Inferring illumination direction estimated from disparate sources in paintings: An investigation into Jan Vermeer's *Girl with a pearl earring*

Micah K. Johnson,[a] David G. Stork,[b] Soma Biswas[c] and Yasuo Furuichi[d]

[a]Department of Computer Science, Dartmouth College, Hanover NH 03755
[b]Ricoh Innovations, 2882 Sand Hill Road Suite 115, Menlo Park CA 94025
and Department of Statistics, Stanford University, Stanford CA 94305
[c]Department of Electrical Engineering, University of Maryland, College Park MD 20742
[d]Kanagawa, Japan

ABSTRACT

The problem in computer vision of inferring the illumination direction is well studied for digital photographs of natural scenes and recently has become important in the study of realist art as well. We extend previous work on this topic in several ways, testing our methods on Jan Vermeer's *Girl with a pearl earring* (c. 1665–1666). We use both model-independent methods (cast-shadow analysis, occluding-contour analysis) and model-based methods (physical models of the pearl, of the girl's eyes, of her face). Some of these methods provide an estimate of the illuminant position in the three dimensions of the picture space, others in just the two dimensions of the picture plane. Our key contributions are a Bayesian evidence integration scheme for such disparate sources of information and an empirical demonstration of the agreement, or at least consistency, among such estimates in a realist painting. Our methods may be useful to humanist art scholars addressing a number of technical problems in the history of art.

Keywords: illumination estimation in art, Jan Vermeer, *Girl with a pearl earring*, computer graphics, Bayesian inference, occluding-contour algorithm, shape from shading, ocular highlights, computer vision and art

1. INTRODUCTION

The problem of estimating the direction of illumination or the location of the illuminants in a digital image or video sequence is well studied in computer vision and pattern recognition.[1] Some of the most advanced methods are used as intermediate steps in three-dimensional shape-from-shading algorithms.[2] Related methods are used when removing shadows from digital images of natural scenes.[3] Recently, the problem of estimating illumination direction arose in a publication by Hockney. In brief, the question arose as to the location of illumination source(s) in the Lorrainese master Georges de la Tour's *Christ in the carpenter's studio*, specifically whether it was indeed at the position of the candle depicted in the tableau or instead "in place of the other figure." [4, page 129] To this end, Stork performed a simple cast-shadow analysis of this painting and others in de la Tour's nocturne œuvre.[5] Shortly thereafter, Stork and Johnson applied the model-free occluding-contour algorithm of Nillius and Eklundh to this painting[6] as well as to de la Tour's *Magdalen with the smoking flame*.[7,8] Most recently, Stork and Furuichi built an elaborate computer graphics model of the tableau in *Christ in the carpenter's studio* and adjusted the locations and brightness of the virtual illuminant until the rendered scene matched the painting most closely; in this way they estimated the location of the illuminant in this painting.[9] Each of these studies showed that the illuminant in that painting was indeed quite close to the depicted candle, thereby rebutting Hockney's claim about artistic praxis starting in the early Renaissance, at least for that painting.

When any algorithm for determining the illumination direction is applied to a given image there will be errors and uncertainties in the estimates—even when applied to a high-quality digital photograph—and the question arises how to integrate these estimates in a principled way. Such errors and uncertainties are likely to be much larger in a hand-executed realist painting, of course. Discrepancies between the directions inferred from different

Send correspondence to Micah K. Johnson, kimo@cs.dartmouth.edu.

visual evidence (e.g., cast shadows versus highlights) might shed light on the particular artist's working methods (Sect. 5). We applied a range of model-free and model-based methods for estimating illuminant position to a realist portrait by the Dutch master Jan Vermeer, who is well-known for executing highly realistic interiors, portraits and genre paintings.[10] We analyzed his *Girl with a pearl earring* (Fig. 1) in large part because it displayed a range of types of evidence about lighting direction, but this evidence here is imprecise and hence integration of such estimates is necessary to reduce overall uncertainty of the estimate.

Figure 1. Jan Vermeer, *Girl with a pearl earring* (c. 1665–1666), 44.5 × 39 cm, oil on canvas, Royal Cabinet of Paintings Mauritshuis, The Hague. The lighting information is contained in the cast shadow of her nose, the highlights on her pearl and eyes, the subtle variations in lightness along the occluding contour of her right cheek and chin, and the overall pattern of form shadows on her face and body. Pearls appear in several Vermeer's paintings—in fact eleven of his women wear them—and art scholars believe that the size of the pearl here may have been exaggerated by Vermeer.

We begin in Sect. 2 by describing the model-free methods for estimating the location of the illuminant: cast-shadow analysis and occluding-contour analysis. In Sect. 3 we describe the model-based methods, which for this painting include models of the pearl and the girl's eyes. We also applied generic methods for inferring illumination direction from a frontal face, and altered and extended them to apply to the face in three-quarters view in Vermeer's painting. In Sect. 4 we describe a Bayesian theory of evidence integration and apply it to the estimates from previous sections. A special consideration here is that some methods provide two-dimensional estimates of the location, while others provide three-dimensional estimates. We conclude in Sect. 5 with a brief discussion of how these methods may be of use in connoisseurship and humanistic scholarship of the visual arts.

2. MODEL-FREE METHODS

We first consider two model-free methods, that is, methods that are based on the image information and little or no explicit modelling of the scene: cast-shadow analysis and occluding-contour analysis.

2.1 Cast-shadow analysis

Cast-shadow analysis of an image is conceptually simple in the case of a single point-source (or small) illuminant. One merely draws a line from a point on a cast shadow through its partner occluding point. This line, if extended, will pass through the two-dimensional location of the illuminant. Multiple such lines (in general form), if extended, should meet near the illuminant.

It is not necessary that the shadow point lie within the image or be visible to be useful, as geometric constraints can limit the illuminant's location, as shown in [11, Figure 2]. Moreover, one can estimate the direction given constraints, as in [11, Figure 4]. This was the method used by Stork to estimate the location of the illuminant in Georges de la Tour's *Christ in the carpenter's studio*; although the actual shadow of Christ's knee lies just outside the frame of the painting, the shadow of his (straight) shin restricts the possible angles well enough to judge the technical question at hand.[5,7]

The best defined cast shadow in *Girl with a pearl earring* is from the lower tip of the center of her nose and its partner shadow point on her cheek. Note that there are some uncertainties due to identifying the exact location of partner points, as shown in Fig. 2, but we find $151° \pm 2°$.

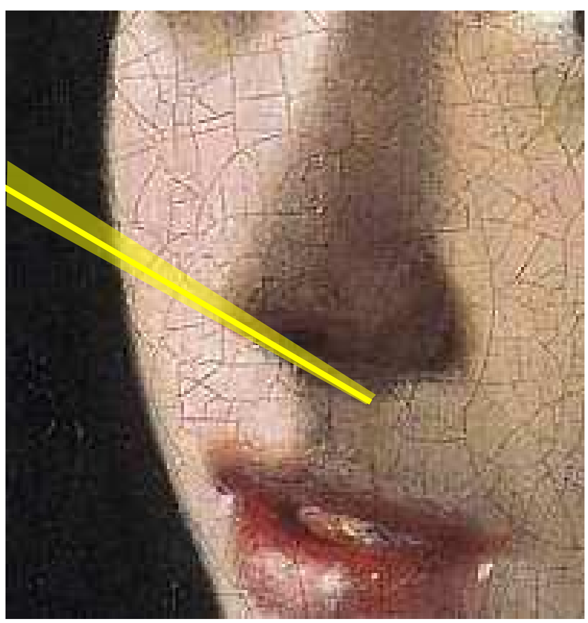

Figure 2. Jan Vermeer, *Girl with a pearl earring* (c. 1665–1666), detail approximately 8.0 × 7.5 cm, Royal Cabinet of Paintings Mauritshuis, The Hague. The direction toward the illuminant based on the cast shadow of the lower tip of the girl's nose along with with its associated error range is marked in yellow.

Note too that there are numerous shadows in this painting that are not cast shadows and thus cannot be exploited as just described. *Attached shadows* or *form shadows*, such as on her smock and left cheek, were not cast by points on *other* objects and thus there is no line defined from a point on the shadow to its partner occluding point. Moreover, such form shadows are generally ambiguous; the same form shadow on a given object can be produced by an illuminant in any of an infinite number of positions and angles.[9] In some cases, however, sets of such shadows can be exploited through model-based methods, such as in Sect. 3.4.

2.2 Occluding-contour analysis

Nillius and Eklundh showed that the lightness along an occluding contour could be exploited to find the direction of illumination in a model-free way.[12] Three conditions had to obtain for their algorithm to yield meaningful estimates:

1. the surface must be of constant reflectivity or albedo
2. the surface must be diffusely reflecting or Lambertian (like cloth or skin, not like glass or metal)
3. the small illuminant must be reasonably distant from the contour so that its light rays are nearly parallel at the contour and the irradiance is nearly the same at each of the patches.

Figure 3 illustrates the basic theory. While the normal vector perpendicular to an arbitrary point on a surface is generally unknown, the normal *is* known along an occluding contour: it is perpendicular to the line of sight and to the (one-dimensional) contour itself (red arrows). The lightness at a short patch L_i along the contour depends on the (unknown) surface albedo ρ, the (unknown) irradiance at the patch, and the (unknown) angle θ_i between the known direction of the normal at i and the unknown direction toward the illuminant.[13] (This angle θ is measured in the plane perpendicular to the line of sight.) The measured light in each patch represents a single equation in three unknowns (ρ, θ_i, L) with suitable normalization, known as Lambert's cosine law:[14]

$$L_i = \rho \max[\mathbf{n}^t \mathbf{s}, 0] = \rho \max[|\mathbf{n}|\ |\mathbf{s}| \cos\theta_i, 0] = \rho \max[\cos\theta_i, 0]. \tag{1}$$

A triple of variables in a single instance of Eq. 1 is of course underdetermined. However if there are three or more patches at a variety of orientations (i.e., angles to the illuminant), then the variables can be determined through a least-squares fit, in particular the k variables θ_i. Once these angles are known the direction to the illumination can be computed easily. This estimated direction is in the plane of the picture (perpendicular to the line of sight), and is thus expressed as a single angle θ with respect to the horizontal, our convention here.

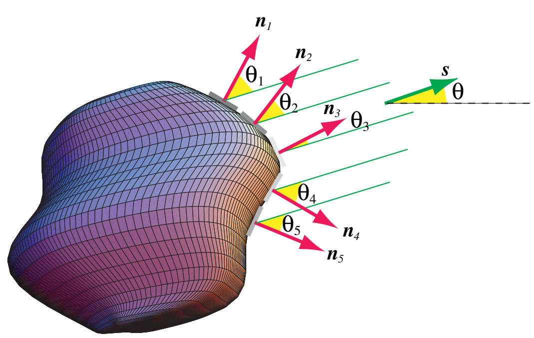

Figure 3. The occluding-contour algorithm takes the measured lightnesses of several patches along an occluding contour and determines the direction of a distant illuminant. Here $k = 5$ patches are shown, along with their known normals (red) and the unknown direction of illumination (green). The angles between the normals and unknown direction of illumination are labelled θ_i. The occluding-contour algorithm takes the measured lightnesses in each patch and the orientations of each normal and computes a best fit (in a least squares sense) to the surface albedo ρ, the irradiance at the surface, and the k unknown angles θ_i ($i = 1, \ldots, k$). Once these angles are known, the direction to the illuminant, θ, can be computed.

Johnson and Farid used this algorithm to expose tampering of digital photographs.[15] Stork and Johnson first applied the algorithm to analyze lighting in realist paintings;[6–8] their sensitivity analysis provided the first objective computer-based test of the ability of realist painters to shade surfaces accurately.

Figure 4 shows occluding contour applied to hand-selected contours along the cheek and chin in *Girl with a pearl earring* (green); the direction toward the illuminant is shown by the yellow arrow.

Figure 4. Occluding-contour algorithm applied to the cheek and chin in *Girl with a pearl earring*, marked by the green contour lines. The best fit angle of the illumination direction is $\theta = 149°$ (measured counter-clockwise with respect to a rightward, horizontal vector), as shown by the yellow arrow. The standard deviation of this estimate was computed from an analysis of the effects of random perturbations to the variables and found to be $\sigma = 4°$.

3. MODEL-BASED METHODS

We now turn to methods for estimating the direction of illumination that are model dependent. Such methods, by definition, must incorporate information or assumptions beyond that provided by the raw image, for instance assumptions about the axial symmetry of the pearl, or generic form or symmetry of human eyeballs, or parameterization of the three-dimensional form of human faces. The most under-constrained model is of the full scene itself, and these can be created using tools from computer graphics.

3.1 Reflections from the pearl

We created a computer graphics model of the pearl by first reading the two-dimensional positions of a dozen points selected along its left contour in a high-resolution scan of the painting, and then fitting a fourth-order polynomial to these points to describe the (planar) contour shape. We then computed the surface of revolution for the wireframe model. Next we texture mapped a white surface to the wireframe, and then adjusted a specular reflection component to the surface (Fig. 5).[13] Finally, we placed the earring in a full computer graphics model (see Sect. 3.4, below) and adjusted the location of the virtual source so as to match the painting. We iteratively adjusted the location of the illuminant and specular component several times, until the image best matched the painting.

We computed an error range for our estimate by adjusting the location of the virtual light away from the perceived ideal until the rendered image did not "look proper." This was, admittedly, a somewhat subjective determination, and was somewhat arbitrarily assigned to be two standard deviations from the mean (cf., Sect. 5). This then gave us an estimate of the error range on this method in this case (cf., Sect. 4).

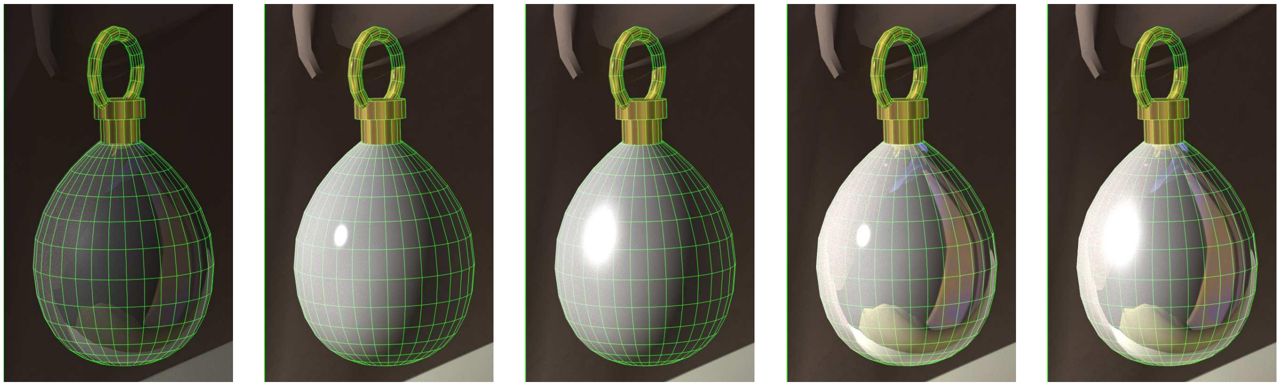

(a) (b) (c) (d) (e)

Figure 5. We created a computer graphics model of the pearl using *Shade 9 Professional*, successively refining the model parameters until the rendered pearl matched that in the painting as closely as possible: (a) the wireframe model in very low illumination, without spot illumination, (b) spotlight but no specular reflection component in the pearl, (c) specular reflection coefficient added, (d) reflections of shoulder and head, but no specular adjustment, and (e) final, full model, with specular coefficient adjusted. The direction of illumination shown the same throughout these images was found through trial and error, as judged by eye, to be $\theta = 155° \pm 4°$.

3.2 Specular highlights on the eyes

Specular highlights on the eye provide an estimate of the direction of illumination. The eye is not a simple sphere, of course, and it is possible to get two specular highlights—one off the cornea, one off the sclera. Johnson and Farid recently developed an algorithm to automatically estimate the direction of illumination from reflections off the cornea, primarily for detecting tampering in digital photographs: if the direction of illumination estimated from one subject's eyes differs significantly from the estimate from another subject's eyes, then it is likely that one of the subjects was digitally composited into the scene.[16]

We selected the iris region by hand, then applied a simplified version of the Johnson/Farid algorithm to the eyes in *Girl with a pearl earring* (Fig. 6). This simplified version assumes the girl's eyes are in different depth planes and are looking along the z-axis of the camera. With this assumed pose, the limbii appear as circles rather than as ellipses. This simplification allows the algorithm to estimate the "focal length" directly from the image and yields separate estimates of the illumination direction from each eye. An estimate of the *position* of the light source is computed by triangulating the estimates from each eye and the final illumination direction is the vector from the midpoint between the eyes to the triangulated light position.

The two-dimensional angle toward the illuminant computed by the algorithm was estimated to be $149.6° \pm 1.22°$ degrees measured with respect to the rightward horizontal direction. The error range was estimated by randomizing the center and radii of the selections by one pixel and running the algorithm on 50 randomized selections. The position of each highlight was not randomized and was estimated as the center of mass of the full region containing the highlight.

3.3 Facial template-based illumination estimation

A significant challenge in automatic face recognition is ensuring invariance to illumination, specifically illumination direction. To this end, several face-recognition algorithms estimate the illuminant as an intermediate step, so as to discount it in the ultimate recognition. Such methods work best for frontal views, where spatial symmetry of the face and templates provide a fair approximation to most faces.

Figure 7 shows the steps of our method. We took a generic average three-dimensional face model and aligned it by hand to match the pose of the girl in Vermeer's painting. Next, we simulated the image formation to get an initial estimate of the illumination. Specifically, we modelled the facial surface as Lambertian, and the diffuse component of the surface reflection used Lambert's cosine law of Eq. 1 above. As we saw in Sect. 2.2 solving for the surface normals, albedo and light source direction from a single image is ill-posed and we need further

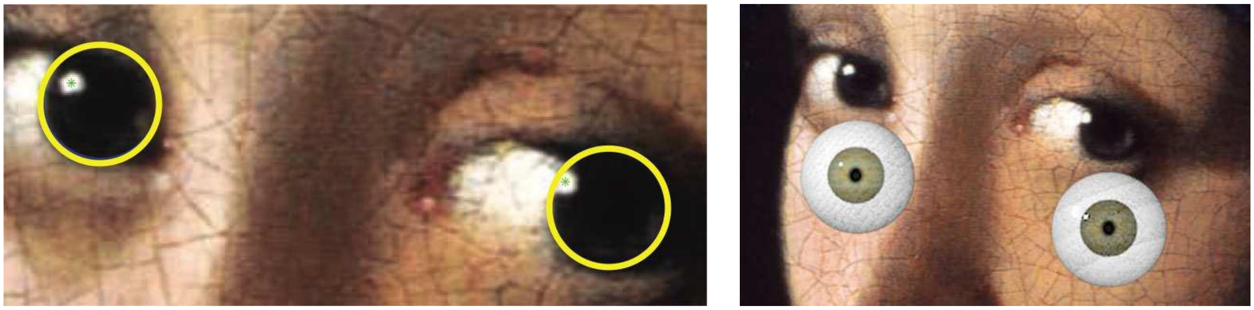

Figure 6. The left part of the figure shows the algorithms's *circular* fit to the iris region, because the girl is looking straight toward the viewer. Each highlight from the cornea was found automatically, as center marked with small green dot. The right part of the figure shows the best fit direction of gaze for the eyes, and associated specular highlights.

assumptions. Here we proceed as follows. We use an average facial model with contant albedo for the skin. The best overall lighting direction is a weighted average over pixels, given by

$$\hat{\mathbf{s}} = \left(\sum_{ij} \mathbf{n}_{ij} \mathbf{n}_{ij}^t \right)^{-1} \sum_{ij} \mathbf{L}_{ij} \mathbf{n}_{ij}. \tag{2}$$

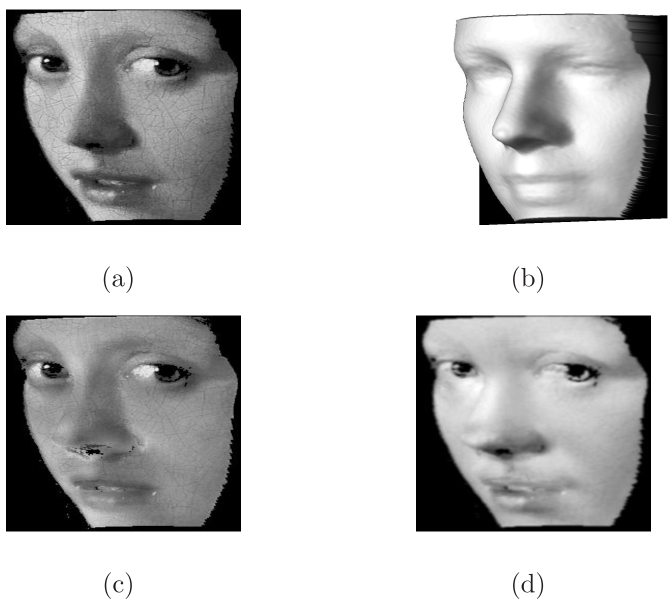

Figure 7. (a) Grayscale image of Vermeer original. (b) Aligned "average face" three-dimensional model. (c) Albedo estimate throughout the face. (d) Albedo-free normalized image.

The constant albedo assumption is merely an initial approximation, of course; any face will have subtle variations in albedo.[17] One can then iterate, refining the albedo model, then the normals (face model), until convergence—a general form of an expectation-maximization algorithm.[18] The angle of the illuminant was in world coordinate $(-0.2778, 0.1, 0.9554)^t$ which corresponds to $160°$, measured with respect to the horizontal rightward direction.

3.4 Computer graphics modelling

The most extreme, indeed complete, model of the painting would involve not only a complete three-dimensional model of the subject, dress, and illumination, but also media and some representation of the artist's working methods and even intentions. Clearly, creating such a complete model is not only impractical (if not impossible), but also unnecessary when addressing the art historical questions at hand. We assume Vermeer executed the portrait from a live model and sought to render reasonably faithfully what he saw—or more specifically, that he had no biases or intentions to misrepresent the direction of illumination in his numerous visual cues.

We created a computer graphics model of Vermeer's subject, her dress, and pearl in *Shade 9 Professional*. The creation of a three-dimensional model from a single two-dimensional picture or projection is formally ill-posed, and is thus part art, part science. We incorporated many natural but unprovable constraints, for instance that the color of her gown, of her headband, and of her yellow scarf were each uniform, that the illumination was from a single, fairly small, distant source with low ambient (omni-directional) component, her face was nearly symmetric, and so on.

Figure 8 shows our computer graphics model with the optimal direction of illumination and nearby directions. The judgement of "best" is admittedly somewhat subjective, but for this painting involved most notably the form shadows on the subject's left cheek and shoulder and the shading of her nose.

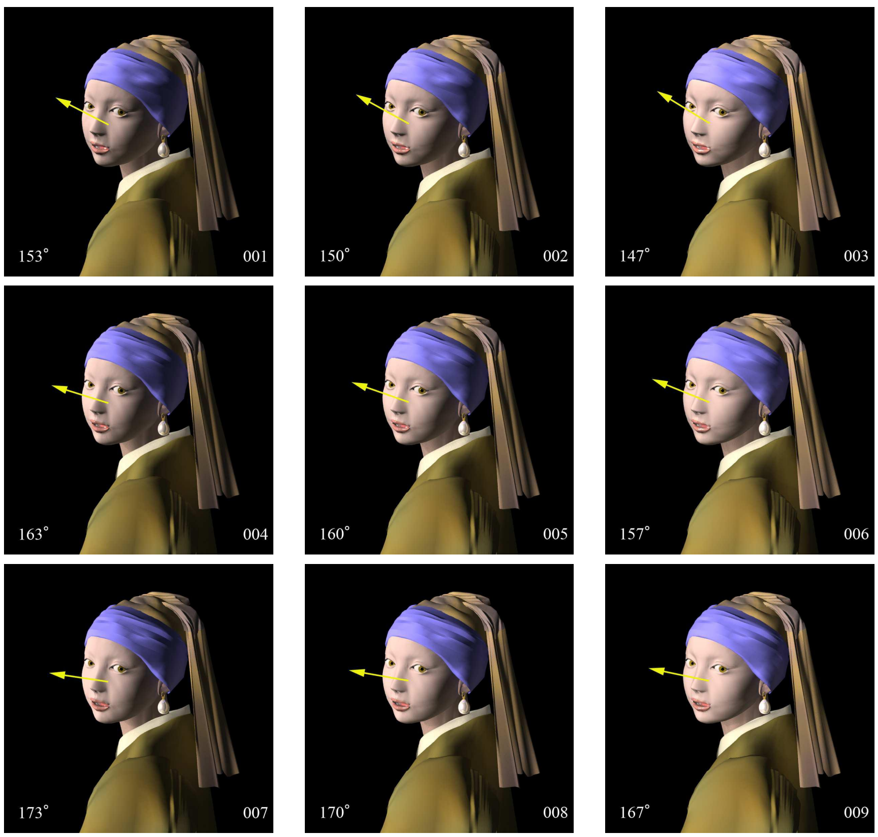

Figure 8. Computer model rendered in *Shade 9 Professional*. The image judged to best match the painting is shown at the center (160°) and just-noticeable deviations differ by roughly 5°, as shown.

4. INFORMATION INTEGRATION

We now consider the method for integrating multiple sources of information. We let S_i be the i^{th} source of information, such as from a single cast shadow, or a single occluding contour, and $p(\mathbf{x}|S_1,\ldots,S_k)$ the probability density the illuminant is in position $\mathbf{x}$ given such information. The "position" $\mathbf{x}$ could be in three, two, or one dimensions, i.e., perpendicular to the line of sight (the case we consider).

4.1 Theory

A full Bayesian theory would assume that the probability density of finding the illuminant in position $\mathbf{x}$ based on k sources $S_1, \ldots, S_k$ would be written $p(\mathbf{x}|S_1, \ldots, S_k)$ and could depend upon an arbitrary function of the evidence. We assume, though, the naive Bayesian model and write the density as: [18, Section 2.12]

$$p(\mathbf{x}|S_1, \ldots, S_k) \propto \prod_{i=1}^{k} P_i p(\mathbf{x}|S_i), \tag{3}$$

where P_i represents the prior probability (reliability) of source S_i, and we omit the obvious normalization condition. Equation 3 relies on the naive Bayes assumption that the sources are approximately independent, an assumption that works surprisingly well in a wide range of problems in pattern classification.[18]

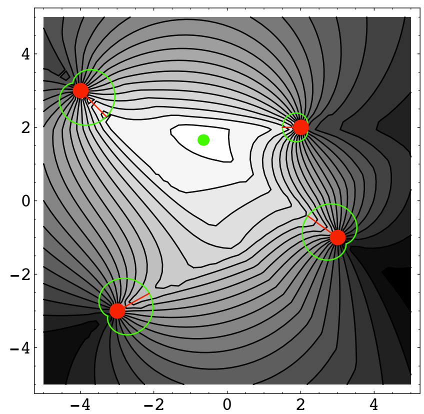

Figure 9. Contours in two dimensions of equal posterior probability for four sources of information about the lighting position, for instance an occluding contour, a cast shadow centered on the red dots. The maximum direction is shown by the short red lines, whose lengths indicate the prior probabilities P_i in Eq. 3. The angular dependency is modeled as a Gaussian function of angle from the this maximum direction, as shown by the green cardiods. The naive Bayesian integration theory. The maximum a posteriori location, $\hat{\mathbf{x}}$ in Eq. 4, here is found analytically and marked with a green dot. In arbitrary cases, there may be several local maxima, and the optimal must be found by gradient descent search. Multiple such maxima might indicate the presence of multiple light sources or systematic biases and differences between the visual evidence in the artist's rendering.

If the estimates are unbiased and informative, then under plausible statistical assumptions the estimated location will become more restricted around the true value the larger the number of information sources. This naive Bayes assumption may or may not be appropriate for a particular painting, and must be explored experimentally more fully (cf., Sect. 5). A full Bayesian theory would include a priori information, for instance that the light is more likely to come from the left than the right, or above than below, and so on (Fig. 9).[18] The maximum a posteriori position is of course

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} p(\mathbf{x}|S_1, \ldots, S_k), \tag{4}$$

and is marked by the green dot in Fig. 9. The analogous case for the one-dimensional orientation perpendicular to the line of sight (cf., Fig 3) is

$$\hat{\theta} = \arg\max_{\theta} p(\theta|S_1, \ldots, S_k). \tag{5}$$

4.2 Application to *Girl with a pearl earring*

We assume equal priors P_i in Eq. 3 and multiply all the estimates, above, and find that the final illumination direction is $155° \pm 3.8°$, as summarized in Table 1.

Method	$\hat{\theta}$	σ
Cast-shadow (nose)	151°	2°
Occluding-contour (cheek)	149°	4°
Physical model (pearl)	155°	4°
Specular highlights (eyes)	150°	1.2°
Shape-from-shading (face)	160°	5°
Computer graphics (total)	160°	5°
Final	**155°**	$\approx$ **3.8°**

Table 1. The estimated light directions by various methods and the final pooled estimation assuming each source of information is equally informative. Although we did not compute the standard deviation for the shape-from-shading method, we use here values typical for other research on digital photography. Also, the estimate of the final σ is based on the assumption that the estimates are independent, which is likely not the case (see text).

The final results stated assumed that the direction estimates were independent and equally informative, that is, they were weighted equally in the computation of the mean. While the estimates from cast shadows and from the lightness along occluding boundary are likely independent, the estimates from the physical model of the pearl and of the face as a whole are likely not. Further exploration of the errors and statistical independence of different methods will be needed to refine our method. Regardless, it is satisfying indeed that the various methods yield estimates that are as commensurate as we find (Fig. 10).

5. CONCLUSIONS AND FUTURE DIRECTIONS

Our results show that realist artists, at least one of the extraordinarily high caliber of Jan Vermeer, can render a portrait such that the lighting information from disparate sources is quite coherent and commensurate. This result, taken alone, is hardly surprising—after all, Stork and Johnson showed that Georges de la Tour rendered cast shadows and highlights with high consistency.[6,7] It is perhaps more noteworthy that the wide range of computer methods—developed separately for applications to digital imagery, implemented here separately—can agree so well on such disparate sources of information within a hand executed painting.

Our work has addressed the estimation of the direction of the illumination perpendicular to the line of sight; that is, this direction could be described by a *single* angle with respect to the horizontal. Model-based methods such as those based on computer graphics reconstructions may provide positions in the three-dimensional picture space, exploited by Stork and Johnson and an area worthy of further research.[6]

There are a number of technical statistical matters that should be explored, for instance the validity of the naive Bayes assumption and hence accuracy of the estimates based upon it. Another problem is to find principled methods for setting priors (P_i in Eq. 3), for instance based on experimental evidence of the *type* of lighting cue (cast shadow versus occluding contour), or based on the amount or strength of information (e.g., length of occluding contour, contrast of a cast shadow, and so on). Another problem would be to develop a principled method for detecting inter-cue inconsistencies, for instance based on statistical hypothesis testing. (Such inconsistencies might indicate multiple sources or reveal artistic choices) by principled statistical hypothesis testing methods rather than by mere heuristics.[1] Likewise, it may be that a large standard deviation in the final estimate may indicate that the illuminant itself is fairly large. A further area is principled methods to integrate inherently one-dimensional, two-dimensional, and three-dimensional estimates.

Our methods, when refined and extended, may help art historians address a number of technical matters:

Working from memory Suppose the paintings of a realist artist generally have highly consistent cues to lighting. If a particular painting reveals large inconsistencies it may indicate that the artist was working from memory and did not have an actual tableau before him as reference.

Figure 10. The final estimated direction to the illuminant (red line) and cardiod describing the relative probability as a function of angle, at given by the data in Table 1.

Detect lighting inconsistencies for attribution or to reveal expressive techniques If the lighting position estimated from different subjects in a given tableau differ, it may indicate that the subjects were painted in separate sittings or campaigns. Likewise, these computer vision tools might shed light on the techniques of surrealists René Magritte, Salvador Dalí and Giorgio di Chirico who in some works alter the lighting information from that dictated by strict realism.

Detect differences in lighting in preparatory studies, underdrawings and final artworks These methods may help scholars detect an artist's alterations in lighting through preparatory sketches or underdrawings and *pentimenti* to help reveal how the artist arrived at the final choice of illumination.

Test for artist's use of optical aids It is securely established that in the 18th century some artists, for instance Canaletto (1697–1768), used optical devices such as the primitive camera obscura as drawing aids. The optical limitations of such devices provide a number constraints upon usable illumination, including its location within the tableau.[4] Stork and Johnson tested for the location of the illuminant, and thereby rejected the claim that an artist used projection optics when executing a particular painting.[5–7] The methods described in this current paper could be likewise applied to test for the use of optics, for instance in the case of Caravaggio (1571–1610) to determine whether the illumination in a tableau is consistent with sunlight streaming through a hole in his studio ceiling, a hole known through contemporary documentary records.[19]

Locate the subject in the studio setting As in the case just described, an accurate estimate of the location of an illuminant might allow scholars to place a rendered figure in a known setting with its associated illumination. For instance, the lighting direction estimated on some of Vermeer's figures might be consistent with a particular known window or lamp position.

More generally, we hope this work, extended and refined technically and informed by knowledge of art history, will find further use in humanistic studies of the visual arts.

ACKNOWLEDGMENTS

We would like to thank the Royal Cabinet of Paintings Mauritshuis, The Hague, for permission to reproduce *Girl with a pearl earring*. Dr. Johnson was supported by a gift from Adobe Systems, Inc., a gift from Microsoft, Inc., a grant from the United States Air Force (FA8750-06-C-0011), and by the Institute for Security Technology Studies at Dartmouth College under grant 2005-DD-BX-1091 from the Bureau of Justice Assistance and Award

Number 2006-CS-001-000001 from the U.S. Department of Homeland Security. Points of view or opinions in this document are those of the author and do not represent the official position or policies of the U.S. Department of Justice, the U.S. Department of Homeland Security, or any other sponsor.

REFERENCES

1. Y. Li, S. Lin, H. Lu, and H.-Y. Shum, "Multiple-cue illumination estimation in textured scenes," in *Proceedings of the 9th IEEE International Conference on Computer Vision (ICCV)*, **2**, pp. 1366–1373, 2003.
2. Q. Zheng and R. Chellappa, "Estimation of illuminant direction, albedo, and shape from shading," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* **13**(7), pp. 680–702, 1991.
3. G. D. Finlayson, S. D. Hordley, and M. S. Drew, "Removing shadows from images," in *Proceedings of the 7th European Conference on Computer Vision (ECCV)*, **2353**, pp. 129–132, Springer, (Berlin, Germany), 2002.
4. D. Hockney, *Secret knowledge: Rediscovering the lost techniques of the old masters*, Viking Studio, New York, NY, 2001.
5. D. G. Stork, "Did Georges de la Tour use optical projections while painting *Christ in the carpenter's studio*?," in *SPIE Electronic Imaging: Image and video communications and processing*, A. Said and J. G. Apolstolopoulos, eds., **5685**, pp. 214–219, SPIE, (Bellingham, WA), 2005.
6. 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," in *Proceedings of the 18th International Conference on Pattern Recognition*, **I**, pp. 255–258, IEEE Press, (Hong Kong), 2006.
7. D. G. Stork and M. K. Johnson, "Computer vision, image analysis and master art, Part II," *IEEE Multimedia* **14**(3), pp. 12–17, 2006.
8. D. G. Stork, "Imaging technology enhances the study of art," *Vision Systems Design* **12**(10), pp. 69–71, 2007.
9. D. G. Stork and Y. Furuichi, "Image analysis of paintings by computer graphics synthesis: An investigation of the illumination in Georges de la Tour's *Christ in the carpenter's studio*," in *Computer image analysis in the study of art*, D. G. Stork and J. Coddington, eds., **6810**, p. 68100J, SPIE/IS&T, Bellingham, WA, 2008.
10. W. Liedtke, *A view of Delft: Vermeer and his contemporaries*, Yale University Press, New Haven, CT, 2001.
11. D. G. Stork, "Mathematical foundations for quantifying shape, shading and cast shadows in realist master drawings and paintings," in *SPIE Electronic Imaging: Mathematics of data/image pattern recognition, compression and encryption with applications IX*, G. X. Ritter, M. S. Schmalz, J. Barrera, and J. T. Astola, eds., **6314**, pp. 63150K1–K6, SPIE, (Bellingham, WA), 2006.
12. P. Nillius and Jan-Olof Eklundh, "Automatic estimation of the projected light source direction," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR01)*, **1**, pp. 1076–1083, 2001.
13. J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes, *Computer graphics: Principles and practice in C*, Addison-Wesley, New York, NY, 1995.
14. B. K. P. Horn and M. J. Brooks, *Shape from shading*, MIT Press, Cambridge, MA, 1989.
15. M. K. Johnson and H. Farid, "Exposing digital forgeries by detecting inconsistencies in lighting," in *ACM Multimedia and Security Workshop*, pp. 1–9, ACM, New York, NY, 2005.
16. M. K. Johnson and H. Farid, "Exposing digital forgeries through specular highlights on the eye," in *9th International Workshop on Information Hiding*, T. Furon, ed., *Lecture Notes in Computer Science* **4567**, pp. 311–325, Springer, (Saint Malo, France), 2007.
17. S. Biswas, G. Aggarwal, and R. Chellappa, "Robust estimation of albedo for illumination-invariant matching and shape recovery," in *IEEE International Conference on Computer Vision (ICCV)*, IEEE, (Rio de Janeiro, Brazil), 2007.
18. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*, John Wiley and Sons, New York, NY, 2nd ed., 2001.
19. C. Puglisi, *Caravaggio*, Phaidon Press, New York, NY, 2000.