

Three-dimensional reconstruction from multiple reflected views within a realist painting: An application to Scott Fraser’s *Three way vanitas*

Brandon M. Smith,^a David G. Stork^b and Li Zhang^a

^aDepartment of Computer Science, University of Wisconsin, Madison WI 53706 USA

^bRicoh Innovations, 2882 Sand Hill Road Suite 115, Menlo Park CA 94025 USA
and Department of Statistics, Stanford University, Stanford CA 94305 USA

ABSTRACT

The problem of reconstructing a three-dimensional scene from single or multiple views has been thoroughly studied in the computer vision literature, and recently has been applied to problems in the history of art. Criminisi pioneered the application of single-view metrology to reconstructing the fictive spaces in Renaissance paintings, such as the vault in Masaccio’s *Trinità* and the plaza in Piero della Francesca’s *Flagellazione*. While the vast majority of realist paintings provide but a single view, some provide multiple views, through mirrors depicted within their tableaux. The contemporary American realist Scott Fraser’s *Three way vanitas* is a highly realistic still-life containing three mirrors; each mirror provides a new view of the objects in the tableau. We applied multiple-view reconstruction methods to the direct image and the images reflected by these mirrors to reconstruct the three-dimensional tableau. Our methods estimate virtual viewpoints for each view using the geometric constraints provided by the direct view of the mirror frames, along with the reflected images themselves. Moreover, our methods automatically discover inconsistencies between the different views, including ones that might elude careful scrutiny by eye, for example the fact that the height of the water in the glass differs between the direct view and that in the mirror at the right. We believe our work provides the first application of multiple-view reconstruction to a single painting and will have application to other paintings and questions in the history of art.

Keywords: painting analysis, three-dimensional reconstruction, multi-view reconstruction, Scott Fraser, *Three way vanitas*, fundamental matrix, essential matrix

1. INTRODUCTION

Algorithms and techniques of computer image analysis are being applied with growing frequency to problems in the understanding of fine art.¹ Such techniques include multispectral imaging to reveal hidden or degraded works,² color modelling and digital “reverse aging” of artworks to recover original color schemes,³ image texture analysis of brushstrokes to authenticate artworks,^{4,5} fractal and image analysis to authenticate drip paintings by Jackson Pollock,^{6–8} feature analysis of marks from worn press platens to date printed artwork,⁹ cast-shadow analysis to estimate the location of illuminants within depicted scenes,¹⁰ rigorous image photogrammetry and perspective estimation to infer artists’ working methods,^{11,12} computer dewarping of distorted images in convex mirrors depicted within paintings to attribute artists to these works,^{13,14} analysis of the pattern of highlights along occluding contours to locate illuminants within depicted scenes,¹⁵ and modelling putative projectors by computer ray tracing to test for artists’ use of optical drawing aids.¹⁶

The problem of reconstructing the three-dimensional picture space from two-dimensional projections is very well studied for the case of digital photographs and videos. Inferring the space from a *single* image is formally ill-posed and is generally addressed using constraints such as that particular lines are presumed parallel or perpendicular in the picture space, that the relative sizes of objects depicted are known, and so on. Inferring the three-dimensional space from two images—stereoscopically—is also well studied, and generally relies on finding corresponding “interest” points in the images and then computing the fundamental matrix \mathbf{F} , which

Send correspondence to David G. Stork, artanalyst@gmail.com.



Figure 1. Scott Fraser, *Three way vanitas* (2006), 89 × 124.5 cm, oil on board © Scott Fraser. The depicted art books here refer to French Fauvist painter Raoul Dufy (1877–1953) and the early Italian Renaissance master Fra Angelico (c. 1395–1455), one of the earliest practitioners of geometrical perspective. The penny coin, hidden behind the egg cup, appears due to two successive reflections; its central position and monetary value may reflect the artist’s worldly pecuniary strivings. Decoding the reflected sight line of this and other objects requires weaving an interesting web through the three-dimensional tableau, surely part of the goal of the artist. For instance, consider the view of the bottom of the egg cup visible near the center of the painting. Working backward from the viewer: the viewer looks at the middle of the large central mirror, then sees the rear reflecting surface of the blue hand mirror on the table, then sees in it the lower-right corner of the large plane mirror at the left, then sees in it the bottom of the egg cup. Visible in the distance in the left mirror a ticking metronome, a reminder of the passage of time, as well as a small photograph of the artist’s daughter, part of his legacy. The drinking glass contains water not yet drunk, symbolizing tasks not accomplished in this corporeal life. The chef’s hat, made by his pre-school son, sits atop a mountain lion skull, a gift to this Colorado-based artist from a fellow artist. The dead moths, pinned on spools of string—one freestanding, one in an egg cup—evoke the beauty an artist strives to leave behind. We can read this painting, then, as the artist’s self-reflective meditation on his career, accomplishments, and ultimate legacy.

then provides the camera models. Reconstruction from three overlapping views relies on computing the trifocal tensor \mathcal{T} , the generalization of the fundamental matrix. Inferring the three-dimensional space from more than three overlapping views involves the quadrifocal tensor which relates interest points visible in four views, but this approach is rarely followed because computation of the tensor is generally unstable. Instead, such four-view reconstruction relies on iterative refinement of successive subsets of images.

The new discipline of computer vision and image analysis in the study of art has addressed a number of questions, including the problem of reconstructing a space from a painting. Criminisi applied single-view uncalibrated methods to reconstruct the space in Hendrick van Steenwijck’s *St. Jerome in his study* (1630) and

the fictive space in Masaccio’s *Trinità* (1427–1428).^{11,17} Stork and his colleagues have built computer models of Georges de la Tour’s *Christ in the carpenter’s studio* (1645)¹⁸ and Jan Vermeer’s *Girl with a pearl earring* (c. 1665–1666)¹⁹ to answer questions about artists’ use of lighting, as well as of Hans Memling’s *Diptych of Maarten van Nieuwenhove* (1487) to test whether the mirror in its left panel was painted *in situ* or instead added later, and thus presumably from the Memling’s imagination.¹⁴ Stork built a computer graphics model of the space in Jan van Eyck’s *Portrait of Giovanni (?) Arnolfini and his wife* (1434) to estimate the focal length of the convex mirror depicted in the painting.²⁰ Wong and Stork built a highly faithful computer graphics model of Caravaggio’s *The calling of St. Matthew* to explore this artist’s number and placement of light sources.²¹




Figure 2. Scott Fraser, *Three way vanitas* (2006), 90.2 × 122 cm, graphite on paper (drawing) © Scott Fraser. Although the geometry in this study is nearly identical to that in the final oil painting of Fig. 1, the compositions differ in a number of small ways: In this study there is neither penny at the center nor metronome nor photograph visible in the left mirror. Likewise, there is no edge of table visible (and hence no droop of string), which gives a sense that this tableau is on a floor, rather than a table.

We are unaware of prior research on the three-dimensional reconstruction of the picture space in a painting from two separate views depicted within a single artwork. Clearly, for accurate reconstruction the painting must be geometrically accurate and highly realistic—possibly based on a photograph as reference, as in the case we consider here. In Sect. 2 we describe the painting we consider, Scott Fraser’s *Three way vanitas*, and then in Sect. 3 we review the basic theory and algorithms of reconstruction of three-dimensional spaces from multiple views, to clarify our notation and conventions. In Sect. 4 we apply these methods to Fraser’s painting and in Sect. 5 we interpret the results both on technical contributions and on how it expands humanistic understanding of this painting. We conclude in Sect. 6 with future technical work on reconstruction from multiple views and speculations on how these techniques might be used in further research in the visual arts.

2. THE PAINTING

Figure 1 shows the work in question, a formal, vanitas still life in oil on board executed in muted colors with spare background, arrayed much like an altar. It was commissioned as part of *The Object Project*,²² in which 15 contemporary artists were given the same objects (hand mirror, moth, bone, spool of string, drinking glass), and asked to create paintings. This work is a very sophisticated two- and three-dimensional composition, evoking a sense of clarity, stasis and reason. The vanitas genre comprises allegorical paintings, generally still lifes, meant to remind the viewer of the transitory nature of earthly life, and as such often contain a skull and a mirror, to metaphorically remind the viewer to see him or herself in the scene.

For most of his works, Scott Fraser works directly from life, but *Three way vanitas* was an exception. He used a photo for preliminary studies (such as in Fig. 2), and “blocked in” (set the overall composition) from a

distance of about 15 feet, then moved closer to observe details, such as the moths, spools of string and mountain lion’s skull. The composition is very finely aligned: moving the viewing point an inch to the side in the studio changes the composition significantly, as we discuss below.

3. RECONSTRUCTION THEORY

We now turn to the theory of reconstructing the three-dimensional scene based on the image in the painting. Figure 3 shows the setup and our conventions. The primary center of projection is C_0 , corresponding to the artist’s eye, or the pupil of the camera taking the reference figure; the image as viewed from this position is denoted I_0 , that is, the painting as we see it. Each of the three plane mirrors— M_1 , M_2 and M_3 , reading left-to-right—has its own effective center of projection, that is, the viewpoint yielding the (reversed) image we see in the painting. For instance, the center of projection C_1 is symmetric to C_0 with respect to the plane of M_1 , on the left, and its corresponding image is I_1 . Likewise, the centers C_2 and C_3 , have associated images I_2 and I_3 . Figure 3 shows the setup and our notation.




Figure 3. The top computer graphics figure shows the tableau and the centers of projection: the primary viewing point (of the artist or reference camera), is C_0 and the virtual centers of projection are C_i , symmetric with respect to each of the three primary plane mirrors, as indicated by the red lines perpendicular to the planes defined by each mirror, and as bisected by the green dots in each such plane. The four images beneath are I_0 , the projection from C_0 and corresponds to the painting we see, and the other images, I_i , corresponding to the views from C_i and the views visible in the mirrors in I_0 . For the three-dimensional reconstructions from multiple views, we use the I_1 , I_2 and I_3 images restricted to the views “through” the panes of the corresponding mirrors.

3.1 One view

Reconstructing a space from its projection to a single image, here I_0 is ill-posed, but can be addressed using geometrical constraints and assumptions. Criminisi pioneered the use of uncalibrated methods from computer vision to reconstruct the three-dimensional space within paintings.^{11,23,24} Here elementary perspective analysis—creation of vanishing points, horizon line, center of projection, and so forth—gives information about the viewing point, relative sizes of objects, and so on, as shown in Fig 4.




Figure 4. Perspective analysis of *Three way vanitas*. The vanishing points defined by the frames of mirrors M_1 and M_3 are VP_1 and VP_3 ; the horizon line is shown. The center of projection is perpendicular to the painting at the center of mirror M_2 at the height of the horizon line. The relative height of the string spool to the water glass can be computed geometrically using the red lines and vanishing point P according to the methods described by Criminisi,¹¹ and found to be approximately 1.04.

3.2 Two views

The basic theory underlying reconstructing three-dimensional spaces from two views is the foundation of stereo perception, and is well understood and includes a number of basic algorithms, such as bundle adjustment.²⁵ The first step is to identify—either by hand, semi-automatically, or automatically—pairs of corresponding interest points in the two images, then compute the fundamental matrix \mathbf{F} . This 3×3 rank-two matrix has seven degrees of freedom and is defined by $\mathbf{x}^t \mathbf{F} \mathbf{x}' = 0$ for each pair of matching interest points $\mathbf{x} \leftrightarrow \mathbf{x}'$ in two images, where $\mathbf{x} = (x, y, 1)^t$ is in one image, and $\mathbf{x}' = (x', y', 1)^t$ in the other image. Each of the k pairs of points $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ ($i = 1, \dots, k$) yields one linear equation in the unknown entries in \mathbf{F} .²⁵ One then derives from \mathbf{F} the two centers of projection and camera models; simple projective geometry then yields the three-dimensional location of these corresponding interest points. There are a number of complexities in this basic algorithm, such as when occlusion obscures one interest point in a matching pair, or when there are multiple locally optimal correspondence solutions. We used these traditional methods, using the Levenberg-Marquart descent to find the global minimum, as outlined in the Appendix. The overall scale is generally not provided by such analysis, but requires some known size, for instance the absolute size of an object in the scene or known distance or angle of view.

The generalization to three-view reconstructions involves three-way point matches, $\mathbf{x} \leftrightarrow \mathbf{x}' \leftrightarrow \mathbf{x}''$, where the role of the fundamental matrix is played by the trifocal tensor \mathcal{T} , which consists of three 3×3 matrices, each corresponding to one of the three pairs of images.²⁵ The rigorous mathematical theory underlying general n -view reconstruction for $n \geq 4$ is not often employed because the estimation of higher-order tensors is very sensitive to errors in the location of interest points. In practice, then, high- n multi-view algorithms iteratively refine the reconstruction based on subsets of the images.²⁵

3.3 Error measure

A principled measure of the quality of fit of a reconstruction is the first-order geometric error or Sampson distance²⁶ is given by:

$$\sum_{i=1}^k \frac{(\mathbf{x}'_i{}^t \mathbf{F} \mathbf{x}_i)^2}{(\mathbf{F} \mathbf{x}_i)_1^2 + (\mathbf{F} \mathbf{x}_i)_2^2 + (\mathbf{F} \mathbf{x}'_i)_1^2 + (\mathbf{F} \mathbf{x}'_i)_2^2}. \quad (1)$$

This measure is the sum over k pairs of interest points of the normalized square of the distance of a point in one view from the location predicted in the other view.

4. APPLICATION TO *THREE WAY VANITAS*

We now turn to the reconstruction of the space from the data in the painting. First we identified by hand “interest” points, such as corners and inflection point of objects, noting for our database in which of the images, I_0, \dots, I_3 , the point was visible, as shown in Fig. 3.

4.1 One view

We assume naturally that the tabletop is horizontal, each mirror frame is rectangular and planar with many vertical edges, that the intrinsic shapes of the flanking mirrors are the same, and so on. We consider first just the single view provided by image I_0 (cf., Fig. 3). Figure 4 shows a perspective analysis of the work including the horizon line, vanishing points, and so on. There are a number of different single-view images and each could be analyzed, but we are primarily concerned with multi-view reconstruction. Nevertheless, there are a few regions that are visible in only one image, so we illustrate the kind of information that can be extracted from those. For example, the spool behind mirror M_1 is visible in only the direct image, I_0 . We can deduce the relative size of the glass and this spool of string, as shown in red.¹¹ The lower red line links the center of the bases of the glass with that of the spool; this line, extended to the horizon line defines the vanishing point P . The upper red line links P to the center of the top of the glass. All objects on the horizontal ground plane whose top touches this red line and bottom touches the other red line are the same height, regardless of distance to the viewer. We can thus read the height of the spool, h_1 , as a proportion of the separation of the red lines, h_2 , and thus find that $h_1/h_2 \approx 1.04$. The overall scale cannot be derived from the form information alone but can be inferred by knowledge of the scale set by the penny coin, for instance.

4.2 Two views

Nearly every algorithm for multiple-view reconstruction rely on finding correspondences between *interest points* (sometimes called *key points*): points visible in two or more views. In applications relying on a large number of interest points, such points are computed automatically for example by SIFT features (shift-invariant feature transform), points of high curvature, terminations, and so on.

Figure 3 shows C_0 as the center of projection of the direct view, the virtual image in the left mirror (M_1) has center of projection C_1 , and likewise for the other two mirrors. The view we see is reversed, left-to-right from those centers. There are $\binom{4}{2} = 6$ pairs of images that can be used in two-image reconstruction. Two views, such as the direct view, and one provided by the central mirror provide the largest number of pairs of interest points.

Figure 6 shows the three-dimensional tableau reconstructed from images I_0 and I_2 , the largest within the tableau. (Other paired reconstructions are similar.) The general spatial arrangement seems plausible (the hand mirror is behind the glass of water, the bone is closer than the moth at its right, and so on), but there are some clear inconsistencies, such as the fact that the bone is significantly closer than the water glass.

5. INTERPRETATION AND ANALYSIS

Our three-dimensional reconstruction reveals the sophisticated composition in a way difficult in a purely two-dimensional analysis. Our model allows us to adjust the viewing position (C'_0) slightly, or the placement of objects. We find that the slightest realignment—by as little as a centimeter in the space of the tableau—can introduce occlusions and awkward compositions in the direct final view, the reflected images in any or all of the mirrors in the final painting. (We also can demonstrate this using our computer graphics model.)

6. CONCLUSIONS AND FUTURE DIRECTIONS

Clearly our work is preliminary, a mere proof of concept that one can reconstruct three-dimensional spaces from multiple viewed depicted in a realist painting. There are a number of future research directions beyond those identified in Sect. 5 above. In the technical algorithm development relevant to art, we envision two key directions, both of which seem to have been overlooked in the computer vision literature:



Figure 5. The small red +s are the hand selected interest points, typically chosen to be point of high intrinsic curvature or extrema on the boundary of objects, or intersections of two curves.

Integrate single-view with multiple-view estimates Our application of computer vision algorithms to this painting highlights a challenge to reconstructing in general realist images: how to integrate information about the separate “camera” models derived from the fundamental matrix (or trifocal tensor) and that estimated from a *single* image. In the present case, this single-image information consists of the perspective of each mirror frame depicted in the painting, which translates, by spatial mirror symmetry, into the location of the virtual centers of projection. A related problem is how to use information from multiple *sequential* reflections (i.e., reflections of reflections) in such reconstructions. Both these problems demand we weight the different spatial errors or uncertainties and propagate them to the final three-dimensional points. We are developing theory and algorithms to address both these general problems.

Incorporating single-view geometric constraints The method described here was purely image based, and relied on finding the stereo correspondences between interest points to infer the Fundamental matrix and thus the effective centers of projection. But note that the monocular view of each mirror frame and the rules of geometrical optics of reflection provides information about the location of the corresponding centers of projection. For instance, the single perspective view (from C_0) of mirror M_2 , assumed to be rectangular, gives an estimate of the location of the center of projection C_2 , up to a scale. One could use this estimate in addition to the image I_2 itself to get an improved three-dimensional reconstruction.

Reconstruct using convex mirror reflections Criminisi, Kemp and Kang showed how to dewarp painted reflections in convex mirror depicted within paintings, such as the *Arnolfini portrait* and *Heinrich von Werl and St. John the Baptist* for authentication, providing a new view.¹³ Stork argued that the high spatial agreement between the direct view and the dewarped mirror view in the *Arnolfini portrait* implied that van Eyck worked from an actual, rather than fictive, space.²⁰ Future research could include using such a *dewarped* image along with the direct view for three-dimensional reconstruction.




Figure 6. The three-dimensional reconstruction of the tableau based on the direct image, I_0 and the reflected image in the large central mirror, I_2 .

Our techniques might be applied to other works in the art historical canon, to provide new views into the spaces depicted in paintings and hence shed light on art historical questions:

Reconstructing space in works depicting plane mirrors There are a number of realist paintings that include plane mirrors, and one might reconstruct parts of the three-dimensional tableau using the techniques we describe: Jean-Auguste-Dominique Ingres’s *Portrait of the Countess D’Haussonville* (1845), Bernardo Strozzi’s *Old woman at the mirror* (c. 1615), and Peter Paul Rubens’ *Venus at the mirror* (1613–1614) or modern photorealists, such as Richard Estes’ *Paris street scene, 1972* (1972) and *Spirit* (1995–1996).

Convex mirror reflections Reconstruct the three-dimensional space of the tableau from the direct view and a dewarped view in a convex mirror, as in Jan van Eyck’s *Portrait of Giovanni (?) Arnolfini and his wife* (1434), Robert Campin’s *Heinrich von Werl and St. John the Baptist* (1438) and other works.

Quantifying spatial composition Much of “stylometry” in the visual arts (the quantification of artistic style) relies on low-level or two-dimensional features such color palette or the shape of brush strokes. Perhaps with refinements and extensions to the methods described here, including uncalibrated single-image methods we can begin to develop a stylometry for three-dimensional composition, one relying on sizes and distances of objects.

Reconstruct spaces from different paintings Just as multiple photographs of a given object or building, taken by different photographers at different times, can be used to reconstruct a three-dimensional model of that building, so too perhaps can multiple realistic paintings of a given object be used to reconstruct a three-dimensional model of a building. This would be particularly intriguing if the building no longer exists, for instance a Greek or Roman building depicted in paintings.

Most importantly, we believe our work builds upon and further demonstrates the value of rigorous computer vision methods in the humanistic study of art, but that its greatest value will be realized when art historians and

conservators understand the power and limitations of these methods and identify problems that can be addressed by these methods.¹

Acknowledgements

We would like to express our deep thanks to Scott Fraser for providing background information and explaining his working methods and providing high-resolution scans of his two versions of *Three way vanitas*. We would also like to thank Yasuo Furuichi who graciously created the computer graphics model for Fig. 3.

APPENDIX

We can extend the Levenberg-Marquardt iteration method to perform a bundle adjustment of all the views at the same time. As with using the Levenberg-Marquardt method for one pair of views, we must supply an initial estimate. A problem arises when we do this, however, since 3D points recovered from one pair of views might have a different scale associated with them than another pair of views. For example, a 3D point from one pair of views might have the form $\mathbf{X}_1 = (a, b, c)^t$ and the same point recovered from another pair of views might have the form $\mathbf{X}_2 = (a/q, b/q, c/q)^t$, that is, the ratios between the coordinates are the same, but the actual points differ by a scale. To make them consistent we can calculate a 3×3 matrix \mathbf{H} such that $\mathbf{X}_1 = \mathbf{H}\mathbf{X}_2$ and \mathbf{X}_2 is updated by \mathbf{H} by $\mathbf{X}'_2 = \mathbf{H}\mathbf{X}_2$.

The key to making this work is to also update the camera matrix \mathbf{P}_2 so that $\mathbf{P}'_2 = \mathbf{P}_2\mathbf{H}^{-1}$ such that

$$x_2 = \mathbf{P}_2\mathbf{X}_2 = \mathbf{P}_2\mathbf{H}^{-1}\mathbf{H}\mathbf{X}_2 = \mathbf{P}'_2\mathbf{X}'_2 = \mathbf{P}'_2\mathbf{X}_1.$$

Thus, after updating we are free to use \mathbf{X}_1 as an initial estimate since it now works for both pairs of views. Thus \mathbf{H} has the form

$$\mathbf{H} = \begin{pmatrix} \mathbf{A} & t \\ v^t & k \end{pmatrix}$$

where v and t are 3-vectors, \mathbf{A} is a 3×3 matrix, and k is a scalar. The 3D points are really only off by a scale, and so we can make the following simplification to \mathbf{H} :

$$\mathbf{H} = \begin{pmatrix} \mathbf{I} & 0 \\ v^t & k \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ v_1 & v_2 & v_3 & k \end{pmatrix}.$$

REFERENCES

1. D. G. Stork and J. Coddington, eds., *Computer image analysis in the study of art*, vol. 6810, SPIE/IS&T, Bellingham, WA, 2008.
2. R. L. Easton, Jr., K. T. Knox, and W. A. Christens-Barry, "Multispectral imaging of the Archimedes palimpsest," in *Proceedings of the 32nd Applied Imagery Pattern Recognition Workshop*, **AIPR03**, pp. 111–116, IEEE, (El Segundo, CA), 2003.
3. R. S. Berns, "Rejuvenating Seurat's palette using color and imaging science: A simulation," in *Seurat and the making of La Grande Jatte*, R. L. Herbert, ed., pp. 214–227, Art Institute of Chicago, Chicago, IL, 2004.
4. S. Lyu, D. Rockmore, and H. Farid, "A digital technique for art authentication," *Proceedings of the National Academy of Sciences* **101**(49), pp. 17006–17010, 2004.
5. M. Shahram, D. G. Stork, and D. Donoho, "Recovering layers of brushstrokes through statistical analysis of color and shape: An application to van Gogh's *Self portrait with grey felt hat*," in *Computer image analysis in the study of art*, D. G. Stork and J. Coddington, eds., **6810**, pp. 6810D–1–8, SPIE/IS&T, Bellingham, WA, 2008.
6. R. P. Taylor, A. P. Micolich, and D. Jonas, "Fractal analysis of Pollock's drip paintings," *Nature* **399**, p. 422, 1999.

7. K. Jones-Smith and H. Mathur, "Fractal analysis: Revisiting Pollock's drip paintings," *Nature* **444**(doi: 10.1038/nature05398), pp. E9–E10, 2006.
8. M. Irfan and D. G. Stork, "Multiple visual features for the computer authentication of Jackson Pollock's drip paintings: Beyond box-counting and fractals," in *SPIE Electronic Imaging*, p. (in press), SPIE/IS&T, Bellingham, WA, 2009.
9. S. B. Hedges, "A method for dating early books and prints using image analysis," *Proceedings of the Royal Society A* **462**(2076), pp. 3555–3573, 2006.
10. 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. Apostolopoulos, eds., **5685**, pp. 214–219, SPIE, (Bellingham, WA), 2005.
11. A. Criminisi, *Accurate visual metrology from single and multiple uncalibrated images*, ACM Distinguished Dissertation Series, Springer-Verlag, London, 2001.
12. A. Criminisi and D. G. Stork, "Did the great masters use optical projections while painting? Perspective comparison of paintings and photographs of Renaissance chandeliers," in *Proceedings of the 17th International Conference on Pattern Recognition*, J. Kittler, M. Petrou, and M. S. Nixon, eds., **IV**, pp. 645–648, 2004.
13. A. Criminisi, M. Kemp, and S.-B. Kang, "Reflections of reality in Jan van Eyck and Robert Campin," *Historical methods* **3**(37), pp. 109–121, 2004.
14. S. Savarese, R. Spronk, D. G. Stork, and A. DelPozo, "Reflections on praxis and facture in a devotional portrait diptych: A computer analysis of the mirror in Hans Memling's *Virgin and Child and Maarten van Nieuwenhove*," in *Computer image analysis in the study of art*, D. G. Stork and J. Coddington, eds., **6810**, pp. 68100G–1–10, SPIE/IS&T, Bellingham, WA, 2008.
15. 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.
16. M. D. Robinson and D. G. Stork, "Aberration analysis of the putative projector for Lorenzo Lotto's *Husband and wife*: Image analysis through computer ray-tracing," in *Computer image analysis in the study of art*, D. G. Stork and J. Coddington, eds., **6810**, pp. 68100H–1–11, SPIE/IS&T, Bellingham, WA, 2008.
17. A. Criminisi, M. Kemp, and A. Zisserman, "Bringing pictorial space to life: Computer techniques for the analysis of paintings," in *Digital art history: A subject in transition*, A. Bentkowska-Kafel, T. Cashen, and H. Gardner, eds., pp. 77–100, Intellect Books, Bristol, UK, 2005.
18. 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**, pp. 68100J–1–12, SPIE/IS&T, Bellingham, WA, 2008.
19. M. K. Johnson, D. G. Stork, S. Biswas, and Y. Furuichi, "Inferring illumination direction estimated from disparate sources in paintings: An investigation into Jan Vermeer's *Girl with a pearl earring*," in *Computer image analysis in the study of art*, D. G. Stork and J. Coddington, eds., **6810**, pp. 68100I–1–12, SPIE/IS&T, Bellingham, WA, 2008.
20. D. G. Stork, "Optics and realism in Renaissance art," *Scientific American* **291**(6), pp. 76–84, 2004.
21. L. Wong and D. G. Stork, "Computer graphics synthesis for analysis of complex lighting in paintings: An application to Caravaggio's *The calling of St. Matthew*," p. (submitted), 2009.
22. M. S. Doherty, *Object Project: Five objects, fifteen artists*, Evansville Museum of Arts, History and Science, Evansville, IN, 2007.
23. A. Criminisi, I. Reid, and A. Zisserman, "Single view metrology," *International Journal of Computer Vision* **40**(2), pp. 123–148, 1999.
24. D. G. Stork, "Optics and the old masters revisited," *Optics and Photonics News* **15**(3), pp. 30–37, 2004.
25. R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*, Cambridge University Press, Cambridge, UK, 2nd ed., 2003.
26. P. D. Sampson, "Fitting conic sections to 'very scattered' data: An iterative refinement of the Bookstein algorithm," *Computer Vision, Graphics, and Image Processing* **18**(1), pp. 97–108, 1982.