

Impact of Shape on Apparent Translucency Differences

Davit Gigilashvili*, Philipp Urban*, Jean-Baptiste Thomas*, Jon Yngve Hardeberg*, Marius Pedersen*

* Department of Computer Science, Norwegian University of Science and Technology; Gjøvik, Norway

† Fraunhofer Institute for Computer Graphics Research IGD; Darmstadt, Germany

Abstract

Translucency is one of the major appearance attributes. Apparent translucency is impacted by various factors including object shape and geometry. Despite general proposals that object shape and geometry have a significant effect on apparent translucency, no quantification has been made so far. Quantifying and modeling the impact of geometry, as well as comprehensive understanding of the translucency perception process, are a point of not only academic, but also industrial interest with 3D printing as an example among many. We hypothesize that a presence of thin areas in the object facilitates material translucency estimation and changes in material properties have larger impact on apparent translucency of the objects with thin areas. Computer-generated images of objects with various geometry and thickness have been used for a psychophysical experiment in order to quantify apparent translucency difference between objects while varying material absorption and scattering properties. Finally, absorption and scattering difference thresholds where the human visual system starts perceiving translucency difference need to be identified and its consistency needs to be analyzed across different shapes and geometries.

Introduction and Background

Apparent translucency is caused by stimuli emitted by an object possessing some degree of subsurface light transport. It depends on extrinsic and intrinsic factors, whereas extrinsic factors are the viewing conditions, in particular the illumination direction. Xiao *et al.* [1] showed that the magnitude of translucency of various materials is perceived differently if viewed under back-lit, side-lit or front-lit conditions. Only for the class of relatively simple materials with an isotropic scattering function (i.e. light scattering within the material is uniform in all directions), the illuminating direction plays a rather small role and the materials can be attributed to be *translucency constant*.

Intrinsic factors are optical material properties used in the *radiative transfer equation*, i.e. spectral refractive index, spectral absorption and scattering index and the spectral scattering phase function. Gkioulekas *et al.* [2] have shown that the scattering phase function plays an important role in how translucency is perceived and that this property alone results in at least two perceptual dimensions. The absorption and scattering coefficients induce another two perceptual dimensions. The larger the absorption and scattering coefficients the smaller is the *mean free path*, i.e. the average distance a photon travels within the material before an absorption or scattering event occurs. Objects that appear almost transparent can be made of a material with a large mean free path or with a small mean free path only if they are thin enough. The reverse applies for objects that appear almost opaque. This makes the geometry of an object to be another intrinsic property affecting

apparent translucency.

The qualitative and quantitative understanding of the impact of geometry on translucency is not only of academic interest, but became recently relevant in various practical applications as well. This is because modern multimaterial 3D printers allow the reproduction of spatially-varying color and translucency in addition to the shape of an object pushing the design freedom further as ever imagined [3]. This is possible by mixing a fully transparent printing material with more opaque printing materials colored in cyan, magenta, yellow, black and white. Since geometry is an intrinsic factor impacting perceived translucency, transferring the apparent translucency from an object made of a specific material to another shape for 3D printing would require a geometry-adaptive adjustment of the printing material mixing ratios considering the limited translucency gamut of the device.

A prerequisite for process control and quality assurance for this mixing, as well as for communicating the necessary appearance information, is a joint color and translucency space possessing a metric reflecting apparent color and translucency differences. Since the object geometry affects the magnitude of apparent translucency, the translucency-relevant part of the metric must depend on distinct geometry features that are unknown to date.

Visual experiments showed that the human visual system does not perform inverse optics to deduce intrinsic optical parameters of the underlying material to judge translucency, but rather relies on simple image cues [4] that consist in luminance contrasts within distinct spatial frequency bands of non-specular object regions [5] and in the information of the occluded scene that shines through the object – either because the object is made of an optically thin material or because it is just thin enough. In addition to that, Gigilashvili *et al.* [6] have demonstrated on physical objects that presence of the thin parts in an object could compensate intrinsic material properties and evoke similar translucency perception between objects of different colorant density.

In this paper we present preliminary work to quantitatively analyze the impact of geometry features (structural thickness) on apparent translucency and to formulate further research hypotheses. We do this for a set of simple virtual materials employed in a recently proposed color and translucency space designed particularly for 3D printing applications [7]. These materials possess an isotropic scattering phase function and vary just in wavelength-independent absorption and scattering coefficients. We render images of differently shaped objects made of these materials by solving the steady-state radiative transfer equation in a virtual viewing booth using a Monte Carlo path tracer. For determining equidistant suprathreshold differences for five different geometries, we use the method of constant stimuli and compare pairs of these physically accurate renderings shown under controlled viewing

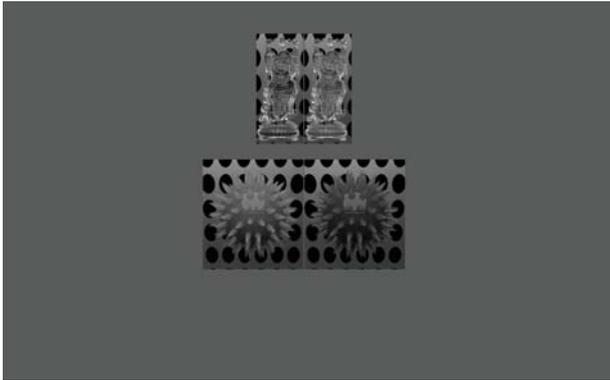


Figure 1: A sample screenshot from the experiment. The top pair represents a suprathreshold anchor pair, while the bottom pair is a test pair judged against it.

conditions on a color calibrated display.

The paper is organized as follows: in the next section we provide a description of the experimental setup. In the subsequent section, the results are introduced and discussed. Finally, we draw conclusions and outline the directions for the future work.

Experimental Setup

Stimuli and Experimental Protocol

The method of constant stimuli [8] has been used to determine suprathreshold translucency differences. A pair of Buddha images¹ with a suprathreshold translucency difference has been used as the anchor pair similarly to Urban *et al.* [7]. This anchor pair was compared to a test pair that was composed of two similarly-shaped objects with different wavelength-independent absorption and scattering coefficients. All materials used in the experiment possess an isotropic scattering phase function. In total five different shapes were used for test pairs. The pairs were placed on a neutral gray background with a 3-pixel gap separating the images within the pair. The pairs were positioned atop of each other. The vertical order between anchor and test pairs, as well as the horizontal allocation of the images within a pair was randomized. Subjects were given the following instruction: "Please, select a pair, either a top, or a bottom one, with higher translucency difference" without explicitly defining translucency. The selection was made using arrow keys of the standard keyboard. A screenshot of the experiment is shown in Figure 1.

All images have been rendered using the Mitsuba Physically Based Renderer [9]. Five control points have been defined in the absorption-scattering space [7]. Afterwards five samples were selected for each of the following four directions:

1. 1-dimensional change - increasing scattering coefficient for control points (CP) 1 and 2; decreasing scattering coefficient for CPs 3, 4 and 5; fixed absorption. Beige in Table 1.
2. 1-dimensional change - increasing absorption for CPs 1 and 2; and decreasing absorption for CPs 3,4, and 5; fixed scattering. Orange in Table 1.
3. 2-dimensional change - increasing absorption and scattering for CPs 1 and 2; decreasing absorption and scattering for CPs 3, 4, and 5. All points were located on a straight line

defined as $y = x$ for CPs 1-3; $y = x - 116$ for CP 4; and $y = x + 116$ for CP 5, where x is scattering, and y is absorption. Marked green in Table 1.

4. 2-dimensional change - increasing absorption and decreasing scattering with points located on the following straight lines: $y = -x + 9$ for CP 1; $y = -x + 155$ for CP 2; $y = -x + 301$ for CP 3; and $y = -x + 185$ for CPs 4 and 5. Marked blue in Table 1.

All pairs of absorption and scattering coefficients used to determine the visual dataset are shown in Table 1. We assume symmetry around control points w.r.t. ΔA differences introduced in [7], and leave studying opposite directions and validation of this assumption for future work. This assumption was inspired by the determination of suprathreshold color-difference tolerances by Berns *et al.* [10] in CIELAB. Each of the samples were compared against the control point creating a test pair for the experiment. Considering that five different shapes and control points have been used, the experiment resulted in 500 comparisons in total: 5 control points \times 4 directions \times 5 samples per direction \times 5 shapes. The samples were selected in relation to the control point so that at least one of the samples had an apparent translucency difference smaller than the anchor pair, while at least one of them had a larger difference. In addition to Alpha differences (ΔA), a trial-and-error approach has been used to select proper absorption and scattering coefficients.

As mentioned above, five shapes have been used for the experiment: a Buddha 3D model used in [7], a perfect sphere and three spheres with spikes of different length and thickness. All spheres were composed of 3,932,160 vertices and 1,310,720 triangles. The topography of the unit sphere was manipulated as follows to create spikes: 100 nearly uniformly distributed antipodally symmetric points [11] have been generated on a unit sphere. Afterwards, the unit sphere has been converted into spherical coordinates and the radius parameter of 100 uniformly distributed vertices, as well as that of other vertices within a fixed circle have been increased to introduce the spikes. The shapes used for the experiment are illustrated in Figure 2. Adding spikes with a simple manipulation introduced high number of thin areas on a compact spherical object leading to significantly different values of medial axis shape descriptor [12]. Comparison among more complex shapes and familiar objects will be studied in the future.

Display and Viewing Conditions

A 24.1 inch EIZO ColorEdge CG246 LCD monitor was used for the experiment. The display was calibrated to CIE D65 white point with gamma equal to 2.2. Luminance of the display was measured using Konica Minolta CS-2000 spectroradiometer on a rendered image of a perfect diffuser patch and was equal to 196 cd/m^2 with color temperature equal to 6542K. The experiment was conducted in a completely dark room with the display being the only light source. The display was placed roughly 60cm away from the observer. The size of each anchor pair image was 148×348 pixels. Height of test pair images was fixed to 348 pixels, while the width varied depending on the shape; the resolution of the largest image was 348×348 pixels. Each anchor pair image occupied 3.81° of the visual field horizontally, and all images occupied 8° of the visual field in height. In width, the sphere without spikes, and the ones with short, medium and long spikes

¹<https://graphics.stanford.edu/data/3Dscanrep/>

Table 1: The table summarizes absorption and scattering coefficients sampled over absorption-scattering space and used to construct the dataset. Each pair of columns represents one control point. Control point coordinates are marked red, while the test samples in four different directions are labeled with yellow, orange, green, and blue cells respectively.

Control Point 1		Control Point 2		Control Point 3		Control Point 4		Control Point 5	
Scattering	Absorption								
4.5	4.5	77.5	77.5	150.5	150.5	150.5	34.5	34.5	150.5
4.7	4.5	78	77.5	148	150.5	145	34.5	30	150.5
6	4.5	82	77.5	141	150.5	135	34.5	20	150.5
7	4.5	90	77.5	121	150.5	125	34.5	15	150.5
10	4.5	110	77.5	100	150.5	110	34.5	5	150.5
20	4.5	130	77.5	71	150.5	95	34.5	0	150.5
4.5	4.8	77.5	80	150.5	148	150.5	30	34.5	145
4.5	6	77.5	100	150.5	140	150.5	23	34.5	130
4.5	8	77.5	120	150.5	125	150.5	15	34.5	120
4.5	12	77.5	150	150.5	100	150.5	7	34.5	100
4.5	20	77.5	200	150.5	80	150.5	0	34.5	50
4.7	4.7	85	85	140	140	146	30	30	146
5.5	5.5	95	95	100	100	139	23	25	141
7	7	100	100	80	80	131	15	17	133
9	9	140	140	60	60	123	7	10	126
12	12	1000	1000	25	25	116	0	0	116
4.3	4.7	72	83	145	156	148	37	30	155
4	5	65	90	135	166	140	45	25	160
3	6	50	105	120	181	130	55	17	168
2	7	40	115	110	191	120	65	10	175
0	9	30	125	90	211	105	80	0	185

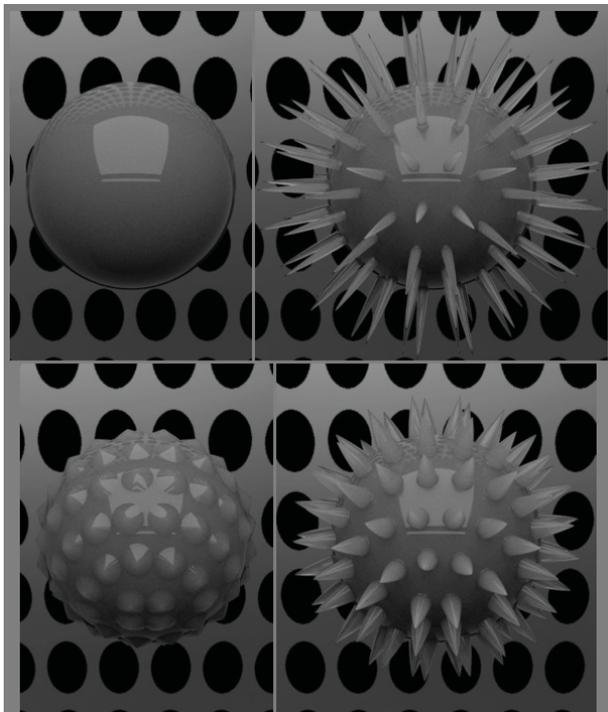


Figure 2: Spikes of various length and thicknesses were added to a unit sphere to introduce three new shapes with thin parts. Illustrated objects are rendered on control point 2 coefficients.

occupied 6.20° , 6.48° , 8.10° , and 8.96° of the visual field, respectively. Minimum 30-minute period for warm up was considered before using the display.

Observers

In total, 27 subjects, 18 males and 9 females, from 20 different nationalities have participated in the experiment including three authors of this paper. The experiment was organized in three sessions and each comparison has been assessed by 20 observers, i.e. no comparisons were assessed by all observers, and not all observers were shown all comparisons. 21 observers had a background in color science, imaging, vision, or related fields. The median age of the observers was 31 years and the standard deviation was 10.97 years. All observers passed a Snellen visual acuity test before taking the experiment.

Analysis

Frequency analysis of the observer responses has been conducted to identify the distance from the control point in the investigated absorption-scattering coefficient direction at which the translucency difference matches the translucency difference of the anchor pair. The frequency the control-test pair's translucency difference was considered to be larger than that of the anchor pair was assigned to the corresponding Euclidean distance from the control point in absorption-scattering space. A probit model was fit to estimate T50 distances from the control points in the investigated direction. The T50 distance is the distance between sample and control point at which 50% of the subjects consider the test pair difference smaller than that of the anchor pair, while the other 50% consider it larger. We compared T50 distances for the investigated shapes.

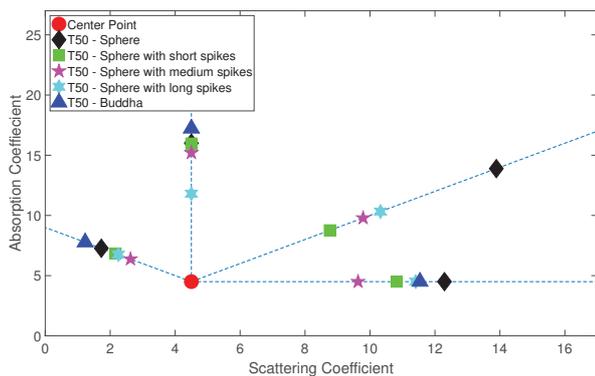


Figure 3: Control Point 1: T50 points for each shape and direction. Colored markers represent T50 points reached within the dataset, whilst gray ones with identical shape represent T50 points estimated with Probit model fitting. Estimated values that either failed the goodness-of-fit test, or are unrealistic, are missing from the figure.

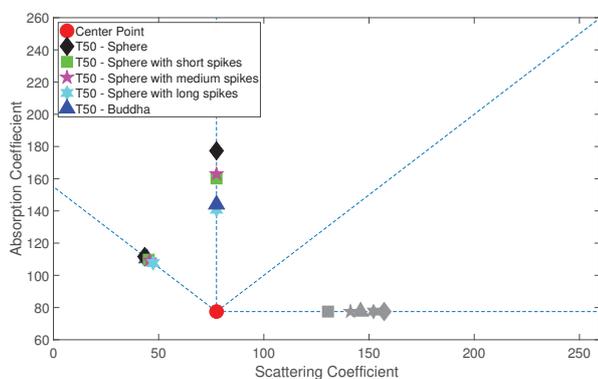


Figure 4: Control Point 2: T50 points for each shape and direction.

Results and Discussion

A Probit binomial model has been fit to the data, where the Euclidean distance in absorption-scattering space between a control point and a test sample ($\Delta_{[\text{absorption}, \text{scattering}]}$) is a predictor variable, and the number of occasions the test pair translucency difference was judged to be larger than that of the anchor pair is considered a success rate. Finally, χ^2 goodness-of-fit test was conducted with α set to 0.05. Estimated thresholds (T50 points) in absorption-scattering space, where subjects start to perceive translucency difference from a control point (referred as a "Center Point" in the figures) has been plotted for each shape and are illustrated on Figures 3 through 7. It is worth mentioning that the T50 point has not been identified in all cases. Estimations that failed a goodness-of-fit test ($\alpha > 0.05$), or with negative estimated absorption-scattering coefficients, are excluded from the plot. If T50 distances were not within the distance range of the control-test pair samples but passed the goodness-of-fit test, then they are shown in grayscale.

The subjects demonstrated highest sensitivity towards translucency changes in optically thin areas, namely, around *Control Point 1* (Figure 3). All T50 points, except for the one in increasing absorption-scattering direction for a Buddha shape, have passed goodness-of-fit test. All T50 points have been reached

within significantly shorter absorption-scattering distance than it is the case for optically thick samples. This could be explained with the fact that cues for transparency and translucency perception are different, and background distortion, that is absent for optically thicker translucent materials, enables very high sensitivity in changes of material properties. While average Euclidean distance in absorption-scattering space needed to reach T50 point decreases as the spike length is increased (9.12, 6.78, 6.49, 6.40 for a sphere, and a sphere with short, medium, and long spikes, respectively), the behaviour is not consistent and further study is needed to validate the hypothesis.

The behaviour changes for the second control point (Figure 4). The threshold is very consistent among all shapes for increasing absorption - decreasing scattering direction. The point is also reached in increasing absorption direction, while, reliable estimations are done in increasing scattering direction. Significant difference among shapes has not been observed for this control point. The average distance (among all shapes) needed to reach T50 point is shortest for the case, when absorption increases and scattering decreases followed by distances in increasing scattering and increasing absorption dimensions (45.65, 67.89, and 79.53, respectively), while it is never reached in case of equidistant increase in both dimensions. The distances to T50 are higher than it was in optically thin area, as already mentioned above. When the observers were asked to reflect on the experiment, the vast majority of them mentioned that in absence of all other cues for optically thick objects, they relied solely on the lightness. Simultaneous increase in absorption and decrease in scattering impacts lightness a lot more than just one-dimensional change. While increase in absorption makes a matter darker, and increase in scattering makes it lighter, these phenomena compensate each other in case of equidistant simultaneous increase in both dimensions, and change in perceived lightness is negligible. This could be an explanation why T50 is never reached in that direction. Preliminary study showed the identical trend for all control points in optically thick areas. Therefore, the samples were selected towards optically thin area, as a hypothesis of never reaching T50 in increasing absorption - increasing scattering direction should be tested in the future.

For *Control Point 3*, the T50 distance could only be determined in decreasing absorption - decreasing scattering direction. Although absorption and scattering estimations for a sphere and the one with medium spikes, are negative, i.e. physically unrealistic. In this case, the selection of the test pairs was the reason (Table 1), as none of the farthest samples in three other dimensions evoked larger-than-suprathreshold apparent translucency difference. This issue in experimental design should be addressed in future work.

Figure 6 summarizes results for *Control Point 4*. The results are consistent among all shapes in decreasing absorption, and decreasing absorption-decreasing scattering directions. Average distance to reach T50 was 34.06, 30.83, 29.16, and 34.21 for a sphere, a sphere with short spikes, the one with medium spikes, and the Buddha, respectively. However, the threshold is not reached in decreasing scattering direction, as the farthest sample in that direction was not considered larger-than-suprathreshold different. The results for a sphere with long spikes have been omitted intentionally, as some computational inaccuracy has been discovered in the rendering. The study of accurate renderings for

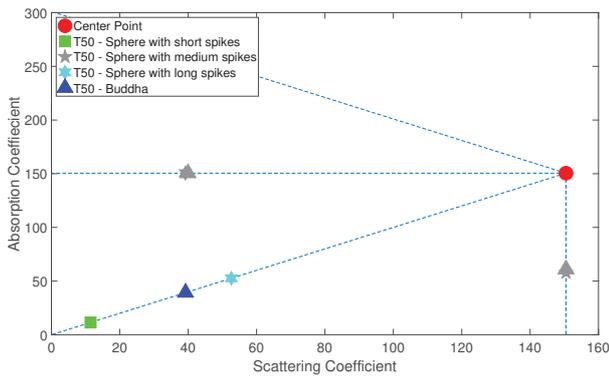


Figure 5: Control Point 3: T50 points for each shape and direction.

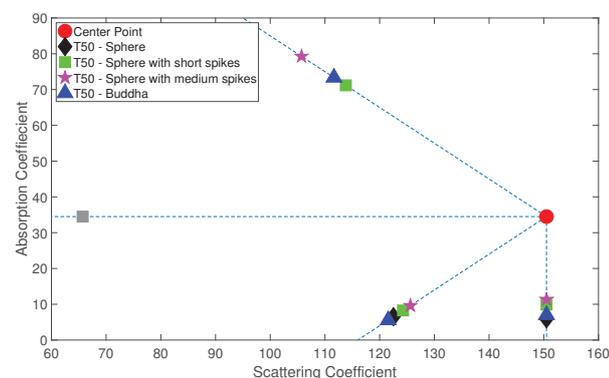


Figure 6: Control Point 4: T50 points for each shape and direction.

this particular shape and control point will follow in the future. Although the threshold is reached for three shapes in increasing absorption - decreasing scattering direction, the distance is rather large.

Finally, the results for *Control Point 5* are shown on Figure 7. Surprisingly, T50 was reached only for one object in all four directions. It is worth mentioning that the majority of the coordinates were estimated as negative values. This can be explained with the proximity of the control point to the scattering axis. High absorption makes objects too dark, and the magnitude of change in scattering is not big enough for perceiving a large enough translucency difference. Interestingly, decreasing absorption direction was easiest one to reach the T50. The estimation for a sphere with long spikes nearly passed a goodness-of-fit test, while the estimation was negative for the object with short spikes. It is worth mentioning that the object with the short spikes should be considered a special case. Regardless of the presence of the spikes, almost the entire surface of the sphere is bumpy, removing see-through cues even for optically thin materials. This effect is illustrated on Figure 8, where the four shapes are rendered using optically thin material. This illustrates that just material properties and a simple shape descriptor, as a histogram of surface-to-medial-axis distances, might not be enough for modeling apparent translucency and that the surface roughness needs to be considered as well.

One of the reasons why T50 distances could not be computed in many cases, is that observers seemed to be very insensitive in

changes of the scattering coefficients in optically thick regions. Obviously the authors were biased by their a priori knowledge of the material properties, in the trial-and-error process of sample selection and assumed a higher sensitivity to perceive subsurface light transport differences in the observer panel. On the contrary to our expectation, the majority of the observers argued that the control-test pairs in optically thick regions of the space were both opaque with an apparent translucency difference smaller than the anchor pair difference possessing strong shine through cues. In other words, the translucency difference of the anchor pair was caused by the difference in shine through cues, while two optically thick samples were considered fully opaque, and just different in lightness due to reflection properties. Several observers described them as "two fully opaque billiard balls with two different colors". This observation supports Fleming's [4] proposal that humans have poor ability to invert optics, and that the impact of shine through cues on perceived translucency differences is much larger than that of luminance contrast cues.

While we assume wavelength-independent absorption and scattering, and limit our study to grayscale images, translucency might also impact chromatic information, introducing new cues and additional degree of freedom. According to Berns [10], stimuli simultaneously varying in chromaticness and lightness need larger effort to be assessed than the ones with just lightness difference. Thus, it is important to study chromatic samples as well.

Conclusion and Future Work

We have discussed a psychophysical experiment where absorption and scattering properties, as well as shape of the objects vary in order to identify the impact of shape, and particularly presence of thin structures, on apparent translucency differences. We have observed that in many cases the presence of the thin regions in an object leads to smaller T50 distances, i.e. more sensitivity towards apparent translucency differences. However, further experiments are needed to confirm this proposal, as well as to quantify this impact. Although indications of impact of shape are not very strong and are primarily limited to optically thin materials, we could still observe several interesting phenomena, and generate research hypothesis for future studies. According to our data, observers tend to be a lot more sensitive to changes in optically thin materials, rather than in optically thick ones. Moreover, considering that we have compared an optically thin anchor pair with

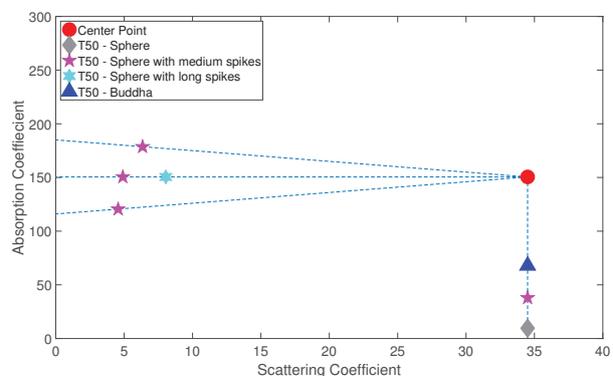


Figure 7: Control Point 5: T50 points for each shape and direction.

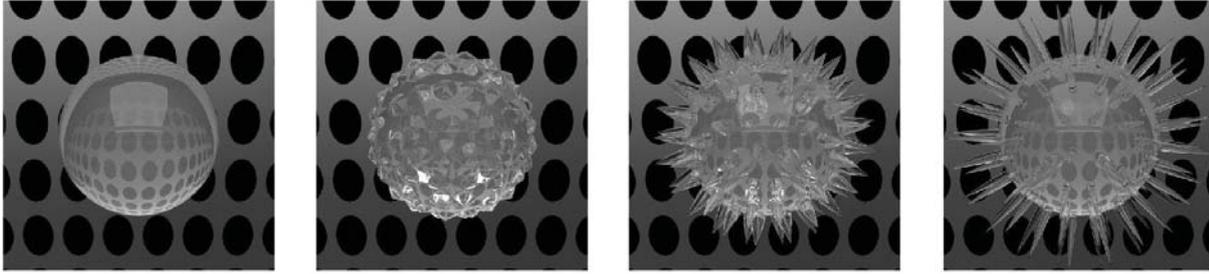


Figure 8: Images of the various shapes corresponding to the optically thin Center Point 1.

optically thick test pairs, this phenomenon might have impacted the entire experiment, and might be responsible for the fact that for some samples, apparent translucency difference among test pairs was never been judged larger than that of the anchor pair. In order to study the impact of this factor, we are currently conducting the identical experiment with an optically thick anchor pair. The comparison with the results introduced above will be reported in future work.

Apart from that, it is likely that for optically thick control points that does not possess any shine-through cues, the perceptual translucency difference toward the optically thick direction is never bigger than that of the anchor pair. This could be studied with larger steps and better sampling in that direction. This observation further highlights the importance of studying the magnitude of the impact of shine-through and luminance contrast cues on translucency difference. Impact of shape might be especially significant for optically thin objects. In the future study, we will correlate apparent translucency differences with object shape descriptors. Particularly, the medial axis [12] of an object will be calculated and a thickness histogram will be constructed based on the surface-to-medial-axis distances. As the presence of thin areas will be quantified as a thickness histogram, the histogram will be correlated with perceptual translucency differences. However, more complex descriptors might be needed, as mentioned above regarding the example shown in Figure 8.

In addition to that, the experiment has several limitations that need to be addressed in future work: first of all, heterogeneous background and its deformation as seen through optically thin media are a strong cue for transparency assessment, while the impact of background-based cues might be mitigated in presence of a homogeneous background. The experiment should be replicated on a homogeneous background. Secondly, the study has been limited to an achromatic environment under one illumination condition (front-lit) and fixed microscopic surface roughness of the objects. We should consider more variables, like chromaticity, different illumination conditions, and several levels of surface roughness, in order to develop a translucency difference formula. Finally, it is worth mentioning that while we limit ourselves to simple manipulations on a unit sphere, the final model should be validated on random shapes of various thickness histograms.

Acknowledgements

The work has been funded by the Measuring and Understanding Visual Appearance - MUVApp project of the Research Council of Norway (project number 250293). The authors also want to thank Tejas Tanksale for his contributions in the stimuli rendering process.

References

- [1] B. Xiao, B. Walter, I. Gkioulekas, T. Zickler, E. H. Adelson, and K. Bala, "Looking against the light: How perception of translucency depends on lighting direction," *Journal of vision*, vol. 14, no. 3:17, pp. 1–22, 2014.
- [2] I. Gkioulekas, B. Xiao, S. Zhao, E. H. Adelson, T. Zickler, and K. Bala, "Understanding the role of phase function in translucent appearance," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 5:147, pp. 1–19, 2013.
- [3] A. Brunton, C. A. Arian, T. M. Tanksale, and P. Urban, "3d printing spatially varying color and translucency," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 157:1–157:13, 2018.
- [4] R. W. Fleming and H.H. Bülthoff, "Low-level image cues in the perception of translucent materials," *ACM Transactions on Applied Perception (TAP)*, vol. 2, no. 3, pp. 346–382, 2005.
- [5] I. Motoyoshi, "Highlight–shading relationship as a cue for the perception of translucent and transparent materials," *Journal of vision*, vol. 10, no. 9, pp. 1–11, article no. 6, 2010.
- [6] Davit Gigilashvili, Jean-Baptiste Thomas, Jon Yngve Hardeberg, and Marius Pedersen, "Behavioral investigation of visual appearance assessment," in *Color and Imaging Conference*. Society for Imaging Science and Technology, 2018, vol. 2018, No. 1, pp. 294–299.
- [7] P. Urban, T. M. Tanksale, A. Brunton, B. M. Vu, and S. Nakauchi, "Redefining A in RGBA: Towards a Standard for Graphical 3D Printing," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 3, pp. 21:1–21:15, 2019 (preprint available at <https://arxiv.org/abs/1710.00546>).
- [8] Peter G Engeldrum, *Psychometric scaling: a toolkit for imaging systems development*, Imcotek, 2000.
- [9] Wenzel Jakob, "Mitsuba renderer," 2010, <http://www.mitsuba-renderer.org>.
- [10] Roy S Berns, David H Alman, Lisa Reniff, Gregory D Snyder, and Mitchell R Balonon-Rosen, "Visual determination of suprathreshold color-difference tolerances using probit analysis," *Color Research & Application*, vol. 16, no. 5, pp. 297–316, 1991.
- [11] Cheng Guan Koay, "A simple scheme for generating nearly uniform distribution of antipodally symmetric points on the unit sphere," *Journal of computational science*, vol. 2, no. 4, pp. 377–381, 2011.
- [12] Daniel Rebain, Baptiste Angles, Julien Valentin, Nicholas Vining, Jiju Peethambaran, Shahram Izadi, and Andrea Tagliasacchi, "LSMAT least squares medial axis transform," in *Computer Graphics Forum*. Wiley Online Library, 2019.