stimuli images used in each of the two psychophysical experiments, overlayed on the corresponding computational and perceptual embeddings.

3.4 Bootstrapping Analysis

As a way of assessing the stability of these embeddings, we performed a bootstrapping analysis: we created 5 perturbations of the “Lucy + Campus” data, by selecting for each of those a random subset of 7200 triplets (24% of the total number, using the 75.94% consistency we report in the main paper), and flipping the corresponding user-supplied ratings; we then used each of the perturbed data sets to learn a two-dimensional embedding, in the same way as described above. Figure 14 compares the embedding using the original data with the embeddings learned from perturbed data. We observe that, despite the differences, the overall structure and ordering of the images in the embedding remain consistent.

4. ANALYSIS

4.1 Generalized Procrustes Analysis

We use generalized Procrustes analysis, described in Appendix B of the main paper, to compute the full Procrustes mean of the embeddings for the nine scenes produced using the cubic root metric (Figure 4 of the main paper). The mean embedding is shown in Figure 15, where it is also compared with the embedding for “Lucy + Campus”. In Tables VII and VIII, we also report the results of the same correlation analysis as in Section 7 of the main paper for parameterizing the two dimensions of the embedding, as applied to the mean embedding.

4.2 Phase Function Distance Metric

In Figure 16, we show the two-dimensional weight function learned in the metric learning experiments of Section 7.2 of the main paper. The function is shown as a matrix, whose dimensions correspond to the axis θ_1 and θ_2 . The diagonal of this matrix is shown in the polar plot of Figure 10(a) of the main paper.

5. DISCUSSION AND APPLICATIONS

5.1 Material Design

In Figures 17-22, we show high resolution versions of the renderings in Figures 1(c) and 12 of the main paper.

5.2 Perceptually Uniform Axes

We show interpolation sequences of images in the form of video sequences, in separate uploaded files: “lucy_linear.mov”, “buddha_linear.mov”, “lucy_quadratic.mov”, and “buddha_quadratic.mov”, containing respectively “Lucy + Campus” linear sampling, “Buddha + Campus” linear sampling, “Lucy + Campus” quadratic sampling, and “Buddha + Campus” quadratic sampling sequences. We observe that the sequences using quadratic sampling have more uniform transitions across the $g \in [0, 0.9]$ interval than those with linear sampling.

REFERENCES

MANTIUK, R., KIM, K., REMPEL, A., AND HEIDRICH, W. 2011. HDR-VDP-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics (TOG)* 30, 4, 40.

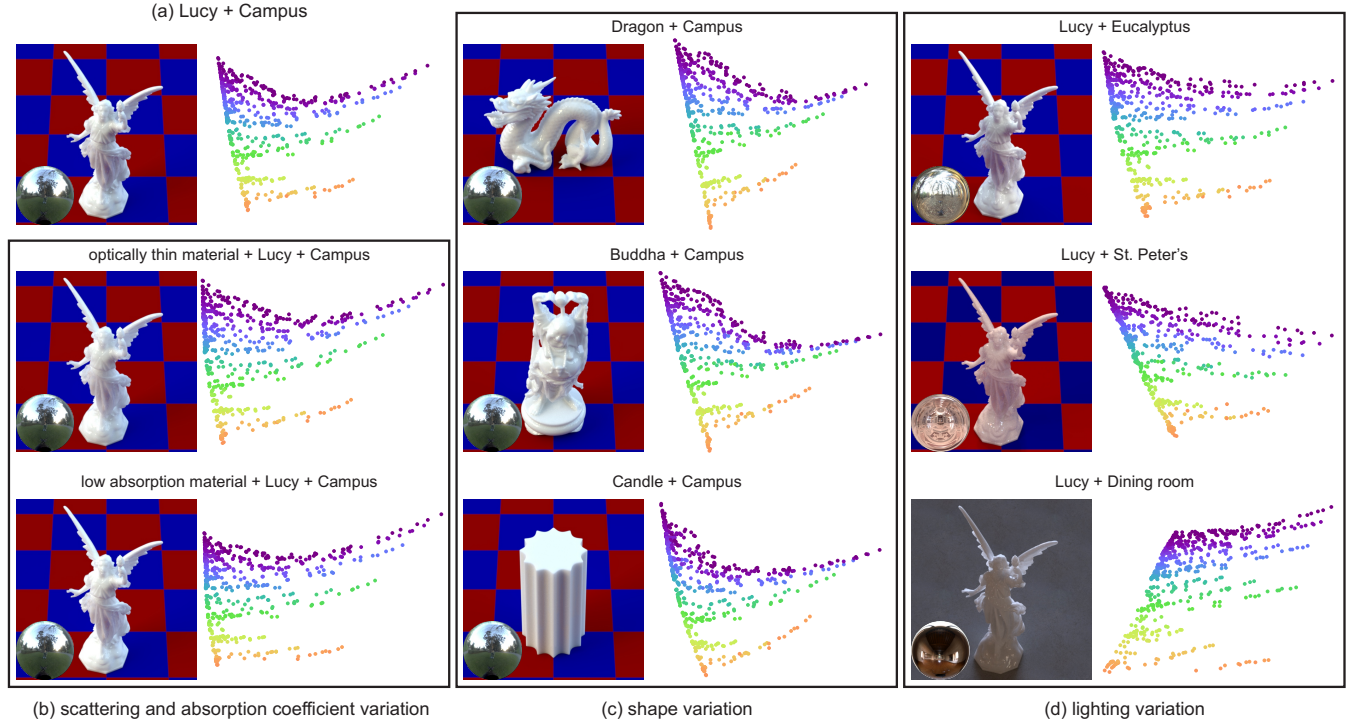


Fig. 3. Two-dimensional embeddings of images rendered with a representative set of phase functions sampled from the physical parameter space, produced using the L_1 -norm image metric. The embedded dots, one per rendered image for a total of 753 dots per embedding, are colored according to the square of the average cosine of the phase function used to render them. To the left of each embedding is shown the scene it corresponds to. Scenes are grouped in terms of the type of change with reference to the “Lucy + Campus” scene shown in (a). Columns left: variation of scattering and absorption coefficients, middle: shape variation, right: lighting variation.

Table I. Pairwise Procrustes distances for the embeddings of Figure 3.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) “Lucy + Campus”	0	0.021	0.035	0.020	0.006	0.018	0.116	0.138	0.415
(2) “Lucy + Dragon”		0	0.006	0.010	0.043	0.066	0.191	0.169	0.540
(3) “Lucy + Buddha”			0	0.010	0.064	0.096	0.214	0.185	0.588
(4) “Lucy + Candle”				0	0.045	0.072	0.171	0.152	0.538
(5) “optically thin material”					0	0.006	0.104	0.151	0.357
(6) “low absorption material”						0	0.106	0.167	0.314
(7) “Lucy + Eucalyptus”							0	0.046	0.189
(8) “Lucy + St. Peter’s”								0	0.321
(9) “Lucy + Dining room”									0

Table II. Pairwise Pearson’s correlation coefficient for the vertical coordinate of the embeddings of Figure 3.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) “Lucy + Campus”	1	0.987	0.971	0.984	0.998	0.992	0.994	0.974	0.971
(2) “Lucy + Dragon”		1	0.994	0.982	0.985	0.966	0.994	0.989	0.953
(3) “Lucy + Buddha”			1	0.980	0.966	0.937	0.980	0.984	0.916
(4) “Lucy + Candle”				1	0.976	0.957	0.977	0.963	0.919
(5) “optically thin material”					1	0.994	0.992	0.970	0.976
(6) “low absorption material”						1	0.983	0.954	0.987
(7) “Lucy + Eucalyptus”							1	0.992	0.974
(8) “Lucy + St. Peter’s”								1	0.951
(9) “Lucy + Dining room”									1

Table III. Pairwise Pearson’s correlation coefficient for the horizontal coordinate of the embeddings of Figure 3.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) “Lucy + Campus”	1	0.997	0.995	0.995	0.997	0.993	0.994	0.955	0.866
(2) “Lucy + Dragon”		1	0.998	0.998	0.990	0.983	0.993	0.973	0.831
(3) “Lucy + Buddha”			1	0.999	0.985	0.976	0.996	0.975	0.812
(4) “Lucy + Candle”				1	0.985	0.976	0.995	0.977	0.811
(5) “optically thin material”					1	0.998	0.987	0.933	0.898
(6) “low absorption material”						1	0.978	0.919	0.915
(7) “Lucy + Eucalyptus”							1	0.966	0.831
(8) “Lucy + St. Peter’s”								1	0.695
(9) “Lucy + Dining room”									1

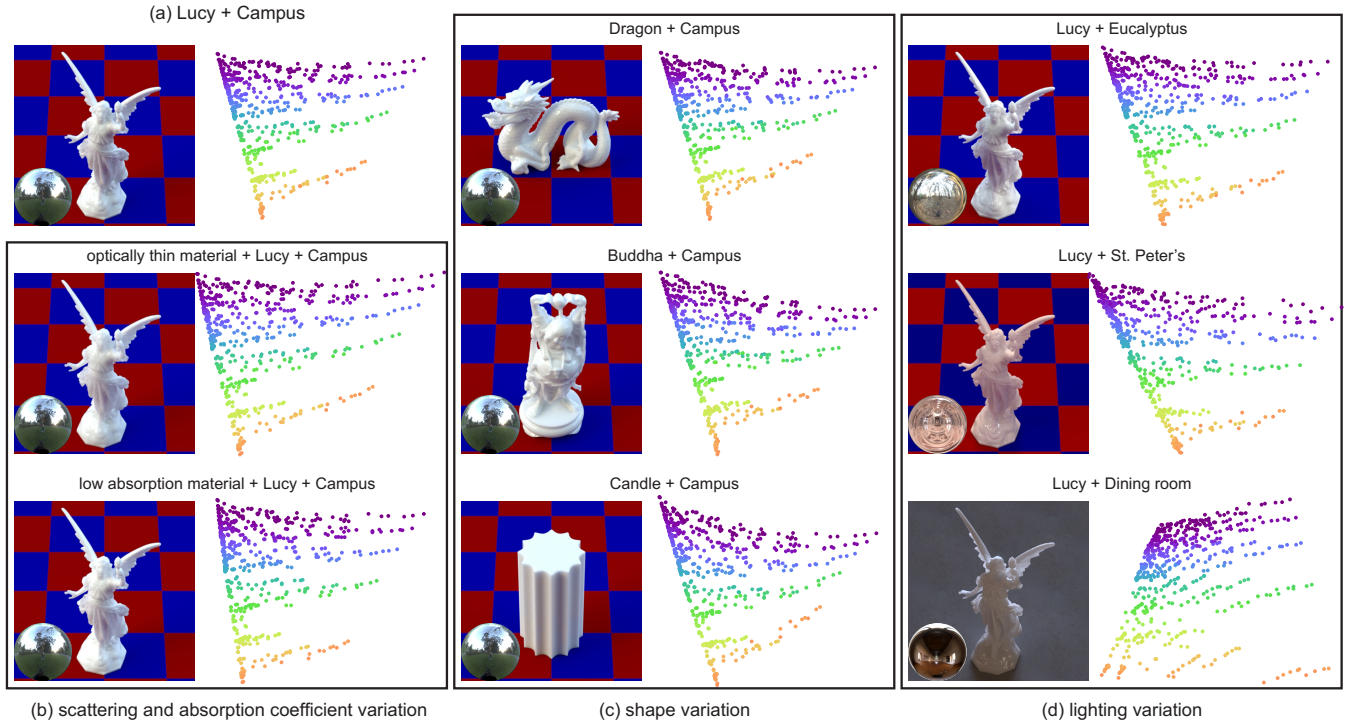


Fig. 4. Two-dimensional embeddings of images rendered with a representative set of phase functions sampled from the physical parameter space, produced using the L_2 -norm image metric. The embedded dots, one per rendered image for a total of 753 dots per embedding, are colored according to the square of the average cosine of the phase function used to render them. To the left of each embedding is shown the scene it corresponds to. Scenes are grouped in terms of the type of change with reference to the “Lucy + Campus” scene shown in (a). Columns left: variation of scattering and absorption coefficients, middle: shape variation, right: lighting variation.

Table IV. Pairwise Procrustes distances for the embeddings of Figure 4.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) “Lucy + Campus”	0	0.006	0.017	0.024	0.005	0.012	0.103	0.101	0.531
(2) “Lucy + Dragon”		0	0.005	0.009	0.020	0.029	0.146	0.124	0.604
(3) “Lucy + Buddha”			0	0.002	0.039	0.054	0.171	0.136	0.648
(4) “Lucy + Candle”				0	0.048	0.064	0.188	0.147	0.669
(5) “optically thin material”					0	0.004	0.084	0.104	0.473
(6) “low absorption material”						0	0.091	0.117	0.457
(7) “Lucy + Eucalyptus”							0	0.033	0.264
(8) “Lucy + St. Peter’s”								0	0.382
(9) “Lucy + Dining room”									0

Table V. Pairwise Pearson's correlation coefficient for the vertical coordinate of the embeddings of Figure 4.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) "Lucy + Campus"	1	0.999	0.993	0.991	0.999	0.997	0.999	0.991	0.985
(2) "Lucy + Dragon"		1	0.997	0.993	0.998	0.995	0.998	0.992	0.978
(3) "Lucy + Buddha"			1	0.995	0.989	0.992	0.995	0.994	0.965
(4) "Lucy + Candle"				1	0.986	0.992	0.994	0.994	0.958
(5) "optically thin material"					1	0.993	0.996	0.985	0.984
(6) "low absorption material"						1	0.999	0.996	0.986
(7) "Lucy + Eucalyptus"							1	0.996	0.983
(8) "Lucy + St. Peter's"								1	0.972
(9) "Lucy + Dining room"									1

Table VI. Pairwise Pearson's correlation coefficient for the horizontal coordinate of the embeddings of Figure 4.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) "Lucy + Campus"	1	0.999	0.998	0.997	0.999	0.992	0.993	0.966	0.850
(2) "Lucy + Dragon"		1	0.999	0.998	0.997	0.989	0.994	0.973	0.835
(3) "Lucy + Buddha"			1	0.999	0.993	0.982	0.997	0.979	0.825
(4) "Lucy + Candle"				1	0.992	0.980	0.996	0.981	0.823
(5) "optically thin material"					1	0.996	0.988	0.952	0.872
(6) "low absorption material"						1	0.973	0.931	0.876
(7) "Lucy + Eucalyptus"							1	0.98	0.832
(8) "Lucy + St. Peter's"								1	0.730
(9) "Lucy + Dining room"									1

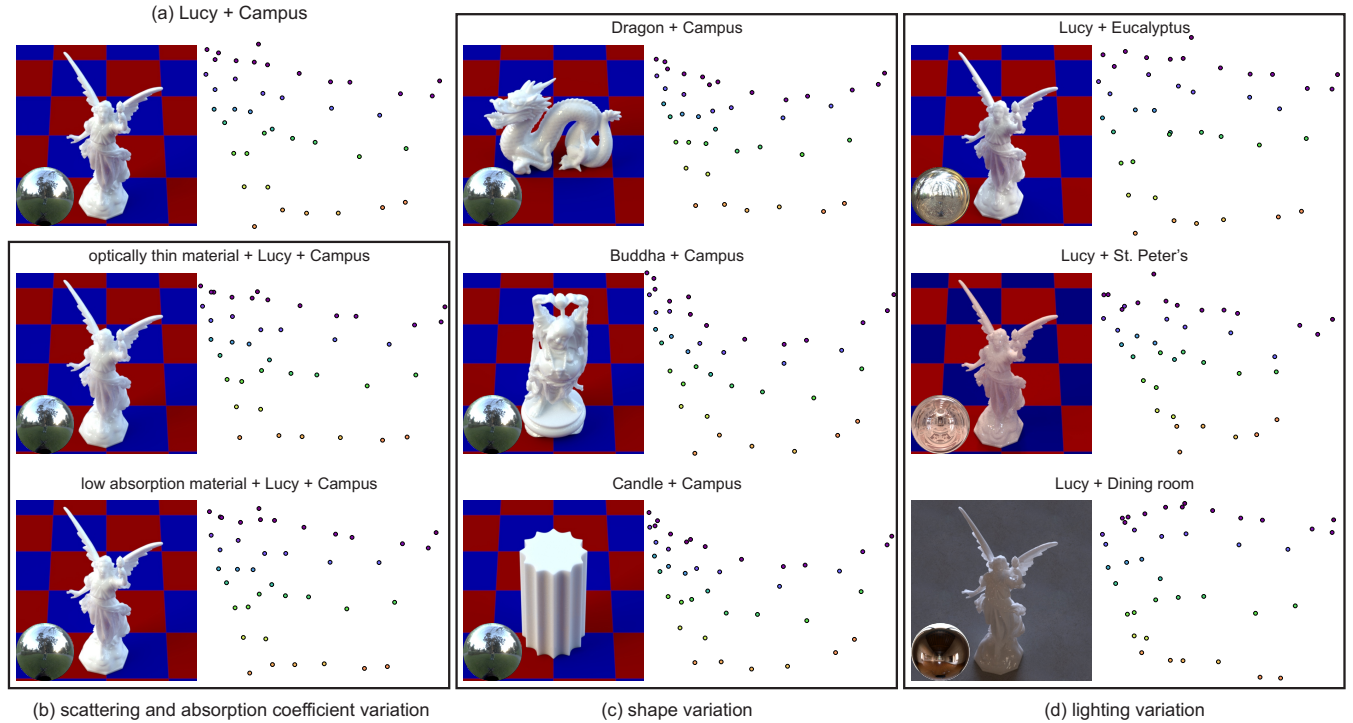


Fig. 5. Two-dimensional embeddings of images rendered with the 40 phase functions selected by the clustering algorithm, produced using the mean-opinion-score image metric. The embedded dots, one per rendered image for a total of 40 dots per embedding, are colored according to the square of the average cosine of the phase function used to render them. To the left of each embedding is shown the scene it corresponds to. Scenes are grouped in terms of the type of change with reference to the "Lucy + Campus" scene shown in (a). Columns left: variation of scattering and absorption coefficients, middle: shape variation, right: lighting variation.

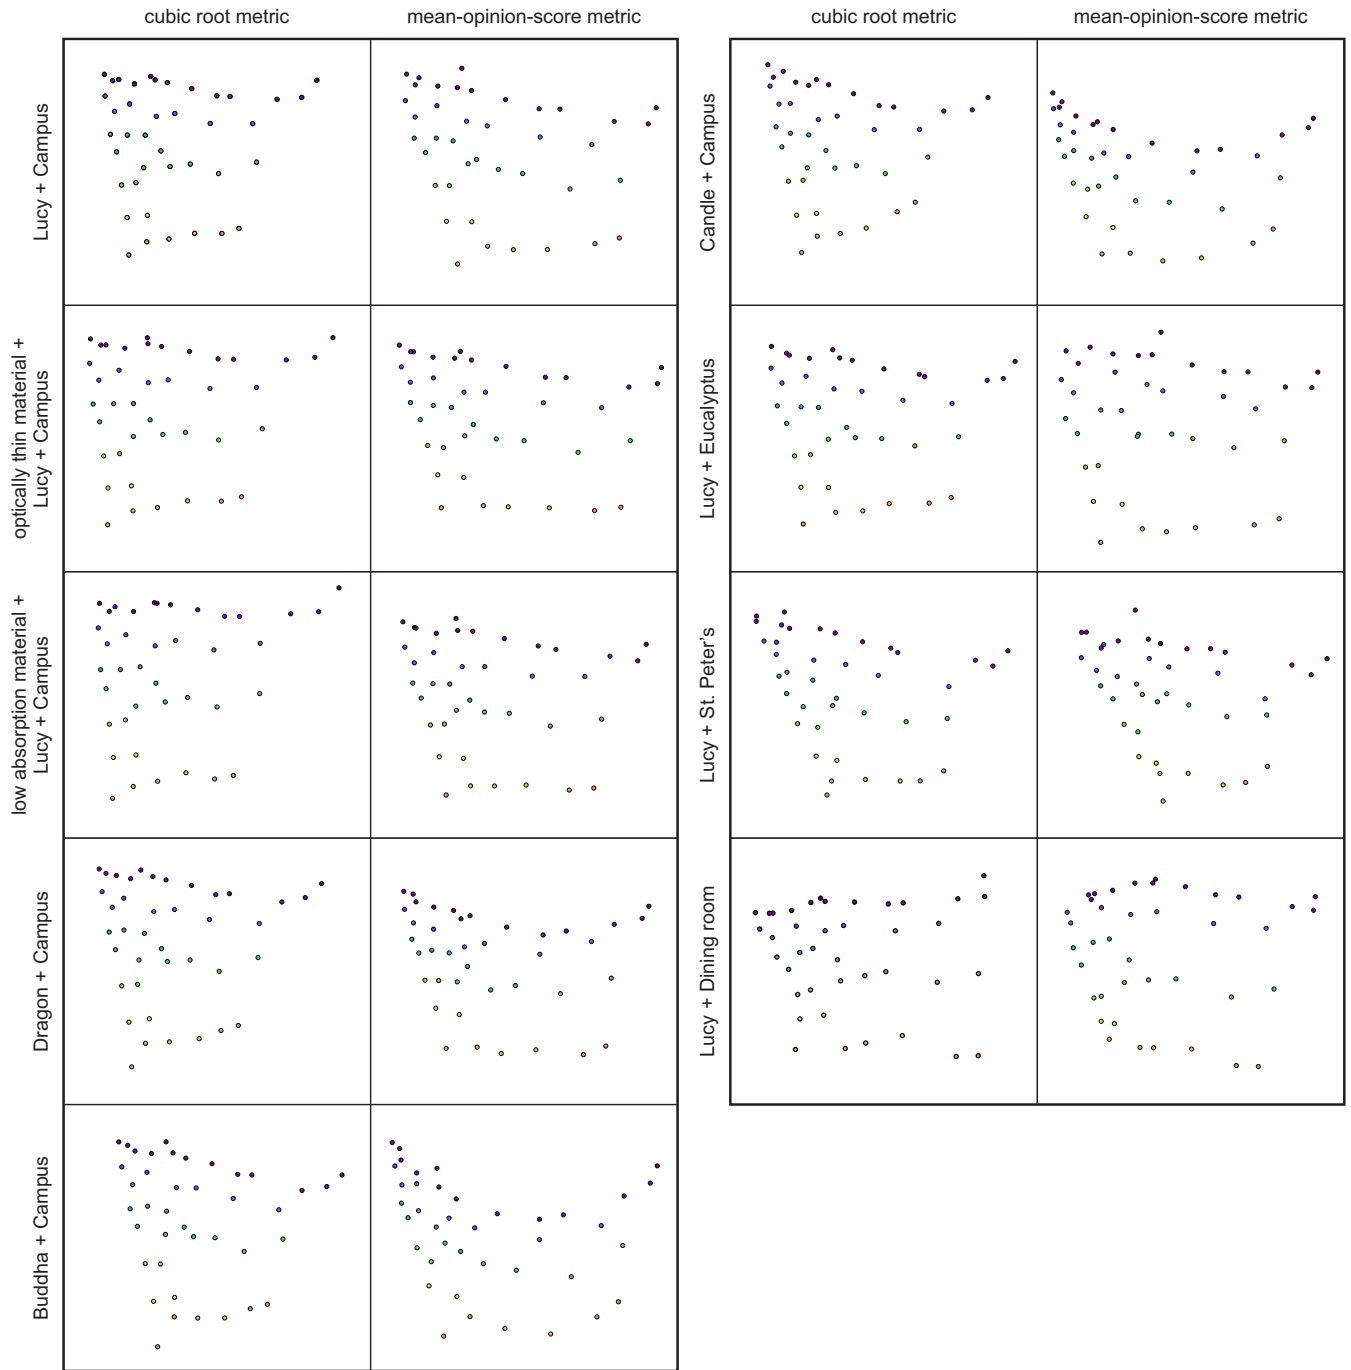


Fig. 6. Comparison of two-dimensional embeddings produced by the cubic root and mean-opinion-score metrics, for images rendered with the 40 phase functions selected by the clustering algorithm. The embedded dots, one per rendered image for a total of 40 dots per embedding, are colored according to the square of the average cosine of the phase function used to render them.




Fig. 7. Member images of representative cluster 1.




Fig. 8. Member images of representative cluster 2.




Fig. 9. Member images of representative cluster 3.



Fig. 10. Visualization of computational embedding for sidelighting experiment.




Fig. 11. Visualization of perceptual embedding for sidelighting experiment.




Fig. 12. Visualization of computational embedding for backlighting experiment.




Fig. 13. Visualization of perceptual embedding for backlighting experiment.




Fig. 14. Bootstrapping analysis: comparison of (a) embedding learned from data provided from the subjects; and (b)-(f) embeddings learned from five random perturbations of the data.

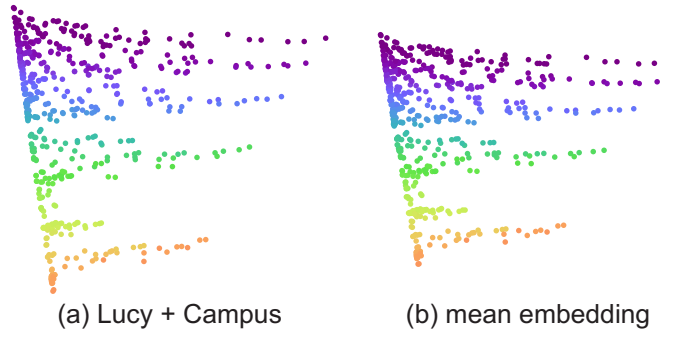


Fig. 15. Comparison between (a) the two-dimensional embedding for the scene “Lucy + campus”, produced using the cubic root metric; and (b) the full Procrustes mean of all the embeddings of Figure 4 of the main paper, as produced by generalized Procrustes analysis.

Table VII. Parameterization of the vertical dimension of the full Procrustes mean of the embeddings of Figure 4 of the main paper. (“Correlation” refers to the absolute value of Pearson’s correlation coefficient.)

scene	correlation with \bar{C}	correlation with \bar{C}^2	best power $a \in [1.5, 2.5]$	correlation with \bar{C}^a
full Procrustes mean	0.9619	0.9942	2.06	0.9943

Table VIII. Parameterization of the horizontal dimension of the full Procrustes mean of the embeddings of Figure 4 of the main paper. (“Correlation” refers to the absolute value of Pearson’s correlation coefficient.)

scene	correlation with $1/\sqrt{1 - M_C}$
full Procrustes mean	0.9109




Fig. 16. Two-dimensional learned weight function for comparing phase functions, visualized as a matrix. Zooming in shows the high values of the diagonal, which is plotted separately in Figure 10(a) of the main paper.



Fig. 17. High resolution version of the marble rendering in Figure 1(c)-left of the main paper.



Fig. 18. High resolution version of the rendering in Figure 1(c)-middle of the main paper.



Fig. 19. High resolution version of the white jade rendering in Figure 1(c)-right of the main paper.



Fig. 20. High resolution version of the marble rendering in Figure 12-left of the main paper.



Fig. 21. High resolution version of the rendering in Figure 12-middle of the main paper.



Fig. 22. High resolution version of the white jade rendering in Figure 12-right of the main paper.