Perceiving translucent materials

Roland W. Fleming[1], Henrik Wann Jensen[2] & Heinrich H Bülthoff[1]

[1]Max Planck Institute for Biological Cybernetics.
[2]Computer Graphics Laboratory, UCSD.

Abstract

Many common materials, including fruit, wax and human skin, are somewhat translucent. What makes an object look translucent or opaque? Here we use a recently developed computer graphics model of subsurface light transport [Jensen, et al., 2001] to study the factors that determine perceived translucency. We discuss how physical factors, such as light-source direction can alter the apparent translucency of an object, finding that objects are perceived to be more translucent when illuminated from behind than in front. We also study the role of a range of image cues, including colour, contrast and blur, in the perception of translucency. Although we learn a lot about images of translucent materials, we find that many simple candidate sources of information fail to predict how translucent an object looks. We suggest that the visual system does not rely solely on these simple image statistics to estimate translucency: the relevant stimulus information remains to be discovered.

CR Categories: J.4 [Computer Application]: Social and Behavioral Sciences; I.3.7. [Computer Graphics] Three-dimensional Graphics and Realism.

Keywords: subsurface scattering; perception; psychophysics.

1. Introduction

Many materials that we commonly encounter are translucent, including leaves and fruit flesh; textiles and papers; various stones, such as agate or marble; soaps; wax; some types of glass; milk and human skin. How do we recognize these materials? What image cues allow us to tell that a surface is translucent rather than opaque? Here we use a combination of psychophysics and computational image analysis to study the image properties that underlie the distinctive appearance of translucent materials.

When light strikes a translucent material, it enters the body of the object, scatters, and re-emerges from the surface. The light that bleeds through translucent objects gives them a characteristic softness and glow. However, although *light* is visible through translucent objects, *form* is generally not. This has profound consequences for the image cues underlying translucency perception, and results from the way that light scatters beneath the surface of an object.

Subsurface light transport

Recent advances in computer graphics [Jensen *et al.* 2001] allow us to simulate translucent materials realistically and efficiently. In translucent materials light scatters below the surface of the object. This phenomenon, called *subsurface scattering*, causes light to spread beneath the surface and reemerge in a region around the point of illumination. In contrast, traditional models for light scattering based on the Bidirectional Reflectance Distribution Function (BRDF) assume that materials are opaque, and that all light is reflected from the point that is illuminated (Figure 1).

(a) Surface reflection (BRDF)
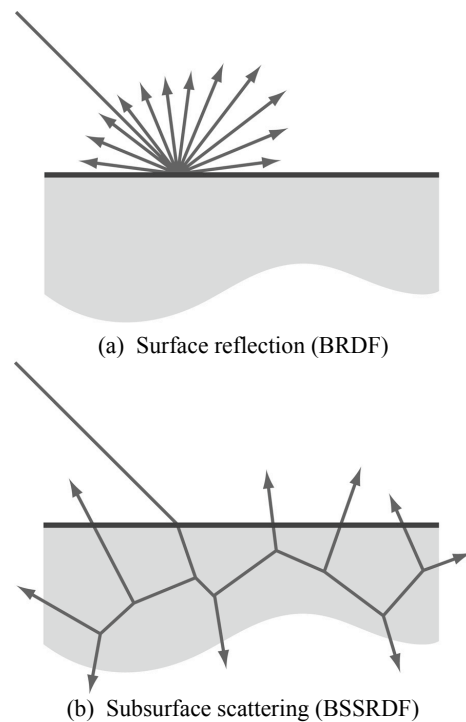

(b) Subsurface scattering (BSSRDF)

Figure 1. The optical behaviour of reflective and translucent materials.

A full simulation of translucency requires solving the radiative transport equation [Chandrasekhar, 1961]. Jensen *et al.* [2001] simplify the light scattering by assuming that the translucent material is homogeneous. In this case the scattering of light can be approximated by a diffusion equation as well as a term for single scattering. These terms form a Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF).

Figure 2. The laughing Buddha rendered with different settings of the BSSRDF. The scattering coefficient progressively increases from left to right, all other parameters are held constant. The apparent material quality ranges from glassy, through jade-like to porcelain. Model courtesy of the Stanford 3D scanning repository [http://graphics.stanford.edu/data/3Dscanrep/].

The parameters in the BSSRDF are the refractive index of the material, the phase function (Jensen *et al*. [2001] use the Henyey-Greenstein [1941] phase function), and the scattering and absorption coefficients. The absorption and scattering coefficients specify the probability that a photon will be absorbed or scattered when traveling a given distance within the material. In most materials these parameters determine the color as well as the degree of translucency. This is particularly the case in highly scattering materials, where the diffusion term dominates [Jensen and Buhler 2002], and the phase function can be mostly ignored.

Koenderink and van Doorn [2001] pointed out that translucent material have an appearance that differs from traditional Lambertian materials, and they investigated the diffusion of light in simple geometric shapes in order to analyze the shape from shading characteristics of translucent objects. They observed that the shape of translucent objects is difficult to analyze as the appearance depends not only on the location of the lighting and the observer, but also on the shape of the object, since most of the lighting is due to scattering within the translucent object.

In this paper we study the effect of the absorption and the scattering coefficient on the perceived appearance of translucent materials. We also study the effect of other components such as highlights and demonstrate how these parameters significantly influence our perception of the material. An example is shown in Figure 2, where a Buddha statue is changed from highly translucent to mostly opaque. Changes in translucency can influence apparent softness, realism and even how edible an object looks.

2. Physical factors that influence the impression of translucency

Although the parameters of the BSSRDF can dramatically alter the appearance of a material, we have found that other factors can also influence how translucent an object looks. Here we discuss physical factors and viewing conditions that affect perceived translucency.
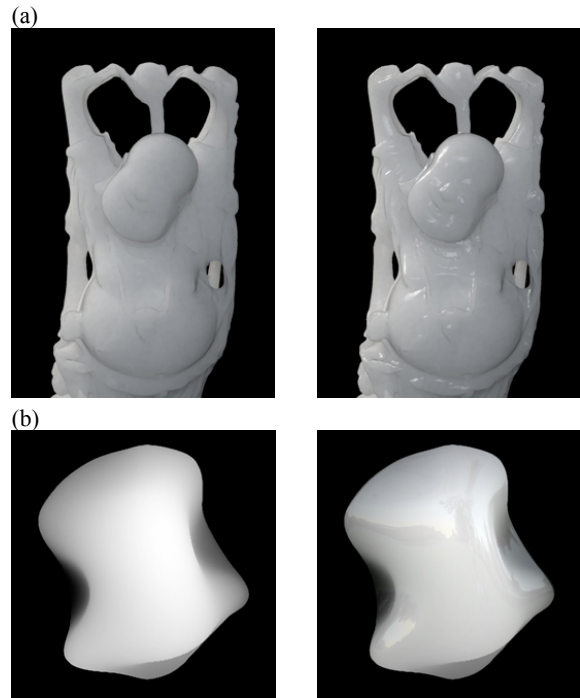
Highlights and translucency

(a)

(b)

Figure 3. The role of highlights in the perception of translucency. Left images have no highlights, right images have highlights. Most observers agree that the specularities make the impression of translucency more compelling.

Highlights occur when light is specularly reflected from the surface of an object. Note that specular reflections are caused by the *interface* between two materials of different refractive index (i.e. the surface itself). As such, highlights are not a direct consequence of light transport within the body of an object. We

might therefore expect the perception of translucency to be quite independent of specular reflections[1].

However, many translucent materials that we commonly encounter are somewhat glossy (e.g. plastic, wax or marmalade). This means that the human visual system may 'expect' translucent materials to exhibit specular reflections.

Interestingly, we have found that highlights can contribute to the visual impression of translucency. Consider, for example, the images shown in Figure 3. Each pair of images has been rendered under identical conditions. The only difference between the images is the presence or absence of specular highlights. Observers generally agree that the glossy surfaces look more like canonical translucent materials than the surfaces without highlights.

It is worth noting that the visual system tolerates physical inconsistencies between the transmitted and reflected components of the image. For example, in image 3b, the translucent component was rendered under a single point light source from in front, while the specular component was rendered under full-scene illumination from the Debevec [1998] panoramic light probe database. The dominant illumination direction in this light probe is from the left. Despite the inconsistency, we readily fuse the two components into a single percept of a glossy, translucent object. The inconsistency in illumination does not hinder the sense of translucency. On the contrary, observers generally agree that the inconsistently illuminated image looks more realistic than its consistent, but highlight-less counterpart.

Effects of light-source direction

Anyone who has spent time playing with translucent materials as a child, will have noticed that they tend to look more translucent when held up to the light. When illuminated from behind, a gemstone or slice of fruit is filled with a distinctive glow, which enhances the sense of the object's translucency.

An example of this effect is shown in Figure 4. A torus is illuminated from six different orientations (three behind and three in front). Note that the apparent translucency of the torus changes depending on the light direction. We have studied this effect systematically, using a psychophysical translucency-matching task.

Fifteen subjects were presented with images of two translucent tori simultaneously. On each trial the image on the left (the "Test" image) was selected by the computer, while the image on the right (the "Match") could be adjusted by the subject. The two objects were illuminated from different directions. The subject's task was to adjust the translucency of the Match image until the torus appeared to be made out of the same material as the Test torus, despite differences in the illumination.

The translucency of the Test object could be one of three different values, ranging from highly translucent—like jade—to highly opaque[2]. We tested the appearance of these three objects under 12 different illumination directions around the object.

---

[1] Note, however, that the binocular depth of highlights is generally *not* on the surface itself [Blake and Bülthoff 1990, 1991]. Here we are dealing with monocular images, but it is possible that adding stereoscopic information would alter the effects described here.

[2] Specifically, the absorption coefficient was held fixed at 0.1 while the scattering coefficient was set to 2.00, 7.13 or 20.0.
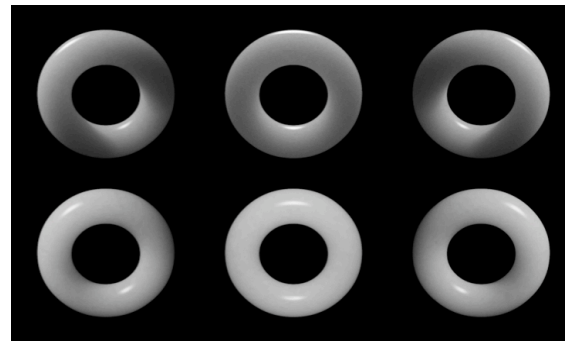

Figure 4. A translucent torus illuminated from six different directions. Top row, illumination from behind. Bottom row, illumination from in front. Note that the apparent opacity of the torus changes depending on the direction of illumination.

The Match object was always illuminated from the same direction (back left), but the subject could adjust the apparent translucency through 128 different values, which spanned a range greater than the three values used for the Test objects. Specifically, the absorption coefficient was held fixed at 0.1 (i.e. the same as the Test objects), while the subject adjusted the scattering coefficient through 128 steps from 0.4 to 60. The step sizes were non-linear to create a more uniform perceptual scale. Subjects readily agreed with the statement "when I move the mouse to the right, the torus appears to change from being relatively translucent to relative opaque, while everything else about the scene appears to stay the same". Therefore, for subsequent discussion, we refer to this modified scattering coefficient scale as "perceived opacity".

All the Test images were shown to the subjects twice in a randomly interleaved sequence. Subjects were given unlimited time to adjust the Match stimulus using the mouse, and once satisfied with the match, could move onto the next trial in the sequence by pressing a button on the keyboard.

Before the main experiment, subjects were given practice trials consisting of the same three Test objects illuminated from 9 different directions. During practice, subjects were asked to say out loud a numerical rating of the perceived opacity on a scale from 1-7 before adjusting the mouse. No feedback was given, the practice was simply intended to help subjects orient to the task.
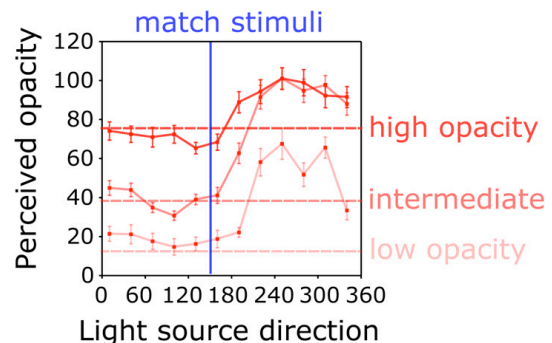

Figure 5. Mean data across 15 subjects. Perceived opacity varies as a function of light source direction. Angles less than 180 degrees corresponding to lighting from behind the object, while angles greater than 180 are illuminated from in front. Error bars represent standard error.

The mean data across all subjects is shown in Figure 5. Note that if the observers were able to accurately estimate the intrinsic parameters of the BSSRDF irrespective of the illumination, the data would fall along the three horizontal lines. This is not the case. Instead, perceived opacity undergoes a dramatic change when the lighting is altered. All three objects appeared significantly more opaque when illuminated from in front than from behind. The effect is most marked for the object of intermediate translucency.

Thus, for the conditions used in our experiment, the visual system seems poor at 'discounting' the effects of light source direction. Objects tend to appear more translucent when illuminated from behind. It is interesting to contrast this with previous work on the perception of other material attributes. For example, Fleming, Dror and Adelson [2002] found that the perception of gloss remains relatively stable across change in illumination, as long as the pattern of illumination is realistic.

One consequence of the current finding is that if we wish to enhance or emphasise the apparent translucency of an object, for example in an animation, we should organize the scene lighting so that the object is illuminated predominantly from behind.

3. Image measurements

If we wish to understand how the visual system estimates translucency, we must identify the image cues that carry information about translucency. To this end, we have measured how various image properties vary with parameters of the BSSRDF. We describe a number of these measurements here.

Colour saturation

How does colour influence perceived translucency? When white light passes through a coloured translucent object, it is progressively filtered, and emerges coloured. Interestingly, hue, saturation and intensity can all vary as a function of the distance travelled by a ray through a translucent material. The BSSRDF can produce a wide range of colour phenomena, depending on the colour values of the scattering and absorption coefficients. A few examples are shown in Figure 6 (see colour plate).
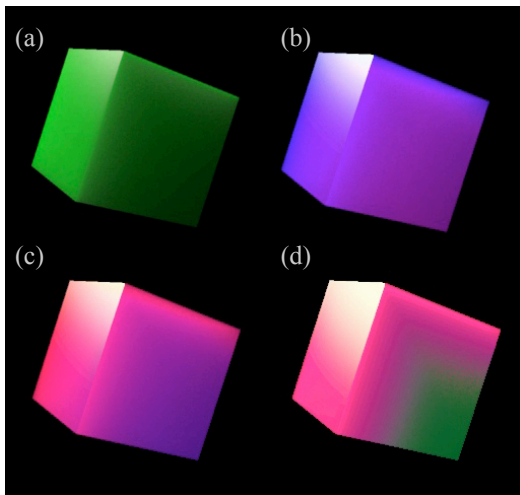

Figure 6. Cubes exhibiting various colour effects using the BSSRDF. (a) Opaque and uniform green; (b) mauve body colour with blue fringe; (c) mauve body with pink fringe; (d) a material unlike those commonly encountered in the world: the colour varies continuously from pink through mauve to green.

We know that colour is not *necessary* for the perception of translucency, because a black-and-white photograph can nevertheless yield a vivid impression of translucency. However, can colour *modify* the sense of translucency when present?

We have found that colour saturation can affect the way a translucent object appears to 'glow'. Consider the two cubes in Figure 7 (see colour plate). The hue and intensity components of the two images are identical, what differs is the saturation component. In (a) the saturation is positively correlated with the intensity image, while in (b) it is negatively correlated. The mean saturation is held constant. Observers generally agree that (a) appears to have a 'warmer' glow, while (b) appears more 'icy' or 'dilute'.
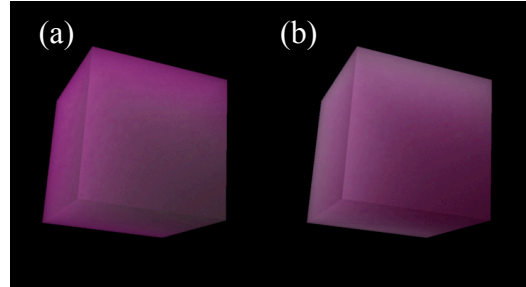

Figure 7. (a) Colour saturation is positively correlated with intensity. (b) saturation is negatively correlated with intensity. Mean saturation is identical for the two images. Most observers agree that (a) appears 'warmer' than (b).

Although saturation variations can affect perceived translucency, they are insufficient on their own to yield an impression of translucency. This is demonstrated in Figure 8 (see colour plate). In (a), saturation is held constant across the image, while intensity varies. In (b), intensity is held constant while saturation varies. Most subjects agree that (a) looks translucent, while (b) does not. This suggests that the saturation component is neither necessary nor sufficient to yield an impression of translucency.
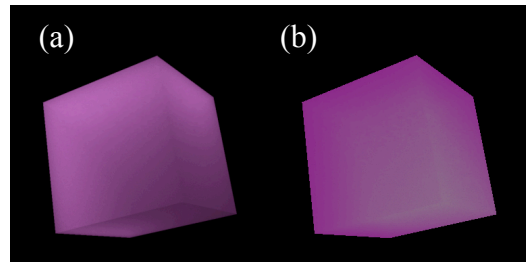

Figure 8. (a) Intensity varies across the image but saturation is constant. (b) intensity is constant but saturation varies. Most observers agree that (a) looks more translucent than (b)

Image contrast

Light diffuses through translucent materials, much like dye diffusing through a fluid. When illuminated translucent objects become 'filled' with light. An important consequence of this is that points on the surface that do not receive any direct illumination (i.e. they are in shadow) can nevertheless receive light from *within* the body of the object. Conversely, regions that receive strong direct illumination tend to dissipate the incident

light by transmitting it to other parts of the object. This has the effect of reducing the overall contrast of translucent objects. Figure 9 shows an example of this. Three objects are shown under the same lighting. Torus A is the most translucent, Torus B is of intermediate translucency, while Torus C is relatively opaque. Note that the range of intensities in the images progressively increases from A to C. It seems reasonable, then, that the human visual system might use image contrast to estimate the opacity of an object. In this section, we discuss the role of image contrast in the perception of translucency.
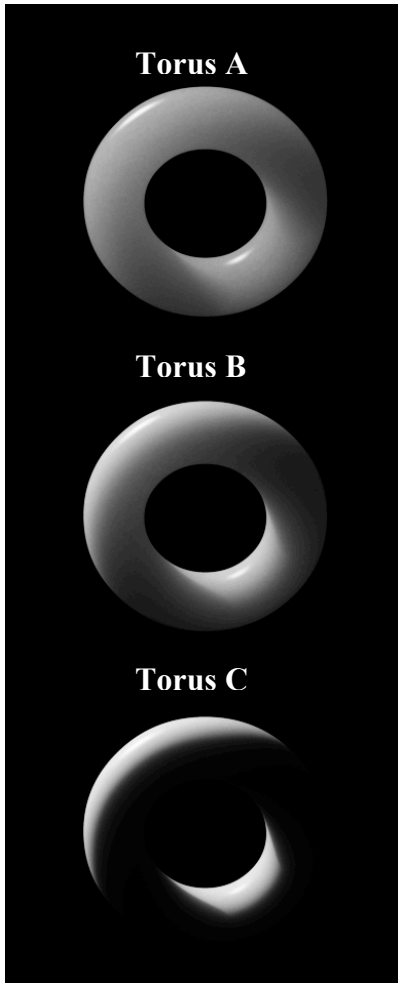

Torus A

Torus B

Torus C

Figure 9. Three tori in increasing order of opacity. Note that the range of intensities in the images increases from A to C.

How should we define contrast? It is common to alter the contrast and brightness of an image using a *linear* (affine) transformation of the image intensities. Specifically, to adjust contrast, intensities are *multiplicatively* scaled, while to adjust brightness, image intensities are *additively* scaled. Can we use the concept of linear transformations to understand the relationship between opaque and translucent objects?

Figure 10 directly compares Torus A with Torus C. If we plot the intensities of the translucent torus on the x-axis, and the intensities of corresponding locations in the opaque torus on the y-axis, we see that the two images are *not* linearly related to one another.
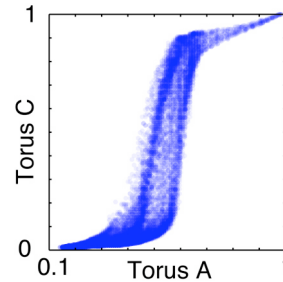

Figure 10. The non-linear relationship between intensity values in Torus A and Torus C.

This is important because it affects our concept of contrast. Specifically, it excludes any definition that assumes linearity. This is highlighted if we take Torus B, and try to adjust the image intensities to make it appear like Torus A or Torus C. If we modify the image intensities according to the best-fitting *linear* transform, we get the results shown in Figure 11.

Torus B linearly adjusted to match contrast of Torus A

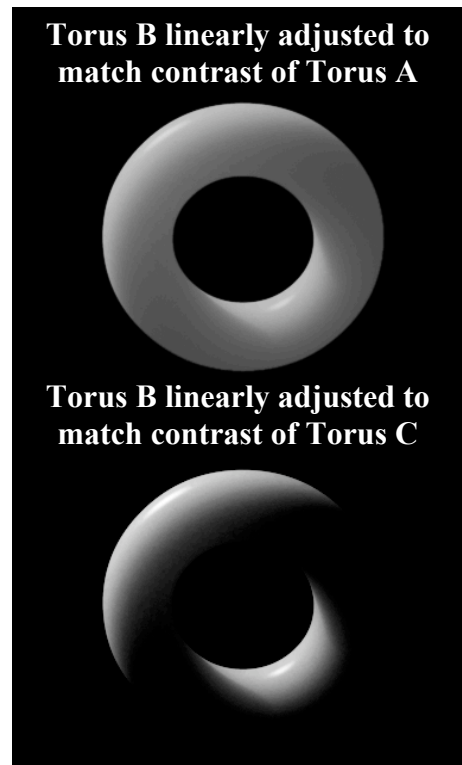**Torus B linearly adjusted to match contrast of Torus C**

Figure 11. Here we linearly adjust the intensities of Torus B to try to make the image appear as translucent as Torus A and as opaque as Torus C.

Although these transforms offer the best compromise between matching the brightness and contrast of the images, we see that they fail to match the apparent *translucency* of the objects. This is especially clear for the low-contrast image, which looks considerably less translucent than the target image (Torus A).

When we gradually alter an object from translucent to opaque, we see that the entire distribution of image intensities changes shape (Figure 12). The mode shifts to lower intensities, while the whole distribution becomes more skewed. This

suggests that our concept of contrast should take into account the entire distribution of intensities, or at least some summary statistics (e.g. mode and skew), that capture these changes.
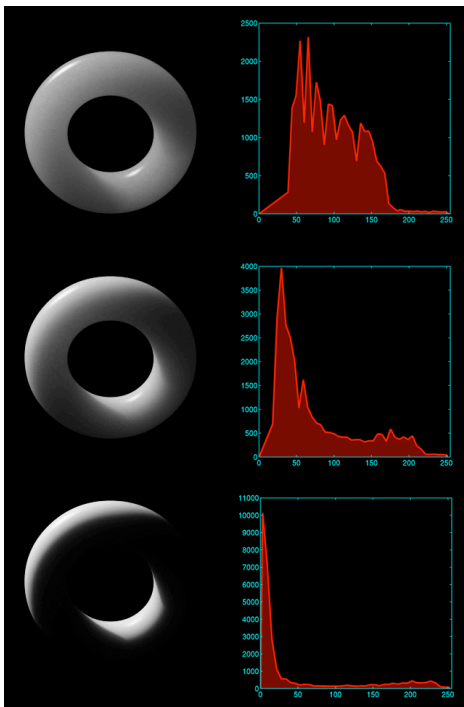

Figure 12. Intensity histograms of the three images. Note that the entire shape of the distribution changes as the torus becomes more opaque.

This is highlighted if we try to match Torus B to the other two images by adjusting the entire distribution of lights and darks. The results of this histogram matching process are shown in Figure 13. Note that the adjusted images look considerably closer to their targets than for the linear transformations in Figure 11.

It is important to note that histogram matching preserves the ordinal relationship between pixels, so the *spatial* structure of the image is not affected by the non-linear transformation. Histogram matching cannot, for example, introduce blurriness to the image, or change the position of highlights. Subsurface scatter does introduce blur to the image, as we discuss below. Thus, it is all the more surprising we can make a translucent object appear opaque simply by changing the intensity histogram.

However, although the intensity histogram captures something important about translucent objects, we should emphasise that it is insufficient alone. For example, in Figure 14, the pixels in Torus A have been scrambled (in a way that keeps the image quite smooth) to create a pattern of random noise with the same intensity histogram as Torus A. Unsurprisingly, this transformation destroys the impression of translucency. It is worth noting that the scrambling also destroys the sense of three-dimensional shape. It is possible that by using a different scrambling procedure that preserves the sense of three-dimensionality, the impression of translucency would persist. Clearly it is not just the distribution of lights and darks, but their spatial relations that inform us that something is translucent.

Furthermore, intensity distributions can, of course, be affected by factors other than the degree of translucency. For example, moving the light source, alters the proportion of the object that is in shadow. This naturally changes ratio of light to dark in the image. As we discussed above, perceived translucency *does* change when the light source moves. Unfortunately, however, these changes cannot be easily predicted solely from the changes in the intensity histogram.
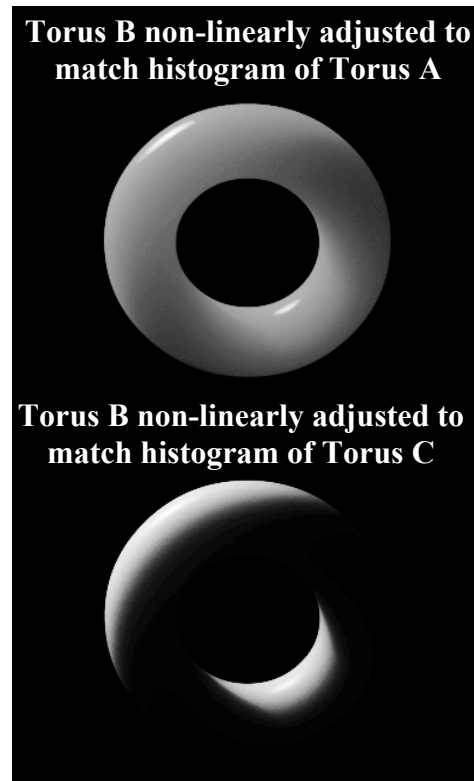

Figure 13. Here we use histogram matching to adjust the intensities in Torus B to make the image appear as Torus A and as opaque as Torus C.

In summary, a simple linear concept of contrast cannot account for the different appearances of translucent and opaque objects. If, instead, we consider the full intensity distribution, we can traverse the space of translucent and opaque objects surprisingly successfully, as long as the lighting conditions are held constant. However, the intensity distribution fails to capture crucial information about the spatial structure of translucent images
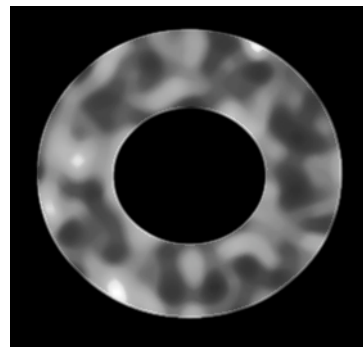

Figure 14 A noise pattern that has the same intensity histogram as Torus A. The image does not appear translucent.

Spatial structure: blur and isophotes

We have argued that it is not solely the proportion of light and dark in an image, but their spatial relationships that makes an object look translucent. In this section we discuss how translucency affects the spatial structure of images.

Subsurface scatter causes light rays to spread out into hazy puffs as they pass through a translucent medium. This has the effect of blurring out sharp details in the image of a translucent object. Consider, for example, the images in Figure 1. Details that are visible in the most opaque Buddha become softened or even invisible in the more translucent objects. Note also the edges of the shadows cast across the tori in Figure 9. These are crisp and pronounced in Torus C, but blurry and diffuse in Torus A. Blur evidently plays a key role in the distinctive soft appearance of translucent materials.

We have found, however, that blur is insufficient on its own to produce a percept of translucency. If the intensity distribution is held roughly constant, adding blur by itself has little effect on perceived translucency (Figure 15). This makes sense given that many other factors can produce blur in images, including depth of field effects and shadow penumbras.
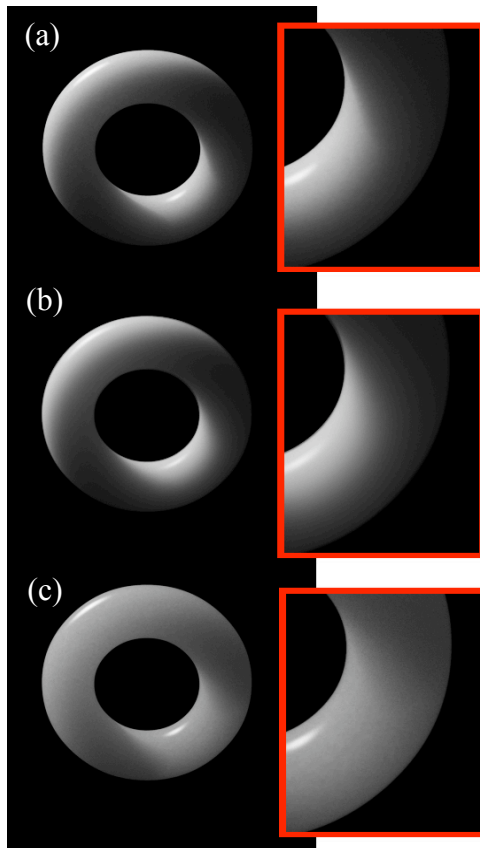

Figure 15. Effects of blur on perceived translucency. (a) Original version of Torus B. (b) Blurred version of Torus B. (c) Blurriness in the more translucent Torus A. Note that the blur has almost no effect on the apparent translucency of (b), although the blur is clearly visible in the close-up.

More generally, it is interesting to study the structure of images of translucent objects. The spatial organisation of lights and darks in an image can be made visible by plotting the *isophotes* (contours of equal luminance), as shown in Figure 16. Note the fact that the isophotes are closer together in Torus B than in Torus A. This is related to the smoothness of Torus A, another way of conceiving of blur.

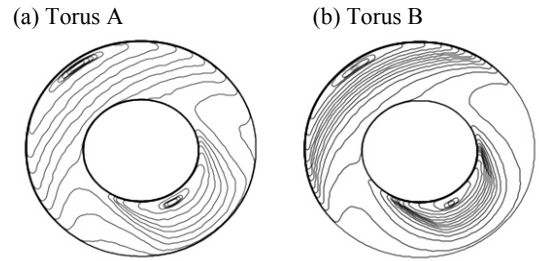(a) Torus A   (b) Torus B

Figure 16. Isophotes of Torus A and B. Note the bunching up around the edge of the shadow in Torus B.

It is important to note that spatial organization alone does not tell us whether an object is translucent or not. This is demonstrated in Figure 17, which shows the photographic negative of Torus A. The photographic negative has exactly the same pattern of isophotes as its positive counterpart, and yet the image does not appear translucent. This is important as it suggests that the visual system attends to the *direction* in which intensity is varying, and not just to the spatial layout of the intensity variations.
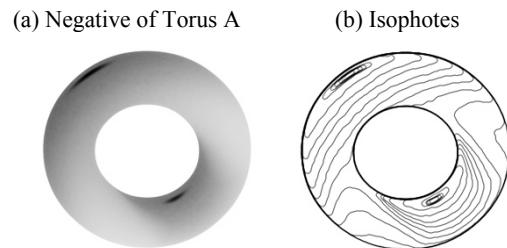

(a) Negative of Torus A   (b) Isophotes

Figure 17. The photographic negative of Torus A has the same isophote pattern, but does not appear translucent

4. Conclusions

What makes wax, cheese or marble look translucent? Here we have used a combination of psychophysics and image analysis to study some of the factors that influence perceived translucency.

To summarize our findings, let us conclude with some suggestions for artists who wish to render translucent materials.

Translucent objects look most translucent when they are glossy and lit from behind. Glossiness also aids the perception of shape, by recovering detail that is lost by the softening effects of subsurface scatter.

If we wish translucent objects to look 'glowing' and 'warm', colour saturation should be positively correlated with intensity. By contrast, if we wish them to look 'icy' or 'dilute', the correlation should be negative.

Translucent objects should be lower contrast than opaque ones. However, the relationship between translucent and opaque versions of an object is generally non-linear. Thus, to portray a translucent object realistically, it is not sufficient simply to reduce

an opaque object's contrast (using Photoshop, for example). It is necessary to modify the entire distribution of intensities.

Sharp cast shadows should be avoided as they make objects appear hard and opaque, while blur and loss of detail gives translucent objects their characteristic soft appearance. However, as with contrast, we cannot make an opaque object appear translucent simply by blurring out the details. Attention must be paid to the global effects of light bleeding through the object if we wish to portray a translucent material.

References

Blake, A. and H.H. Bülthoff: 1990. Does the brain know the physics of specular reflection? *Nature* 343, no. 6254, 165-168.

Blake, A. and H.H. Bülthoff: 1991. Shape from Specularities: Computation and Psychophysics. *Philosophical Transactions of the Royal Society (London)* Series B 331, 237-252.

Chandrasekhar, S. 1960. *Radiative Transfer*. Oxford University Press.

Debevec, P. E. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. Proceedings of SIGGRAPH 98 (Orlando, Florida, July 19-24, 1998). In *Computer Graphics* Proceedings, Annual Conference Series, 1998, ACM SIGGRAPH, pp. 189—198.

Fleming, R. W., Dror, R. O., & Adelson, E. H. 2003. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, *3* (5), 347-368, http://journalofvision.org/3/5/3/, doi:10.1167/3.5.3.

Heyney, L. and Greenstein, J. 1941. Diffuse radiation in the galaxy. *Astrophysics Journal* 93, 70—83.

Jensen, H. W., Marschner, S. R., Levoy, M., and Hanrahan, P. 2001. A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001*, ACM Press / ACM SIGGRAPH, New York. E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Seies, ACM, 511—518.

Jensen, H. W. and Buhler, J. 2002. A Rapid Hierarchical Rendering Technique for Translucent Materials. ACM Transactions of Graphics, *(SIGGRAPH)*.

Koenderink, J and van Doorn, A. 2001. Shading in the case of translucent objects. *Proceedings of SPIE*, *4299*, 312—320.