MEASURING ART

Reflections of Reality in Jan van Eyck and Robert Campin

ANTONIO CRIMINISI
Microsoft Research, Ltd.
Cambridge, UK

MARTIN KEMP
Department of History of Art
University of Oxford

SING BING KANG
Microsoft Research, USA

Abstract. There has been considerable debate about the perspectival/optical bases of the naturalism pioneered by Robert Campin and Jan van Eyck. Their paintings feature brilliantly rendered *convex* mirrors, which have been the subject of much comment, especially iconographical. David Hockney has recently argued that the Netherlandish painters exploited the image-forming capacities of *concave* mirrors. However, the secrets of the images within the painted mirrors have yet to be revealed. By means of novel, rigorous techniques to analyze the geometric accuracy of the mirrors, unexpected findings emerge, which radically affect our perception of the way in which the paintings have been generated. The authors focus on Jan van Eyck's *Arnolfini Portrait*, and the Heinrich von Werl Triptych, here reattributed to Robert Campin. The accuracy of the images in the convex mirrors depicted in these paintings is assessed by applying mathematical techniques drawn from computer vision. The proposed algorithms also allow the viewer to "rectify" the image in the mirror so that it becomes a central projection, thus providing a second view from the back of the painted room. The plausibility of the painters' renderings of space in the convex mirrors can then be assessed. The rectified images can be used for purposes of three-dimensional reconstruction as well as for measuring accurate dimensions of objects and people. The surprising results presented in this article cast a new light on the understanding of the artists' techniques and their optical imitation of seen things and potentially require a rethinking of the foundations of Netherlandish naturalism. The results also suggest that the von Werl panels should be reinstated as autograph works by Campin. Additionally, this research represents a further attempt to build a constructive dialogue between two very different disciplines: computer science and art history. Despite their fundamental differences, the two disciplines can learn from and be enriched by each other.

Keywords: accuracy of paintings, machine vision, opticality theory, projective geometry

The advent of the astonishing and largely unprecedented naturalism of paintings by the van Eyck brothers, Hubert and Jan, and of the artist called the Master of Flémalle (now generally identified with the documented painter Robert Campin), has always been recognized as one of the most remarkable episodes in the history of Western art. Various explanatory modes have been developed to explain the basis of the new naturalism and its roles in the religious and secular societies of the Netherlands in the early fifteenth century. Among the new technical factors, the perfection and innovative use of the oil medium are clearly seminal. In terms of meaning, the naturalistic integration of objects within coherent spaces allowed a profusion of symbolic references to be "concealed" within what seems to be the normal ensemble of a domestic or ecclesiastical interior. It has also been long recognized that the coherence of the spaces in Netherlandish painting did not rely upon any dogmatic rule (e.g., the convergence of orthogonals to the "centric," or vanishing point) that provided the basis for Italian perspective in the wake of Brunelleschi (see fig. 1). The straight edges of forms perpendicular to the plane of the picture converge in a broadly systematic manner but are not subject to precise optical geometry. We concentrate on two paintings, the so-called *Arnolfini Portrait* by Jan van Eyck (figs. 2a, b) and the *St. John the Baptist with a Donor* from the Heinrich von Werl Triptych (figs. 2c, d), both of which demonstrate how convincing spaces could have been created by imprecise means.

The effects and meaning of the new naturalism have been the subject of more attention than the questions of the visual or optical resources used to compile the images and how such revolutionary naturalism could be forged in the face of resistant conventions. The English artist David Hockney (2001) has argued that optical projections provide the key. Initially disposed to argue that the artist relied very directly on images projected onto white surfaces by lenses or concave mirrors, Hockney now places more emphasis on the projected image as the key breakthrough in revealing what a three-dimensional, or 3-D, array looks like when flattened by projection. In any event, he rapidly came to see that an image such as Jan

FIGURE 1. Convergence of Orthogonals in Paintings: (a) Masaccio, *Trinity*, c. 1426, Santa Maria Novella, Florence. (b) The dominant orthogonals converge at a single vanishing point. (c) J. van Eyck, *Madonna in the Church*, 1437–1439, Staatliche Museen zu Berlin, Gemäldegalerie, Berlin. (d) The orthogonals (marked) do not converge at a single vanishing point.

van Eyck's would not have been projected as a whole for literal imitation, not least because the overall perspective of the paintings is consistently inconsistent. Rather, Hockney advocated that the components in the image were studied separately in optical devices, to be collaged together in what he has called a "many windows" technique. The arguments between Hockney and his detractors have become increasingly bogged down in personalized polemic that dogmatically uses varieties of technical evidence in the service of predetermined stances, without much sense of improvisatory procedures that typify artists' uses of the tools at their disposal.

Our intention is to inquire about Netherlandish naturalism along a different route. Here, we undertake a close analysis of the remarkable convex mirrors in the two paintings, using techniques of computer vision to "rectify" the curved images in the mirrors. We want to know what the rectified images tell us about the spaces that the painter has portrayed—looking, as it were, into the interiors from the other end—and what implications emerge for our understanding of how the painters composed their images. We find unavoidable implications for their approach to stage-managing their ensembles of figures and objects. Along the way, we suggest that the position of the von Werl Triptych in relation to Campin's oeuvre and the works of his followers needs to be reassessed. Our conclusions will have a bearing on the Hockney hypothesis, though in a permissive rather than conclusive manner.

First, we want to introduce the two paintings that will be serving as our witnesses.

The Arnolfini Portrait

The painting by Jan van Eyck in the National Gallery, London, executed in oil on an oak panel, is generally (though not universally) identified as portraying Giovanni di Niccolò Arnolfini and Giovanna Cenami. Arnolfini was a successful merchant from Lucca and a representative of the Medici bank in Bruges, who satisfied Philip the Good's eager demands for large quantities of silk and velvet. The double portrait has reasonably been identified as representing an event, namely, Giovanni and Giovanna's marriage or, more probably, their betrothal, as witnessed by the two men reflected in the mirror. Jan himself appears to have been one of the witnesses, because a beautifully inscribed graffito on the back wall records the date, 1434, and the fact that "Johannes de eyck fuit hic" (Jan van Eyck was here).

The St. John Panel

The other painting, the *St. John with a Donor* in the Museo del Prado, Madrid, uses the oil medium no less brilliantly than does van Eyck. It is one of the two surviving wings of an altarpiece commissioned, as the inscription tells us, by Heinrich von Werl in 1438. Von Werl was a Franciscan theologian from Cologne who became head of the Minorite Order (a Conventual branch of the Franciscans). The right wing depicts St. Barbara in an interior within which light effects are rendered with notable virtuosity. The main panel, probably containing the Virgin and Child with saints (one of whom would almost certainly have been St. Francis), is no longer traceable. It is likely that von Werl had chosen St. John as his personal saint, to act the intercessor who grants him the privilege of witnessing the Virgin through the door that opens to the sacred realm of the central panel. Two Franciscans, visible in the painted mirror, witness the scene. The play of domestic and sacred space would have been crucial in the reading of the triptych's meaning and in signaling the donor's

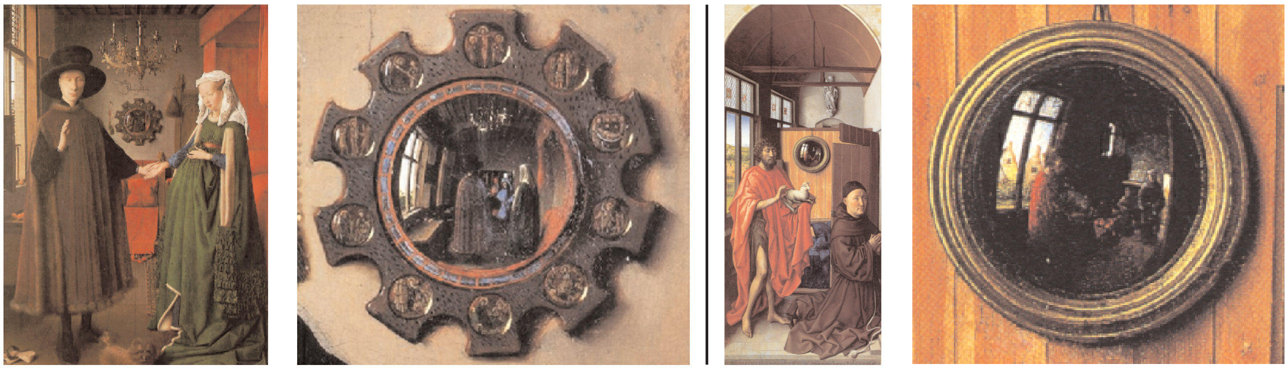

a b c d

FIGURE 2. Original Paintings Analyzed in Text: (a) J. van Eyck, *Arnolfini Portrait,* panel, 1434, National Gallery, London. (b) Enlarged view of the convex mirror in (a). (c) R. Campin, *St. John with a Donor,* panel, 1438, Museo del Prado, Madrid. (d) Enlarged view of the convex mirror in (c).

hoped-for translation to the heavenly dwelling of the Virgin and saints after his death.

In the following discussion, we provide a geometric analysis of the two painted mirrors on the basis of the optics of curved mirrors and transform the images into normalized rectilinear views of the painted interiors. The rectified images are subject to accuracy analysis, and the artists' possible manipulation of the effects is considered. Finally, we spell out the surprising and radical implications that emerge from the study of two such apparently subsidiary aspects of the complex images.

GEOMETRIC ANALYSIS OF PAINTED CONVEX MIRRORS

In this section, we present some simple but rigorous techniques for the analysis of the geometric accuracy of convex mirrors in paintings. The algorithms used here build upon the vast computer vision literature, with specific reference to the field of compound image formation, termed *catadioptric imaging* (Benosman and Kang 2001; Nayar 1997). A catadioptric system uses a camera and a combination of lenses and mirrors.

The Basic Mirror-Camera Model

Given a painting of a convex mirror (e.g., figs. 2b, d), we model the combination mirror-panel (or, similarly, mirror-canvas) as a catadioptric acquisition system comprising a mirror and an orthographic camera as illustrated in figures 3a, b.

It is well known that curved mirrors produce distorted reflections of their environment (see figs. 3c, d). For instance, straight scene lines become curved (see the window frames in figs. 2b, d). However, if the shape of the mirror is known, then the reflected image can be transformed (warped) in such a way as to produce a perspectively corrected image—that is, an image as would be formed by a camera with optical center in the center of the mirror.[1] Whereas this is true for real mirrors, the case of painted mirrors is more complicated. In fact, one should not forget that a painting is not a photograph but rather the product of the hand of a skilled artist; even when it appears to be correct, its geometry may differ from that produced by a real camera.

In this article, we propose novel techniques for (a) assessing the geometric accuracy of a painted mirror and (b) producing rectified perspective images from the viewpoint of the mirror itself.

The basic algorithm may be described in general terms as follows:

1. Hypothesizing the shape of the convex mirror from *direct* analysis of the distortions in the original images.
2. Rectifying the input image (distorted) to produce the corresponding perspective (corrected) image.
3. Measuring the discrepancy between the rectified image and the "perspectively correct" image to assess the accuracy of the painted mirror.

As will subsequently be clarified, these basic geometric tools will help us cast a new light on the analyzed paintings and their artists. To proceed, we need to make some explicit assumptions about the shape of the mirror. We explore three different rectification assumptions:

- *Assumption 1*: The mirror has a parabolic shape.
- *Assumption 2*: The mirror has a spherical shape.
- *Assumption 3*: The geometric distortion shown in the reflected image can be modeled as a generic radial-distortion process.

For each of those three basic assumptions, from the direct analysis of the image data alone, our algorithm estimates the best geometric transformations that can "remove" the

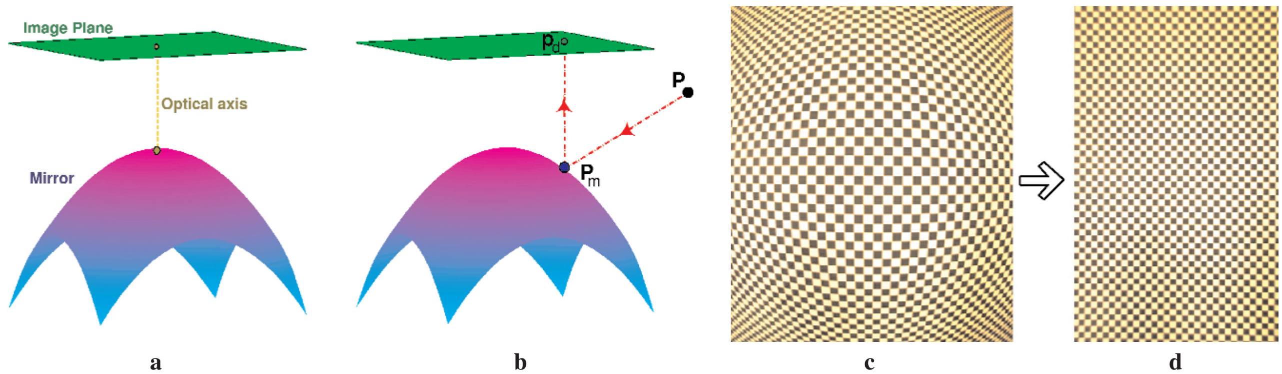

FIGURE 3. Catadioptric Imaging Model: (a) The mirror-panel pair is modeled as a catadioptric acquisition system, comprising a curved mirror and an acquiring camera. (b) A ray emanating from a point P in the scene is reflected by the mirror onto the corresponding image (panel) point p_d. The ray is reflected off the mirror surface in the point P_m. (c) A square grid reflected by a convex mirror appears barrel-distorted. (d) One of our goals is to "rectify" the image to produce a perspective image. In the case of fig. (c), we wish to recover an image of the regular grid, where all the black-and-white patches are perfect squares.

distortion in the reflected image and rectify it into a perspective image. A detailed discussion of the geometry of mirrors and the rectification algorithm appears in the appendix.

It must be noted, though, that from a historical point of view the parabolic mirror shape can be discarded because the standard convex mirror of the time was cut from a blown glass *sphere*. However, the parabolic mirror remains useful as a model in the following analyses.

Rectification Results

Results of the rectification of Campin and van Eyck's mirrors are illustrated in figure 4 and figure 5, respectively. An initial observation is that the three different mirror-shape assumptions lead to very similar (but not identical) rectifications, which can be explained if one considers the fact that the central portions of the parabolic and spherical mirrors are very similar in shape. In fact, it is conventional in optics to make parabolic assumptions for lenses and mirrors when small portions of their spherical surface are being considered. When the visible parts of the mirrors are limited central portions, as in these two cases, it is not surprising that the rectification results are quite similar. Furthermore, the fact that both the spherical assumption and the parabolic one yield such similar rectification results validates the assumption of substantial distances between the 3-D scene points and the mirror, an assumption that needs to be made in the case of the spherical mirror. The comparison between the rectifications obtained by the parabolic and spherical assumptions is an indirect way of assessing the validity of the rectification approach we have undertaken.

Note that horizontal flipping of the rectified images in figures 4b, c, d and figures 5b, c, d would produce the perspective images that an observer would see if he or she stood in the place of the mirror in the depicted scene (e.g., the large window would be on the right-hand side).

The rectification process produces novel perspective images from a different vantage point than that of the artist, thus, in effect, giving rise to stereo views of the depicted scene. Given two views of the same scene from two different viewing positions, it is conceivable to apply standard 3-D computer vision techniques (Faugeras 1993; Hartley and Zisserman 2000) to reconstruct complete virtual models of the depicted environment. Alternatively, single-view reconstruction (Criminisi 2001) may be applied twice: to the image from the front (the painting itself) and to the image from the back (the rectified mirror reflection), with the two resulting 3-D reconstructions merged together to create a single complete shoebox-like virtual model of the scene.

Comparing Mirror Protrusions

As described in the appendix, in the case of the spherical-mirror assumption, to rectify the input (reflection) images it is necessary to compute the radius r of the spherical mirror itself (see table A1 in the appendix).

Because we are working directly on the image plane with *no* additional assumptions on absolute scene measurements, it is *not* possible to measure radii in absolute metric terms but rather only as relative measurements. For instance, we can measure the ratio between the radius r and the radius of the *disk*, the visible part of a spherical mirror on the image plane, in each painting (see fig. 6a).

Figure 6a shows a cross-section of a spherical mirror protruding from a wall with the disk radius clearly marked. Figure 6b shows a comparison between two spherical mirrors with the second mirror (mirror 2) bulgier than the first one. We now define a measure of the *protrusion* (bulge) of

a b c d

FIGURE 4. Rectification of the Convex Mirror in Campin's St. John Panel: (a) Original (distorted) mirror from fig. 2d. (b, c, d) Rectified images. (b) Result of rectification using the parabolic assumption. (c) Result of rectification using the spherical assumption. (d) Result of rectification using the generic radial model fitting. Note that in all the rectified images the edges of the window and door are straight.

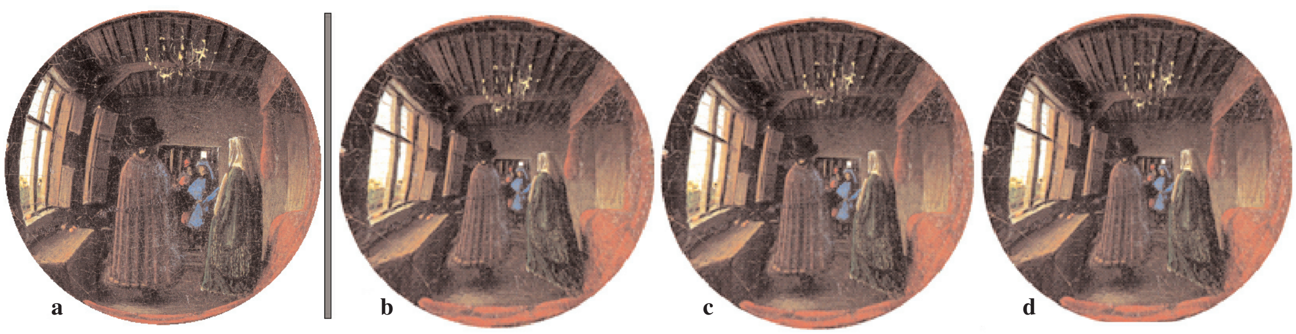

a b c d

FIGURE 5. Rectification of the Convex Mirror in van Eyck's Portrait: (a) Original (distorted) mirror from fig. 2b. (b, c, d) Rectified images. (b) Result of rectification via the parabolic assumption. (c) Result of rectification via the spherical assumption. (d) Result of rectification via generic radial model fitting.

a spherical mirror as the ratio P between the disk radius and the radius of the sphere.

$$P = \frac{r_{disk}}{r} \quad (1)$$

If $r_2 < r_1$ and the disk radius r_{disk} is the same, then $P_2 > P_1$ (see fig. 6b).

From the protrusion measures reported in table 1 for Campin's and van Eyck's painted mirrors, we can conclude that the two mirrors are very similar in terms of shapes (similar protrusions P), but the former is slightly bulgier (greater value of P) than the latter.

This result should be taken with some care. In fact, to be completely certain of the degree of protrusion of each mirror, we need to be certain about the accuracy level of its rendered geometry. The following section will show that although we can be quite confident about the geometric correctness of Campin's mirror, the same cannot be said of van Eyck's.

The protrusion measurements of table 1 can also be used to compute the focal lengths of the two mirrors. In fact, first-order optics teaches us that the radius of curvature r of a mirror is twice its focal length. Therefore, from equation (1), if we knew r_{disk}, we could compute r and, therefore, $f = r/2$ would be the focal length. Note, though, that an absolute metric value for r_{disk} can be estimated only after we make further assumptions on shapes and sizes of objects that appear in the scene. These kinds of assumptions would make sense if, as in the case of a few paintings by Jan Vermeer, some depicted objects still exist and can

TABLE 1. A Comparison of the "Bulge" of the Two Mirrors

Painting	Protrusion factor (P)
Campin's *St. John*	$P = 0.83$
van Eyck's *Arnolfini Portrait*	$P = 0.78$

be measured (Kemp 1990; Steadman 2001). But in the case of the two paintings analyzed in this article, we feel that assumptions not supported by strong physical evidence may conduce to misleading results on the value of the mirrors' focal lengths.

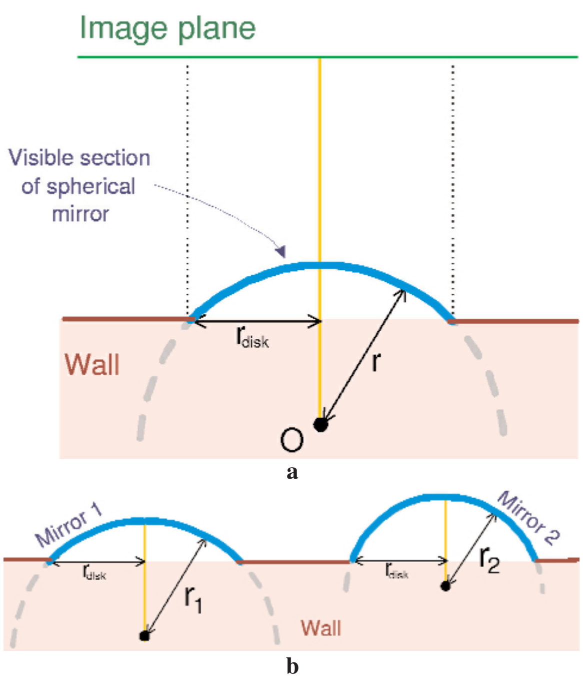

FIGURE 6. (a) The visible part of a spherical mirror is called a *disk*. The quantity r_{disk} is the disk radius. (b) A cross-section of two spherical mirrors protruding from a wall. Because $r_1 > r_2$ and r_{disk} is the same for both mirrors, the second mirror sticks out more and thus is "bulgier" (has a larger protrusion) than the first mirror.

Accuracy Analysis

We have already described how direct analysis of the image of a convex mirror can lead to the removal of the inherent optical distortion in the original images and the generation of the corresponding perspectively correct images. We have also been able to measure the "bulge" of the two mirrors. In the following section, we analyze instead the rectified images to assess their geometric consistency.

Accuracy of Campin's St. John Panel

Straightening of curved edges. Figure 7a (identical to fig. 4c) shows the rectification of Campin's mirror obtained from the assumption of the spherical mirror. Images of straight scene lines have now become consistently straight (see fig. 7b). Furthermore, the edge of the door and the three vertical edges of the window meet, quite accurately, in a single vanishing point (far above the image). These results demonstrate the extraordinary accuracy of the rendered geometry and the high level of skill exercised by Robert Campin. In fact, if it is difficult to paint in a perspectively correct way, then painting a curved mirror with the degree of accuracy demonstrated in Campin's *St. John* requires an extraordinary effort.

Measuring heights of people. The accuracy in Campin's painting is sufficient for us to apply single-view metrology techniques (Criminisi 2001) to measure the ratio between the heights of the two monks. Criminisi's height-estimation algorithms applied to the image in figure 7a produce the ratio $h_1/h_2 = 1.38$ between the heights of the two monks, thus proving that the difference in the monks' imaged height is due not only to perspective effects (reduction of farther objects) but

FIGURE 7. Accuracy of Campin's Mirror: (a) Rectified image of mirror from spherical assumption. (b) Straight lines have been superimposed on (a) to demonstrate the accuracy of the rectification process. (c) The accuracy of the rectified image is sufficient for extracting metric information from the painting. By applying single-view metrology techniques, the ratio between the heights of the two monks is computed to be $h_1/h_2 = 1.38$.

also to a genuine difference in their height. This discrepancy in height could be explained by the possibility that the farther monk is kneeling, but the different scales of their heads suggest that Campin's otherwise meticulous scaling of objects in the mirror image has gone awry in this small detail.

Accuracy of van Eyck's Arnolfini Portrait

Inconsistency of mirror geometry. As can be observed in figures 5b, c, d, whatever the parameters used (e.g., the sphere radius), it does not seem possible to simultaneously straighten all the edges in Jan van Eyck's mirror (see also fig. 8). The most striking error can be observed in the leftmost edge of the large window in the rectified images. In fact, whereas it is possible to straighten up the central and farthest vertical edges of the window frame, together with the edges of other objects in the scene, the leftmost window edge remains curved. Figure 8 illustrates this concept more clearly. In figure 8b, some of the relevant edges have been marked. Note that after rectification, the leftmost edge of the window remains visibly curved. This can be interpreted as a lack of geometric consistency: that is, it is not possible for a physical parabolic or spherical mirror to produce the kind of image observed in the original painting (fig. 2 a, b).

Figure 9 shows the effect of rectifying the reflected image by the spherical model with different values of the sphere radius (the mirror radius r). As can be seen, none of these values can straighten all the edges at the same time. The same problem arises with the parabolic or generic radial model assumptions. Incidentally, as predicted by the laws of optics, changing the radius of the mirror has the effect of changing the focal length of the mirror-camera system, that is, the zooming effect (see fig. 9b).

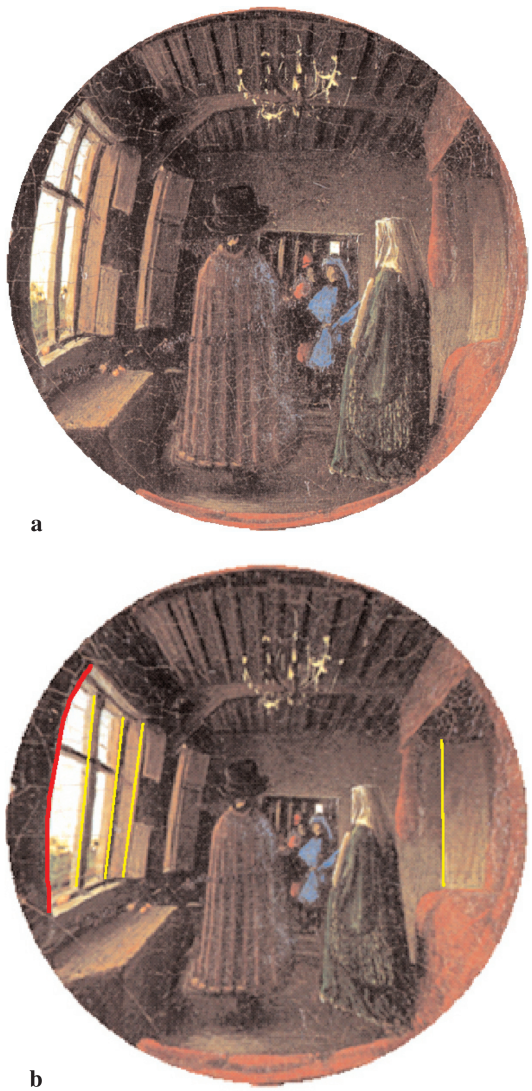

a

b

FIGURE 8. Accuracy of J. van Eyck's Mirror: (a) Original image. (b) The best possible rectification with spherical mirror assumption. Most edges have been correctly straightened, but the leftmost edge of the window has remained noticeably curved, indicating an inconsistency in the geometry of the painted mirror.

Correcting inaccuracies in van Eyck's Portrait. Further inconsistencies characterize the geometry of Jan van Eyck's mirror. In figure 8b, we observe that the bottom edge of the wooden bench and the bottom edge of the woman's gown do not appear to be horizontal (as they are likely to be). Instead, these edges are roughly horizontal in the original painting (fig. 8a). However, this fact is physically impossible: these edges, because of their proximity to the boundary of a convex mirror, should appear curved and oriented at an angle. To explain this concept in easier visual terms, we have run the experiment illustrated in figure 10, taking the following steps:

1. Rectify the original mirror (fig. 10a) using the process described in the previous section and spherical assumption. The resulting rectified image (fig. 10b) shows the window, bench, and gown artifacts mentioned earlier.
2. Straighten these three edges manually via an off-the-shelf image-editing software. The resulting image (in fig. 10c) is perspectively and physically "correct" in that the projected straight 3-D edges are now straight and the bench and gown edges (circled with a dotted line) are horizontal.
3. Warp the image in figure 10c using the inverse of the spherical transformation employed in the first step to obtain the image in figure 10d.

The resulting image (fig. 10d) illustrates what the artist "should have painted." When compared with the original painting, the image in figure 10d is more consistent with the laws of optics applied to a spherical convex mirror and general assumptions about straight lines in an indoor environment. In figure 10d, the bottom edge of the bench more realistically follows a sloped curve, as does the bottom of the woman's gown. Furthermore, the curvature of the leftmost edge of the large window is much less pronounced

FIGURE 9. The Effect of Different Radii: Rectifying J. van Eyck's mirror by means of spherical assumption and varying the radius of the sphere. Despite all efforts, there is no single value of the mirror radius that can straighten all the edges at the same time.

than in the original painting. Overall, the image in figure 10d looks more consistent. In the interest of clarity, figure 11 shows some of the images in figure 10, with the edges of interest marked.

The preceding analysis suggests the possibility that the artist has deliberately altered the geometry of some objects in the painted mirror. The question is: why would Jan van Eyck have done so? Plausible answers may be that (a) the artist accentuated the curvature of the leftmost edge of the window to convey a stronger sense of the bulge of the mirror (our analysis shows that Campin's mirror is bulgier than that of van Eyck's) and (b) the artist may have painted the edges of the bench and gown as straight to make them look less strange, that is, to accentuate reality.

CONCLUSION

We have conducted a rigorous geometric analysis of the convex mirrors painted in the St. John panel from the Heinrich von Werl Triptych and Jan van Eyck's *Arnolfini Portrait*. It must be stressed that the results were obtained from direct analysis of the original paintings, without unnecessary (and often unconvincing) assumptions about the represented scene, such as the position and size of depicted objects or heights of people. Finally, all the assumptions involved in this analysis have been made explicit, cross-validated, and verified from a scientific as well as a historical point of view. The computer vision algorithms we used have allowed us to analyze the shapes of the mirrors, compare them, and transform the reflected images into ones that correspond with orthodox perspective. We are thus provided with new views from the back of the depicted rooms. A rigorous comparison between the two analyzed paintings leads to the conclusions that

- the geometry of Campin's mirror is astonishingly good and is considerably more accurate than that of Jan van Eyck, which is nevertheless quite impressive;
- Campin's mirror appears to be bulgier than van Eyck's;
- it appears that Jan van Eyck has purposely modified the geometry of the image in certain parts of the painted mirror.

A series of notable consequences in art history flow from these findings, especially with respect to the St. John panel.

The St. John Panel

The viewpoint used by Campin to paint the mirror is located on a vertical axis a small distance outside the right boundary of the mirror (see fig. 12a). The viewpoint's height and distance are more problematic. The rectified view of the mirror produces a point of convergence for the windows at a horizontal level equivalent to eye level of the standing St. John and nearest Franciscan.[2] However, the most prominent verticals in the rectified image (in fig. 12b) undergo upward convergence toward a second, single vanishing point, suggesting that the painter's eye level was below the central axis of the mirror. Experimenting with an actual spherical mirror has confirmed that a lower viewpoint, relatively close to the mirror, does indeed produce the effect of a convergence for the orthogonals on or close to the central horizontal axis of the mirror. This lower viewpoint is consistent with that of the interior as a whole as represented on the panel, which is clearly seen from a position closer to the eye level of the kneeling donor and perpendicular to a point that lies decisively within the lost central panel (i.e., well outside the right boundary of the St. John panel). What we cannot tell decisively is how close the artist's viewpoint was located to the mirror in the actual painting.

The foregoing observations suggest the following model for the artist's procedure. To paint the staged scene in his convex mirror, Campin moved closer to the mirror than is inferred from the original setup as portrayed as a

FIGURE 10. Correcting J. van Eyck's Portrait: (a) The original painted mirror. (b) The rectified mirror (see text for details). (c) The leftmost window edge is manually straightened. The bottom edges of the wooden bench and that of the woman's gown are manually made horizontal. (d) Image (c) is warped back by the inverse of the spherical model used to generate image (b). The resulting image (d) is the "corrected" version of van Eyck's mirror, one that more closely obeys the laws of optics.

whole in the panel. He observed it from a position just sufficiently outside the right margin of the mirror to avoid the problem of his head's occluding significant portions of the image. It may be that he used the door as a "shield" from which he could observe the image with a minimum of intrusion. He also observed the view in the mirror from an angle below its horizontal midline, that is, effectively tilting the picture plane or "camera" so as to cause the upward convergence.

The fundamental veracity of the painted reflection extends to such details of the small portion of the donor's trailing drapery visible beyond the edge of the door, though the donor himself is hidden. Whereas it may be possible to envisage an imaginary view in a spherical mirror, it is inconceivable that such consistently accurate optical effects could have been achieved by a simple act of the imagination. We are drawn to what seems to be the inescapable conclusion that the artist has directly observed and recorded the

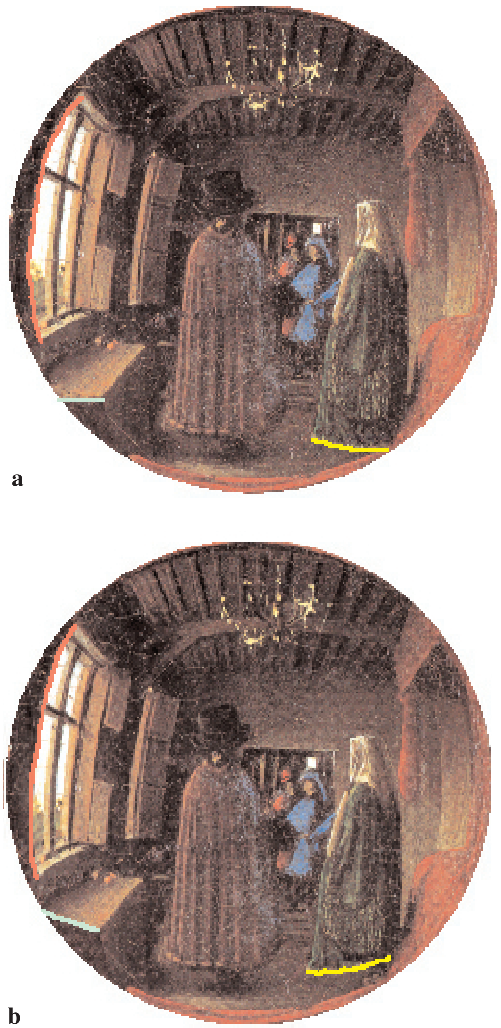

FIGURE 11. Correcting J. van Eyck's Portrait: (a) The original painting (see also fig. 10a) with three marked edges. (b) Our corrected version (see also fig. 10d) with the corrected edges marked. The leftmost edge of the large window is less pronounced in the corrected version, and the bench and gown bottom edges are correctly curved, consistent with the laws of optics.

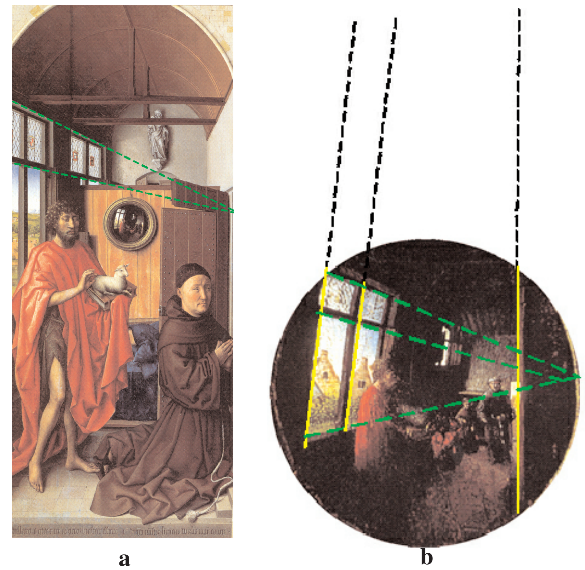

FIGURE 12. Vanishing Points in Campin's St. John: (a) The edges of the left window suggest a vanishing point a short distance to the right of the right edge of the painting. (b) In the image of the rectified mirror, the horizontal window edges converge roughly to a point on the right boundary of the mirror, whereas the verticals converge upward. The upward convergence suggests a tilting of the "camera/panel."

effects visible when actual figures and objects are located in a specific interior. Such a result means that, at some point, models must have been posed in exactly those positions occupied in the painting while dressed in appropriate costumes—with the exception of the two Franciscan witnesses, who are in their normal habits. If this particular picture was achieved by means of such a stage-managed arrangement, we may reasonably believe that other Netherlandish masterpieces of naturalism were accomplished in the same way.

The question of how the painter achieved such optical accuracy remains. To some extent, given the condensing of an image within the small compass of the mirror, and with the circular frame serving as a ready point of precise reference, the painter is faced with an easier task than the one he faces when he surveys the whole scene at its normal scale. On the other hand, if the optical image in the mirror could itself be optically projected, the process would be facilitated. If the mirror were to be projected from the distance implied in the painting, the definition of the image with contemporary equipment would have come nowhere near delivering the required results. However, if, as suggested, the artist portrayed the mirror in a separate act, from a closer position, an optical projection remains a possibility. We may note in passing that such an astonishing achievement suggests a major mind and talent at work. The recent tendency to see the von Werl wings as the works of a contemporary pasticheur of Campin's and van Eyck's seems entirely unjustified.

The Arnolfini Portrait

Judged on its own merits, the accuracy of the image in van Eyck's convex mirror is striking. It is only by comparison with Campin's that it appears less than notable. Even with its faults, it is remarkable enough to support (though less conclusively) the hypothesis that the image in the mirror was painted directly from a stage-managed setup in an actual space. The most notable of the artist's departures from optical accuracy, the line of the nearest edge of the bench by the window, and the hem of the woman's dress, can both be explained as instinctive responses to the extreme effects visible in the mirror. In other words, the painter has very selectively played to what we expect the image to look like rather than precisely following its actual appearance. Such moves are common enough in Italian perspective pictures. It may also be that the border regions of the mirror used by van Eyck were less optically sound than the middle sections. It should also be noted that the mirror reflection in his painting is portrayed from a viewpoint perpendicular to or very near the center of the mirror, which suggests that he did not adopt the expedient we have suggested for Campin, that is, moving closer to the mirror and slightly to the side to portray the reflected scene more readily.

The Hockney Hypothesis and Wider Considerations

The finding that Campin and van Eyck seem to have worked with a tableau vivant, using posed figures, actual objects, and real interiors, does nothing to negate Hockney's hypothesis that optical projection from a concave mirror or lens was used to make the paintings. But neither does it prove it. We are still left with the possibility that the painters "eyeballed" their scenes with miraculous accomplishment. In any event, it remains to be demonstrated that the images in the painted convex mirrors could have been projected, using contemporary equipment, with a quality such as to provide such mini-masterpieces within each masterpiece. On the other hand, the fact that the Campin mirror may have been represented using a different viewpoint from the interior in the panel as a whole lends some support to Hockney's "many windows" theory. Our findings may be regarded as broadly permissive for Hockney's theory, but they stop short of providing incontrovertible support. The main implications of our work lie elsewhere.

The chief of the wider considerations increases the radical distance we can discern between the wonders of illusion being accomplished in the Netherlands and Italy at the very same time. Italian perspectival pictures, as pioneered by Masaccio, were achieved through the synthetic construction of geometrical spaces according to preconceived designs. Within the geometrical containers, the Italian painters then inserted to scale separately studied figures or small groups (sometimes beginning with nude or near-nude studies). By contrast, the leading Netherlandish reformers of representation may be seen to be working as literally as possible with what they could see. They are true champions of what Ernst Gombrich has described as "making and matching." Such literalism of matching, much to the taste of nineteenth-century critics like Ruskin, has not found favor with recent art historians, who prefer more complicated explanations. It is almost as if we are precociously entering the realm of Courbet, the nineteenth-century realist, who declared, "Show me an angel and I will paint one." It seems that five centuries earlier, Campin seemed to be saying, "To paint a St. John, I need to see one."

NOTES

Color versions of the illustrations in this article will soon be available online at http://www.heldref.org

1. Straight scene edges are imaged as straight lines by a pinhole camera. In addition, strictly speaking, the corrective transformation can be done with the knowledge of only the mirror shape if the imaging system has a single point of projection. Otherwise, additional knowledge such as scene depth has to be known as well.

2. The horizontal edges of the window (in fig. 12b) converge roughly at a point on the right boundary of the mirror, approximately at the eye level of the reflections of St. John and the closest Franciscan in the mirror.

REFERENCES

Benosman, R., and S. B. Kang, eds. 2001. *Panoramic vision: Sensors, theory, and applications*. New York: Springer-Verlag.

Canny, J. F. 1986. A computational approach to edge detection. *Pattern Analysis and Machine Intelligence* 8(6): 679–98.

Criminisi, A. 2001. *Accurate visual metrology from single and multiple uncalibrated images*. Distinguished Dissertation Series. London: Springer-Verlag.

Devernay, F., and O. D. Faugeras. 1995. Automatic calibration and removal of distortion from scenes of structured environments. In *Proceedings of SPIE*, vol. 2567, 62–72. San Diego, CA.

Faugeras, O. D. 1993. *Three-dimensional computer vision: A geometric viewpoint*. Cambridge, MA: MIT Press.

Hartley, R. I., and A. Zisserman. 2000. *Multiple view geometry in computer vision*. Cambridge: Cambridge University Press.

Hockney, D. 2001. *Secret knowledge: Rediscovering the lost techniques of the old masters*. New York: Viking Studio.

Kemp, M. 1990. *The science of art*. New Haven: Yale University Press.

Lin, S. S., and R. Bajcsy. 2001. True single view point cone mirror omnidirectional catadioptric system. In *Proceedings of the International Conference on Computer Vision*, vol. 2, 102–7. Vancouver, BC.

Nayar, S. 1997. Catadioptric omnidirectional camera. In *Computer Vision and Pattern Recognition*, 482–88. Puerto Rico.

Press, W., B. Flannery, S. Teukolsky, and W. Vetterling. 1988. *Numerical Recipes in C*. Cambridge: Cambridge University Press.

Steadman, P. 2001. *Vermeer's camera: Uncovering the truth behind the masterpieces*. Oxford: Oxford University Press.

Svoboda, T., T. Pajdla, and V. Hlavac. 1998. Epipolar geometry for panoramic cameras. In *European Conference on Computer Vision*, 218–31. Freiburg, Germany.

Swaminathan, R., M. D. Grossberg, and S. K. Nayar. 2001. Non-single viewpoint catadioptric cameras: Geometry and analysis. Technical report, Dept. of Computer Science, Columbia University.

APPENDIX

This section investigates the geometry of mirrors and presents an algorithm for rectifying reflected images according to the three different mirror models presented in the body of the article.

Assumption 1: Parabolic mirror

The parabolic shape assumption is convenient because it allows us to reuse a considerable number of techniques developed in the field of catadioptric imaging (Benosman and Kang 2001; Nayar 1997). In fact, it is possible to treat the painted mirror as part of a parabolic mirror and orthographic camera catadioptric system (fig. 3a). Such a system behaves as a single-optical-center acquisition device,[1] and the corresponding rectification equations are straightforward.

The diagram in figure A1 describes a cross-section of the basic parabolic mirror-camera system; simple algebra leads to

appendix continues

the basic equations for undistorting the reflected image. Given the pixel coordinate u_d of the distorted point in the original image, we can compute the coordinate u_c of the corresponding corrected point by applying the equations that follow:

$$u_c = \frac{\Delta}{w} u_d \text{ with } w = \frac{h^2 - r^2}{2h}, \qquad \text{(A1)}$$

where Δ is the distance of the camera/panel from the mirror focus, h is the parabolic parameter, and r is the distance between the focus and the point P_m.

Note that the above equations have been derived for the 2-D mirror cross-sections only, but because the mirrors are assumed to be perfectly symmetrical surfaces of revolution, the derivation of the formulas for the complete 3-D case is straightforward.

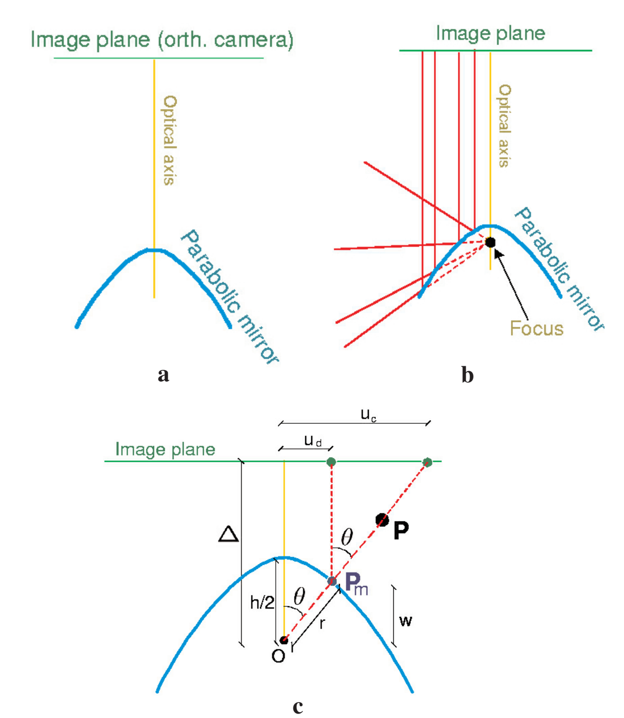

FIGURE A1. The Geometry of Parabolic Mirrors: (a) Cross-section of a parabolic mirror and the image (panel) plane. (b) All the rays coming from the orthographic camera are reflected by the mirror as rays that intersect in a single point, namely, the *focus* of the mirror. The existence of a single focus makes parabolic mirrors particularly convenient (see text for details). (c) A point P in the 3-D scene is reflected by the mirror into its corresponding image point at u_d distance from the center of the image. In other words, u_d is the observed (distorted) coordinate. The corresponding undistorted position is labeled u_c. The center of projection (focus) is denoted O and the parabolic parameter *h*.

Assumption 2: Spherical mirror

Unlike parabolic mirrors, spherical mirrors do not present a single point of projection but rather a whole locus of viewpoints, called a *caustic surface* (Swaminathan, Grossberg, and Nayar 2001). The rectification process becomes more complicated here, for it is now necessary to know the 3-D shape and depth of the surrounding environment.

However, the equations for rectifying the image reflected by a spherical mirror simplify if we assume that the visualized points (e.g., the point P in fig. 3b) lie at infinity —or, more realistically, far away from the mirror itself. Obviously, from a historical point of view, this latter assumption is more plausible and effectively operates in our two chosen examples.

From an analysis of the diagrams in figure A2 and some simple algebra, we can derive the following rectification equations:

$$u_c = u_d + (\Delta - \gamma)\frac{2u_d\gamma}{r^2 - 2u_d^2} \text{ with } \gamma = \sqrt{r^2 - u_d^2}, \qquad \text{(A2)}$$

where Δ is the distance (large) of the camera/panel from the center of the sphere, r is the radius of the sphere, u_d is the coordinate of the distorted image point, and u_c is the coordinate of the corresponding corrected point.

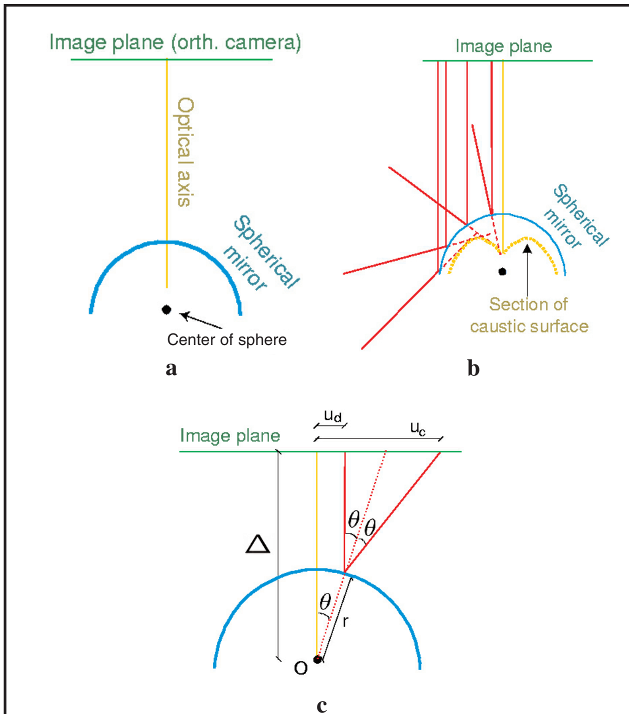

FIGURE A2. The Geometry of Spherical Mirrors: (a) Cross-section of a spherical mirror and the image (panel) plane. (b) All the rays coming from the orthographic camera are reflected by the mirror as rays that do not intersect in a single point, but rather in a *caustic surface.* This characteristic makes spherical mirrors harder to use, unless the 3-D point P is assumed to be an infinite distance (or very far) from the center of the mirror. (c) u_d and u_e are the observed and corrected coordinates, respectively.

Assumption 3: The convex mirror induces radial distortion on the image plane

In this discussion, rather than reasoning about the 3-D shape of the mirror, we try to model directly the distorting effect that a convex mirror produces on the reflected image.

As shown in figures 3c, d, straight scene edges appear curved in the image reflected by a convex mirror. Here, we make the plausible assumption that the distorting transformation can be modeled as a standard radial distortion model. Radial distortion typically occurs in cameras with very wide angle lenses (Devernay and Faugeras 1995). Cheap Web cameras are a good example.

A generic radial model can be described by the following equations (see fig. A3):

$$p_c = c + f(l)(p_d - c), \tag{A3}$$

where p_d is a 2-D point on the plane of the original image and p_c is its corresponding corrected point. The quantity $l = d(p_d, c)$ is the distance in the image plane of the point p_d from the center of the image c, and the function $f(l)$ is defined as

$$f(l) = 1 + k_1 l + k_2 l^2 + k_3 l^3 + k_4 l^4 + \ldots .$$

Therefore, if the k_i parameters of the above model were known, this radial transformation could be applied to the reflected image to "correct" the distortion and obtain the corresponding perspectival image. It can be shown that a generic radial model subsumes both the parabolic and spherical models, by expanding the right-hand sides of equations (A1) and (A2).

We can choose the complexity of the radial model by selecting how many and which of the k_i parameters to use. To keep the model simple, we have chosen to use only the k_1 and k_2 parameters, as is shown in the discussion that follows.

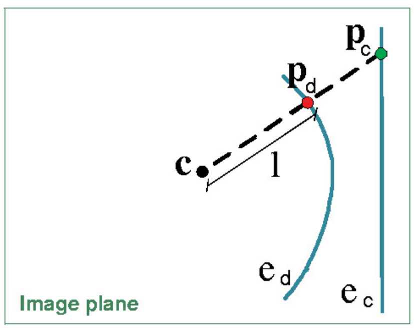

FIGURE A3. Mirror Rectification via a Generic Radial Model: A radial transformation maps a point p_d in the original image into the corresponding point p_c in the rectified (corrected) image by applying a radial transformation centered in the image center c (see text for details). The complete rectified image is obtained by applying this point-based transformation to all the points in the original (distorted) image. Note that the radial transformation straightens curved edges—for example, the edge e_d in the original image is transformed in the corresponding edge e_c (straightened) in the rectified image. (See also figs. 3c, d.)

Image rectification

Estimating the parameters. In the previous three cases, to rectify the reflected image it is necessary first to estimate optimal values for the parameters of the described mathematical models. The parameters that need to be estimated in the three cases are summarized in table A1.

We expect the rectification algorithm to straighten the curved images of straight scene edges such as window frames and door frames. Therefore, the best set of rectification parameters may be computed as the set of values that best straightens some carefully selected edges in the original (distorted) images. Our semiautomatic rectification algorithm proceeds as follows:

1. a Canny (1986) edge detection algorithm automatically detects sub-pixel accurate image edges;
2. a user manually selects some curved edges that correspond to straight 3-D scene lines;
3. a numerical optimization routine, such as Levenberg-Marquardt (Press et al. 1988), automatically estimates the parameter values that best straighten the selected edges.

TABLE A1. Mirror Shapes and Corresponding Parameters

Assumption	Parameters to be computed
A1. parabolic mirror	h (parabolic parameter)
A2. spherical mirror	r (radius of sphere)
A3. generic radial model	k_1, k_2

NOTE

1. The interested reader may refer to Shih-Schon Lin and Ruzena Bajcsy (2001) and Tomas Svoboda, Tomas Pajdla, and Vaclav Hlavac (1998) for discussions on other single viewpoint catadioptric systems.