Dorée Duncan Seligmann
Avaya Labs

Computer Vision, Image Analysis, and Master Art: Part 2

David G. Stork
Ricoh Innovations

Micah K. Johnson
Dartmouth College

As we saw in the first part of this series,[1] we can apply computer vision and image analysis algorithms to problems in the understanding of art, specifically realist paintings. We saw how new, uncalibrated methods for estimating perspective transformations let us transform and view from different positions objects depicted in realist paintings. Such algorithms let us compare perspectively aligned passages within a single painting and thus test a claim about the working methods of early Renaissance painters, specifically Jan van Eyck in his masterpiece, *Portrait of Arnolfini and His Wife* (1434). In this second part of the series, we turn to another recently developed algorithm, this one for analyzing shadows and lighting for determining the illuminant's location in a depicted scene. We'll see that this, too, helps address problems in the history of art.

Algorithms for inferring illuminant direction

The computer vision community has long studied the problem of estimating the direction of illumination in a scene.[2] Inferring this direction from cast shadows presents little or no problem—all we need is to connect pairs of points, one on an opaque object, the other on its corresponding cast shadow. Multiple such lines should intersect at the illuminant. Note that the shadow point need not lie on a flat surface (such as a floor or wall), but can lie on an arbitrary 3D surface. Thus, a person's fingertip and its shadow cast onto the curved leg of another person define a line that points to the illuminant. Figure 1 illustrates this simple procedure applied to a painting by the Lorrainnese Baroque painter Georges de la Tour (1593–1652). The cast shadow lines generally converge on the candle at the center, a fact to which we shall return.

For many years, most algorithms based on shading and highlights required an assumption that the surface was Lambertian—that is, it scattered light throughout a range of angles the way cloth, wood, or skin do, but not as glass or shiny metal do. Furthermore (and more problematically), these algorithms required knowledge or assumptions of the object in question's 3D form. Spheres and other regular solids are particularly simple cases. Algorithms for inferring the illuminant from images of human faces are fairly accurate, but this is due in large part to the fact that the 3D structure of faces doesn't differ much between faces and can be modeled fairly well. Such research on faces has been driven in large part by the goal of discounting or compensating for the illuminant so as to improve face recognition methods.

If we know, or can assume we know, the 3D form of an object we can compute the orientation of each patch of its surface, described as the direction of the line perpendicular to the surface, called the *normal*. The normal can be in any of the 4π steradians of the solid angle, as can the direction of illumination. That patch's lightness depends on the surface reflectance or *albedo* and the angle between the normal and the illumination's direction. These algorithms estimate the illumination direction that best explains the lightness on a number of surface patches. (Estimating the positions of multiple point sources is a direct extension to these algorithms, but we won't consider them here.)

Editor's Note

We can determine if some photographs have been doctored and then invalidate it as proof that two people were in the same place at the same time. But how true is realist art? In this second article of a three-part series, David G. Stork and Micah K. Johnson describe how they have adapted techniques of forensic analysis of digital images to analyze paintings. By examining light and shadows, they uncover what techniques the artist used to render scenes more "realistic," shedding even more light on problems in art history.

—Dorée Duncan Seligmann

1070-986X/06/$20.00 © 2006 IEEE

Published by the IEEE Computer Society

Figure 1. Cast-shadow analysis of Georges de la Tour's **Christ in the Carpenter's Studio** *(1645), oil on canvas, 137 × 101 cm. Each yellow wedge marks an occluder–shadow pair for light shining from right to left (toward St. Joseph), the angular subtense determined by the accuracy with which the researcher can determine shadow direction. The one shadow direction determined by the far left corner of the beam and shadow on the floor is the least well defined as both the occluder and shadow are in a dark region and close together (magnifying any angular uncertainties). Nevertheless, we include this shadow to show how our integration methods can accommodate an uncertain direction. There are several other cast shadows, but these can't be used because corresponding points on the occluders can't be identified accurately. (Courtesy of Musée du Louvre, Paris.)*

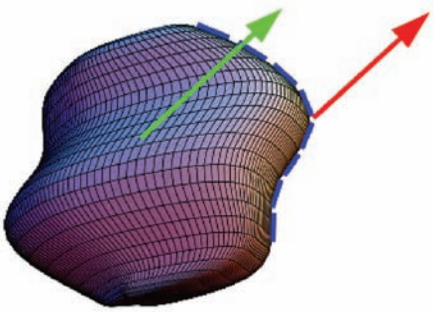

Figure 2. Traditional algorithms for estimating the direction of illumination use a 3D model of the Lambertian surface and yield a 3D vector toward the illuminant (green arrow). The occluding contour algorithm of Nillius and Eklundh uses the lightness function along an occluding contour (blue segments) and yields a 2D vector in the plane perpendicular to the line of sight (red arrow).

Occluding contour algorithm

But suppose we wish to find the direction of illumination when we have few or no cast shadows or little information about the objects' 3D structure in a scene (the general case for a photograph of a complicated object such as a sculpture by Henry Moore and for a typical realist painting)? Nillius and Eklundh had a clever insight into this problem.[3] They realized that for surface patches close to an edge or *occluding contour*, the normal becomes perpendicular to the line of sight (see Figure 2). For instance, each normal to a sphere along its circular occluding boundary is perpendicular to the line of sight and points away from the sphere's center. They also realized that they could restrict consideration of the illuminant direction to the 2π radians of angle illumination perpendicular to the line of sight—that is, in the photograph's plane. Together, these restrictions or constraints make the problem of inferring the direction to the illuminant mathematically tractable.

In brief, their method divides the continuous occluding contour into a number n of short segments. The measured lightness on each segment depends on the unknown surface reflectivity (albedo) as well as the angle between the measured normal direction and the unknown illuminant direction. We choose contours that have a uniform albedo and we assume the illuminant is far away, so the directions toward the illuminant are the same. (Their algorithm can be modified for nearby illuminants as well.) These assumptions reduce the number of variables and make the problem tractable. It's a simple least-squares estimation problem to find the illuminant direction most consistent with the measured lightness along the contour.

Johnson and Farid applied a slight variant of this algorithm to the problem of detecting tampering in digital photographs.[4] Was that photograph of Brad Pitt with Cher that graces the cover of the tabloid at the supermarket checkout counter genuine, or was it created by a digital artist using graphics software such as Adobe Photoshop? If the direction of illumination on Brad

Figure 3. Georges de la Tour, **Magdalen with the Smoking Flame** *(c. 1640), oil on canvas, 117 × 92 cm. (Courtesy of Los Angeles County Museum of Art, gift of the Ahmanson Foundation. Photograph © 2006 Museum Associates/LACMA.)*

differs significantly from that on Cher, then it's quite likely that the photograph was tampered. Indeed, the occluding contour algorithm can often expose such tampering without the need for a 3D model of faces or other objects in the scene.

Finding the illuminant in realist art

As we'll soon see, it's sometimes important to determine an illuminant's location in a scene depicted in a realist painting. Perceptual psychologists have found that most people are somewhat unreliable in judging the direction of illumination in an arbitrary scene or detecting inconsistencies in lighting (we suspect that artists, art historians, and image analysts are more accurate than the general public).

For this reason, we recently applied the occluding contour algorithm to a number of Renaissance and Baroque realist paintings of particular interest to art historians.[5] While we were confident that cast shadows were a reliable source of information about the location of an illuminant, we were unsure whether the abilities of even a master painter were sufficient to shade contours accurately enough for the occluding contour algorithm to work. We applied the occluding contour algorithm to four hand-selected contours in Georges de la Tour's *Magdalen with the Smoking Flame* (see Figure 3): the left knee, crown of the skull, left shoulder, and forehead. Each of the four lines representing the computed illuminant direction, passed close to the candle, and the overall estimate of the illuminant position was at the center of the glass, just beneath the flame itself. This excellent agreement between the computed directions and the candle itself shows that de la Tour executed the work with the candle in the tableau.

We explored the sensitivity of the occluding contour algorithm on *Christ in the Carpenter's Studio* (Figure 1). We artificially replaced the actual measured lightness along the contours in principled ways, ran the algorithm again, and observed the agreement among the computed illuminant directions. If we replace the lightness function by a *constant*, the computed directions are quite haphazard and inconsistent. If we replace the lightness function by a *Gaussian*, with the mean and variance of the measured data, then the computed directions are just partially consistent. But if we replace the lightness function by a two-component spline fit, then the computed directions are quite consistent and agree with the candle's location. In short, it seems that de la Tour could render the lightness along a contour up to the accuracy provided by a spline fit.

Our research might be the first to provide an objective method for judging the fidelity of contour shading in realist art. We hasten to stress that many realist artists exaggerate form, color, and shading to make their art appear more real. We provide an objective measure of one aspect of artistic ability—not an aesthetic one. Even the most committed realist artists will alter color, shading, and form away from a photographic ideal for expressive ends or perhaps to make the painting appear more real.

Integrating direction estimates

Suppose we wish to find the location of the illuminant most consistent with several occluding contour estimates. We make a number of plausible assumptions—concerning, for instance, how the probability depends upon angle—and form a maximum-likelihood estimate of the illuminant location.[6] In this model, the probability that the illuminant is in a particular location is merely the product of the individual probabilities computed from each contour. (There are several subtleties that we can't go into here—for example, methods for incorporating overall confidence of each contour estimate.) We can apply the same probabilistic method to estimates from cast shadows as well.

If we wish to integrate the estimates of the cast-shadow and occluding-contour methods, then we confront a central difficulty. In the probabilistic framework, how do we weigh the relative contributions of the multiple estimates from the different methods? In the simplest naïve Bayesian approach, we merely multiply the probabilities derived from each source (contour or cast-shadow line). While it's unlikely this is the optimal method, we've found that for the half-dozen or so paintings we've explored, these two classes of estimates agree quite well and the location doesn't depend much on which relative weighting we use. Nevertheless, it remains a theoretical and experimental research problem to determine the best method for integrating all estimates.

Painting praxis

How might our methods be used to address problems in art history? In 2001, the celebrated contemporary artist David Hockney came up with a bold theory—actually two theories.[7] Perusing the development of western painting, he was intrigued by the fairly dramatic rise in realism in the *ars nova*, or new art, at the dawn of the Renaissance, around 1430, particularly in the Low Countries of present-day Netherlands and Belgium. The works of Jan van Eyck, Robert Campin, Rogier van der Weyden, and others possessed what Hockney called "the optical look."[7] What was the source of this new realism? Hockney's primary thesis is that, for the first time in history, artists of the early Renaissance saw images optically projected onto a screen or wall and that these provided a new "optical ideal," a new type of image, to be imitated in paint. This is a claim about artistic influence.

It's hard to know how to test this claim given that the evidence is impressionistic and can support a number of explanations. This claim is also weakened by the fact that historians of optics and of art find no documentary records that anyone saw or were influenced by images projected onto a screen at that time.

But Hockney's secondary thesis is in some sense far bolder: He claimed such artists *actually used* optical devices during the execution of their works, specifically projecting images onto their

Hockney's primary thesis is that, for the first time in history, artists of the early Renaissance saw images optically projected onto a screen or wall.

canvas or oak panel "supports," tracing these images and then applying paint.[7] This is a claim about artistic practice, or *praxis*. Here is where computer vision and image analysis can complement more traditional curatorial and art historical methods.[8]

Historians of art and of optics have voiced grave doubt about Hockney's theory of praxis. No evidence exists from that time that shows that artists, mirror makers, or optical scientists had even seen the image of an illuminated object projected onto a screen by a concave mirror or converging lens, let alone traced over such images. Nor is there evidence that the concave mirrors of the time—Hockney's preferred instrument—could project an image of sufficient quality for the process.[9] (Most rebuttals apply to claims about glass lenses as well.) The earliest record we have suggesting anyone might trace such an image was around 1550 from the pen of Giambattista della Porta,[10] and the earliest secure record we have that anyone in fact created a picture using tracing is 1603, when the great scientist and astronomer Johannes Kepler traced the projected image of a sunlit wooded scene. (Alas, this milestone work in the development of image making has not survived.) Nevertheless, Hockney's tracing theory warrants careful technical scrutiny.

Illumination in the tracing theory

One of the key technical difficulties in Hockney's tracing theory is that the images are much dimmer than the scene imaged, by a factor of roughly A/f^2, where A is the area of the lens or mirror and f its focal length. For typical focal lengths and areas of lenses or concave mirrors consistent with other contemporary evidence, the reduction in brightness can be roughly 1/250. Hence, only direct solar illumination or an unusually bright

(a) (b) (c)

Figure 4. Georges de la Tour, Christ in the Carpenter's Studio. *(a) The posterior probability contours of finding a point source based on the evidence in five occluding contours on St. Joseph. (b) The probability based on the seven cast shadows for light shining right-to-left of Figure 1. (c) The probability based on all evidence, here computed using the simple product of the probabilities given in (a) and (b). The full information tightly constrains the illuminant to the candle's position, not to the position of Christ and certainly not outside the picture. (Courtesy of Museé du Louvre, Paris.)*

collection of oil lamps could yield an image sufficiently bright for the artist to see and trace.[11]

If we leave the early Renaissance and move much later, into the Baroque, perhaps we might find some of these artists' traced projections. (There is secure evidence that much later, in the 18th and 19th centuries, painters such as Canaletto, William Hyde Wollaston, and others used projections.[12]) Indeed, Hockney proposes that Georges de la Tour traced projected images during the execution of his dramatic nocturne paintings. But then there is still the problem of inadequate illumination, which Hockney addresses thus (p. 129):

> In France, the most famous of all Caravaggio's followers was probably Georges de la Tour, who produced a number of stunning candlelight scenes. Take this painting, *St. Joseph the Carpenter* of 1645, for example. Could that candle really produce all that light? Once again, the source of light seems to be outside the picture, as it must be if you use optics—if the source of light were in the setting, it would cause flare in the lens. Joseph and the girl [sic] were probably painted separately, each lit by a shielded light source in place of the other figure.[7]

Hockney is referring to the painting as it appears (not an invisible underdrawing or preliminary sketch) and shows two half paintings to illustrate the tableau during the two projections he believes de la Tour used.

We applied the cast-shadow and occluding contour methods to check Hockney's claim about the illuminant's location.[5] Figure 4 shows some of our results. Figure 4a is the estimate based on the occluded contour algorithm, Figure 4b is based on the cast shadows, and Figure 4c is based on both methods. Note that this final estimate is spatially constrained and close to the candle—neither in the position of Christ, nor outside the picture. The same analysis, flipped and applied to Christ, also shows the candle as the most likely position of the illuminant for the other purported "exposure." Other shadow evidence also argues against the use of projections—for instance, the fact that the cast shadows are crisp, indicating a small source (like a candle) and not a large one (many bright lamps).

It thus appears that the (dim) candle was indeed the illumination source, and hence that de la Tour used neither very bright sources nor optical projections. A similar analysis of *Magdalen with the Smoking Flame* (see Figure 3) and other of this master's works argue against his use of projections.

Future directions

We can make several improvements and extensions to the model-free algorithms we've described. The first is to find a more principled or experimental model of the angular error in each occluded contour estimate. The second is to find improved methods for integrating estimates from multiple sources, such as weighing different estimates. Another is to derive improved methods for estimating the distance to the illuminant from a single contour. And finally, we need principled methods for integrating the estimates from multiple sources, particularly when some are occluding contours and others cast shadows.

But what other problems in art history might profit from the use of our methods? Recall the problem of the possibly tampered photograph of Brad Pitt and Cher. A related problem arises in paintings done collaboratively by masters and apprentices. Which figure in a painting was painted by whom? Perhaps the methods we described can help detect inconsistent lighting on different subjects, thus indicating that they might have been painted at different times and in different studio conditions and thus possibly by different artists.

Image analysts must learn enough about the problems and methods of art historians to develop software tools that make a genuine contribution to problems in art history and art conservation. Someday, these software tools might be so powerful and useful that the next generation of art historians will use them. **MM**

Acknowledgments

We thank Hany Farid for insight and support. Micah K. Johnson was supported by gifts from Adobe Systems and Microsoft and through a grant (2000-DT-CX-K001) from the US Department of Homeland Security, Science and Technology Directorate. (Points of view in this article are those of the authors and do not necessarily represent the official position of the US Department of Homeland Security or the Science and Technology Directorate.) We thank the Musée du Louvre, Paris, and the Los Angeles County Museum of Art for permission to reproduce *Christ in the Carpenter's Studio* and *Magdalen with the Smoking Flame*, respectively.

References

1. D.G. Stork, "Computer Vision, Image Analysis, and Master Art: Part 1," *IEEE MultiMedia*, vol. 13, no. 3, 2006, pp. 16-20.
2. A. Pentland, "Finding the Illuminant Direction," *J. Optical Soc. of America*, vol. 72, no. 4, 1982, pp. 448-455.
3. P. Nillius and J.-O. Eklundh, "Automatic Estimation of the Projected Light Direction," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition* (CVPR), IEEE CS Press, 2001, pp. 1076-1083.
4. M.K. Johnson and H. Farid, "Exposing Digital Forgeries by Detecting Inconsistencies in Lighting," *Proc. ACM Multimedia and Security Workshop*, ACM Press, 2005, pp. 1-9.
5. D.G. Stork and M.K. Johnson, "Estimating the Location of Illuminants in Realist Master Paintings: Computer Image Analysis Addresses a Debate in Art History of the Baroque," to be published in *Proc. Int'l Conf. Pattern Recognition* (ICPR), IEEE CS Press, 2006.
6. R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification*, 2nd ed., Wiley, 2000.
7. D. Hockney, *Secret Knowledge: Rediscovering the Lost Techniques of the Old Masters*, Viking Studio, 2001.
8. D.G. Stork, "Optics and Realism in Renaissance Art," *Scientific American*, vol. 291, no. 6, 2004, pp. 76-83.
9. S. Schechtner, "Between Knowing and Doing: Mirrors and Their Imperfections in the Renaissance," *Early Science and Medicine*, vol. 10, no. 2, 2005, pp. 137-162.
10. G. della Porta, *Magiae Naturalis, Sive de Miraculis Rerum Naturalium Libri IIII* [On the Wonders of Natural Things], Apud Matthiam Cancer, Naples, 1558.
11. D.G. Stork, "Illumination and Color in the Hockney Theory of Renaissance Painting Praxis: A Critical Evaluation," *Proc. IS&T/SID 12th Color Imaging Conf.*, Soc. Information Systems & Technology/Soc. Information Display, 2003, pp. 11-15.
12. J.G. Links, *Canaletto*, Phaidon Press, 1999.

Readers may contact David G. Stork at artanalyst@gmail.com.

Contact Artful Media editor Dorée Duncan Seligmann at doree@avaya.com.