

COMMENTARY

Perceiving transparency from opaque surface materials

James Todd^{a,1}

It is easy to marvel at the abilities of master painters and sculptors to depict draped cloth, especially in comparison with lesser artists whose depictions of cloth may often appear not quite right. Draped cloth is difficult to represent because its visible appearance is influenced by many factors, including the properties of the cloth itself (e.g., its thickness and stiffness) and also the shape of the object over which it is draped. The depiction of transparent cloth is even more difficult, because the artist must convey the impression of two distinct surfaces (Fig. 1, *Left*). Painters are able to manipulate color and shading to achieve that effect, and the perceptual consequences of those manipulations have been studied extensively (1–4). However, those tools are not available to sculptors. The appearance of transparency in a sculpture must be based entirely on the three-dimensional (3D) shape of the sculpted object. The article by Phillips and Fleming (5) on “The Veiled Virgin illustrates visual segmentation of shape by cause” is a scientific investigation of how human observers are able to interpret the presence of transparency in solid shapes composed entirely of opaque materials.

One of the central hypotheses that guided this research is that the shape of an object can provide information about the morphogenic processes with which it was created. This idea was first proposed over a century ago by the great Scottish naturalist D’Arcy Wentworth Thompson (6), and it has been developed more fully using group theory in a series of recent books by Michael Leyton (7, 8). Phillips and Fleming (5) speculated that this approach might be useful for understanding observers’ perceptions of veiled objects like the ones shown in Fig. 1. If observers can discern the causal origins of different shape features, they reasoned, then it should be possible to identify which features are primarily caused by the veil, and which ones are primarily caused by the underlying object.

A Psychophysical Investigation of Draped Cloth

To test this, they presented untrained observers with an image of a sculpture by Giovanni Strazza called *The*



Fig. 1. Images of objects with draped cloth. (Left) A sculpture of *The Veiled Virgin* by Giovanni Strazza, which appears as a human face with a transparent veil draped over it. **(Right)** A computer simulation of a sphere with an opaque cloth draped over it. Note, in both cases, that it is possible to determine which regions are in contact with the underlying object and which ones are not.

Veiled Virgin (ca. 1850) (5) that is shown in Fig. 1, *Left*. The image was shown on a tablet device, and observers were asked to paint on the tablet to indicate regions where the veil appeared to be in contact with the underlying face, and regions where the veil appeared to be elevated above the face. These two types of judgments were performed on separate screens that could be toggled back and forth. The results revealed that observers are able to make these judgments with a high degree of reliability, thus indicating that they are able to distinguish features that appear to be part of the veil from those that appear to be part of the face.

An important disadvantage of using images of famous sculptures for psychophysical research is that it is impossible to measure the ground truth, so there is no way to assess the accuracy of observers’ judgments. Moreover, after carefully examining this particular sculpture, the authors recognized that it is not a faithful depiction of how a draped veil would behave in the natural environment. What the artist created was a series of alternating regions, some that depict the

^aDepartment of Psychology, The Ohio State University, Columbus, OH 43235

Author contributions: J.T. wrote the paper.

The author declares no competing interest.

Published under the [PNAS license](#).

See companion article, “The *Veiled Virgin* illustrates visual segmentation of shape by cause,” [10.1073/pnas.1917565117](https://doi.org/10.1073/pnas.1917565117).

¹Email: todd.44@osu.edu.

shape of the veil, and others that depict the underlying surface as if it were unobstructed by the cloth. This is particularly noticeable in the eye. I will return to this issue later in the commentary to suggest that this technique may have been essential to produce the effect Strazza was trying to achieve.

To overcome these limitations, the authors performed a second experiment using the same task but with different stimulus materials that they created (5). These stimuli depicted three unfamiliar 3D objects with multiple samples of draped cloth for each one, and a 3D scanner was used to measure the shapes of the underlying objects as well as the structure of each sample of draped cloth. The results revealed that observers' judgments about the contact between the veil and the underlying surface were highly correlated with the physical distance between the veil and the underlying surface. In other words, their judgments were highly correlated with the physical ground truth. Finally, observers were also asked to estimate the depth profiles along horizontal scan lines of the depicted surfaces for both the veil and the underlying object. The results indicate that observers are able to discern that a composite shape has multiple layers and that they are able to judge the individual shapes of each one.

There are several reasons why I believe that this work makes an important contribution to the field. Although there is a substantial literature on the perception of transparency (1–4), there is nothing in that prior research that sheds any light on the phenomenon of *The Veiled Virgin*. What is particularly interesting about this phenomenon is that the basic technique has been known to sculptors for centuries, but Phillips and Fleming (5) have identified its scientific significance. This paper also raises some important issues that are certain to inspire future research. One is to uncover the specific mechanisms by which observers are able to segment the underlying surface from the veil. The authors have done an impressive series of analyses in an effort to address that, but, unfortunately, they did not provide a definitive explanation. Another more general issue includes the ability of observers to recognize cloth materials from their shapes. Given that modeling software is now widely available for simulating draped cloth, this is a problem that is ripe for future investigation.

Despite my high regard for this research, I am not convinced that adopting the language of causal structure is the best way to think about the results. It is interesting to note that the authors did not use that language in their instructions to the experimental participants (5). They did not ask whether surface regions were primarily caused by the veil or the underlying object. What they asked, instead, was whether or not the veil appeared to be in contact with the underlying surface. I think that question is much closer to what the observers actually perceive for these stimuli. The real causal process by which the structure of the veil is created involves a complex optimization of minimizing the overall height of the cloth relative to the ground in response to gravitational forces, and minimizing elastic deformations of the cloth in response to surface tension at points of contact. It is clear that this level of physical causality is not what the authors have in mind when they use the phrase “shape by cause” in the title, but it is not at all clear how that language provides any insight into the phenomena they are investigating.

I suspect it is the case that observers are able to segment the layers in these stimuli because cloth appears fundamentally different from other types of 3D objects. Imagine an opaque object that is mostly contained within a larger opaque object, but some parts of the smaller object stick out from the larger one. Assuming that both objects are made from the same material, it is unlikely

that this would produce the perception of transparency. Why then does this occur for a veil? I believe that the answer to this question is best understood using differential geometry.

The Differential Geometry of Cloth

To develop this point, it is useful to begin with a very brief tutorial (9). At any point on a smooth surface, there are two principal directions of curvature that are always orthogonal to one another: one where the curvature is larger than in any other direction, and another where the curvature is smaller than in any other direction. The product of these two principal curvatures is

The article by Phillips and Fleming on “The Veiled Virgin illustrates visual segmentation of shape by cause” is a scientific investigation of how human observers are able to interpret the presence of transparency in solid shapes composed entirely of opaque materials.

called Gaussian curvature. For planes or cylinders, it has a value of zero; for spherical surfaces, it has a positive value; and, for saddle-shaped surfaces, it has a negative value. The sign of Gaussian curvature is an intrinsic property of surface structure, in that it can be measured without ever having to leave the surface.

Many of the surfaces we observe in the natural environment, like faces, contain regions of both positive and negative Gaussian curvature, and these are always separated by regions with zero Gaussian curvature, although the widths of those regions are often vanishingly small. Draped cloth is quite different, however. Let's first consider a cloth material with negligible elasticity. Although it may appear curved from an extrinsic perspective, its intrinsic structure is homogeneously planar—that is to say, it has zero Gaussian curvature at every point. Because it has negligible elasticity, it can only conform to an underlying surface in limited regions before it is forced to form folds. This is easily confirmed by trying to upholster or giftwrap a basketball. It is simply not mathematically possible without introducing cuts or folds in the wrapping material. As the elasticity of a material is increased, so is its ability to conform to an underlying surface. In that case, the cloth will take on the intrinsic curvature of the underlying surface in regions of contact, but it will remain intrinsically planar everywhere else. Fig. 1, *Right* shows a very simple example of cloth that is draped over a sphere. Note that the cloth appears to have a positive Gaussian curvature where it is in contact with the sphere, but, in the regions without contact, it forms a pattern of ridges with zero Gaussian curvature.

These observations suggest a possible source of information about which regions of cloth are in contact with an underlying surface and which ones are not. Those with a positive or negative Gaussian curvature are likely to be in contact with an underlying surface, whereas those with zero Gaussian curvature are likely to be suspended in air without any surface contact. Phillips and Fleming (5) make a similar point in their discussion. One complication with this type of analysis is that some parts of the underlying object (e.g., the lips, or the bridge of the nose) could have Gaussian curvatures that are close to zero, and those regions could be in contact with the overlying cloth. This suggests that nonzero Gaussian curvature is more diagnostic about contact points than zero Gaussian curvature is about noncontact points.

Of course, the intrinsic curvature of a surface region can only be informative to human perception if observers are able to perceptually distinguish the different types of Gaussian curvature, but there is some psychophysical evidence to suggest that they can (10).

There are also some topological constraints on draped cloth. In general, the regions of contact will all be convex, and concave regions of the underlying surface will be visually obscured. This constraint can be overcome by purposely tucking the cloth inside concavities, but that will not occur if cloth is casually tossed over an object. This may be the reason why Strazza did not attempt to simulate the behavior of real veils in sculpting *The Veiled Virgin*. By creating some local regions that depict the shape of the

underlying surface as if the veil were not present, he was able to provide visual information about the surface concavities that would not have been possible with a more physically accurate depiction.

It is important to keep in mind that all of these observations are completely speculative, and additional research is obviously necessary to see whether any of them have any merit. It should also be pointed out that they were all inspired by reflecting upon Phillips and Fleming's (5) article in PNAS. Although it may be a long time before we fully understand the perception of draped cloth, their insightful research on this topic has taken a crucial step in that direction.

-
- 1 F. Metelli, An algebraic development of the theory of perceptual transparency. *Ergonomics* **13**, 59–66 (1970).
 - 2 F. Metelli, The perception of transparency. *Sci. Am.* **230**, 90–98 (1974).
 - 3 J. Beck, R. Ivry, On the role of figural organization in perceptual transparency. *Percept. Psychophys.* **44**, 585–594 (1988).
 - 4 B. L. Anderson, A theory of illusory lightness and transparency in monocular and binocular images: The role of contour junctions. *Perception* **26**, 419–453 (1997).
 - 5 F. Phillips, R. W. Fleming, The *Veiled Virgin* illustrates visual segmentation of shape by cause. *Proc. Natl. Acad. Sci. U.S.A.*, 10.1073/pnas.1917565117. (2020).
 - 6 D. W. Thompson, *On Growth and Form* (Cambridge University Press, 1917).
 - 7 M. Leyton, *A Generative Theory of Shape* (Springer, Berlin, Germany, 1992).
 - 8 M. Leyton, *Process Grammar: The Basis of Morphology* (Springer, Berlin, Germany, 2012).
 - 9 J. J. Koenderink, *Solid Shape* (MIT Press, Cambridge, MA, 1990).
 - 10 F. Phillips, J. T. Todd, Perception of local three-dimensional shape. *J. Exp. Psychol. Hum. Percept. Perform.* **22**, 930–944 (1996).