Reflections of Reality in Jan van Eyck and Robert Campin

A. Criminisi†, M. Kemp‡ and S. B. Kang⋆

†Microsoft Research, UK, ‡University of Oxford, UK, ⋆Microsoft Research, USA

antcrim@microsoft.com

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. Using novel, rigorous techniques to analyse the geometric accuracy of the mirrors, unexpected findings emerge, which radically affect how we see the paintings as being generated. We focus on Jan van Eyck's Arnolfini Portrait, and the Heinrich von Werl Triptych, here re-attributed to Robert Campin. The accuracy of the convex mirrors depicted in these paintings is assessed by applying mathematical techniques drawn from computer vision. The proposed algorithms allow us also to "rectify" the image in the mirror so that it becomes a normalised projection, thus providing us with a second view from the back of the painted room. The plausibility of the painters' renderings of space in the convex mirrors can be assessed. The rectified images can be used for purposes of three-dimensional reconstruction as well as measuring accurate dimensions of objects and people. The surprising results presented in this paper cast a new light on the understanding of the artists' techniques and their optical imitation of seen things, and potentially require a re-thinking of the foundations of Netherlandish naturalism. They also suggest that the von Werl panels should be re-instated as autograph works by R. Campin. Additionally, this research represents a further attempt to build a constructive dialogue between two very different disciplines: computer science and history of art. Despite their fundamental differences, the procedures followed by science and art history can learn and be enriched by each other.*

1. Introduction

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 Flemalle (now generally identified with the documented painter, Robert Campin), has always been recognised 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 15th century. Amongst 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 looks the normal ensemble

a b c d

Figure 1: Convergence of orthogonals in paintings: **(a)** Masaccio, *Trinità con la Vergine e San Giovanni*, approx. 1426, Museo di Santa Maria Novella, Firenze, Italy. **(b)** The dominant orthogonals (superimposed in yellow) converge at a single vanishing point. **(c)** Jan van Eyck, *Madonna in the Church*, 1437-1439, Staatliche Museen zu Berlin, Gemaldegalerie, Berlin. **(d)** The orthogonals (marked) do *not* converge at a single vanishing point.

of a domestic or ecclesiastical interior. It has also been long recognised that the coherence of the spaces in Netherlandish painting did not rely upon any dogmatic rule, such as the convergence of orthogonals to the "centric" or vanishing point, which provided the basis for Italian perspective in the wake of Brunelleschi (*cf.* 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. The two paintings on which we shall concentrate, the so-called *Arnolfini Wedding* by Jan van Eyck (figs. 2a,b) and the *St. John the Baptist with a Donor* from the Heinrich van Werl triptych (figs. 2c,d), both demonstrate the way that convincing spaces could be 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. Recently, David Hockney has argued that optical projections provide the key. Initially disposed to argue that the artist relied very directly on images projected on to white surfaces by lenses or concave mirrors, Hockney now places more emphasis upon the projected image as the key breakthrough in revealing what a 3-D array looks like when flattened by projection. In any event, he rapidly came to see that an image like Jan 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 he 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 personalised polemic which dogmatically uses varieties of technical evidence in the service of pre-determined stances, without much sense of improvisatory procedures which typify artists' uses of the tools at the disposal.

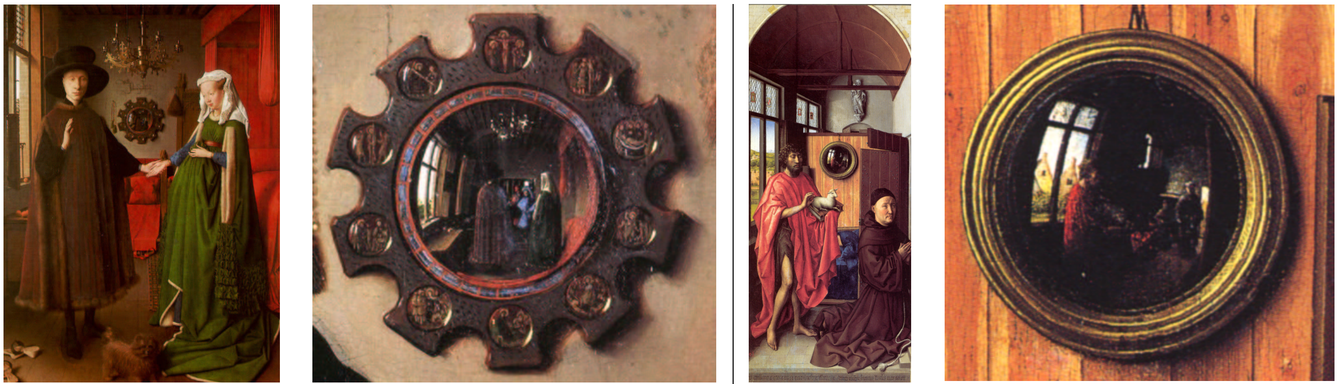

a b c d

Figure 2: Original paintings analysed in this paper: (a) J. van Eyck, *Arnolfini Wedding*, panel, 1434, National Gallery, London. **(b)** Enlarged view of the convex mirror in (c). **(c)** R. Campin, *St. John with a Donor*, panel, 1438, Museo del Prado, Madrid. **(d)** Enlarged view of the convex mirror in (a).

Our intention is to enquire about Netherlandish naturalism along a different route. We will be undertaking 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 will be asking 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 how we conceive the painters as composing their images. What emerges from our analysis seems to us to have unavoidable implications for how they stage-managed their ensembles of figures and objects. Along the way, we will 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.

Before we begin our analysis, it will be a good idea to introduce the two paintings which will be serving as our witnesses.

1.1. The Arnolfini Wedding

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 Nicolao Arnolfini and Giovanna Cenami. Arnolfini was a successful merchant from Lucca and 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 their 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, since a beautifully inscribed graffito on the end wall records the date, 1434, and the fact that "Johannes de eyck fuit hic" (Jan van Eyck was here).

1.2. The St. John Panel

The other painting, the *St. John with a Donor* in the Prado Museum in Madrid, uses the oil medium no less brilliantly than van Eyck. It is one of the two surviving wings of an altarpiece commissioned, as the inscription tells us, by Heinrich van 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 Heinrich 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 signalling the donor's hoped-for translation to the heavenly dwelling of the Virgin and saints after his death.

In the technical analyses that follow, we provide a geometric analysis of the two painted mirrors on the basis of the optics of spherical and parabolic mirrors, and transform the images into normalised 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.

2. Geometric analysis of painted convex mirrors

This section presents some simple but rigorous techniques for the analysis of the geometric accuracy of convex mirrors in paintings. The algorithms employed here build upon the vast computer vision literature, with specific reference to the field of compound image formation, termed catadioptric imaging [1, 10]. 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 composed by a mirror and an orthographic camera as illustrated in figs. 3a,b.

It is well known that curved mirrors produce distorted reflections (*cf.* figs. 3c,d) of the environment they are in. For instance, straight scene lines become curved (*e.g.* 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 to produce a perspectively corrected image; *i.e.* an image as if it was acquired by a camera with optical centre in the centre of the mirror[1]. But, while 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 the product of the hand of a skilled artist; and, even when it appears correct, its geometry may differ from that produced by a real camera.

In this paper we propose novel techniques for: i) Assessing the geometric accuracy of a painted mirror and ii) Producing rectified, perspective images from the viewpoint of the mirror itself.

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

[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.

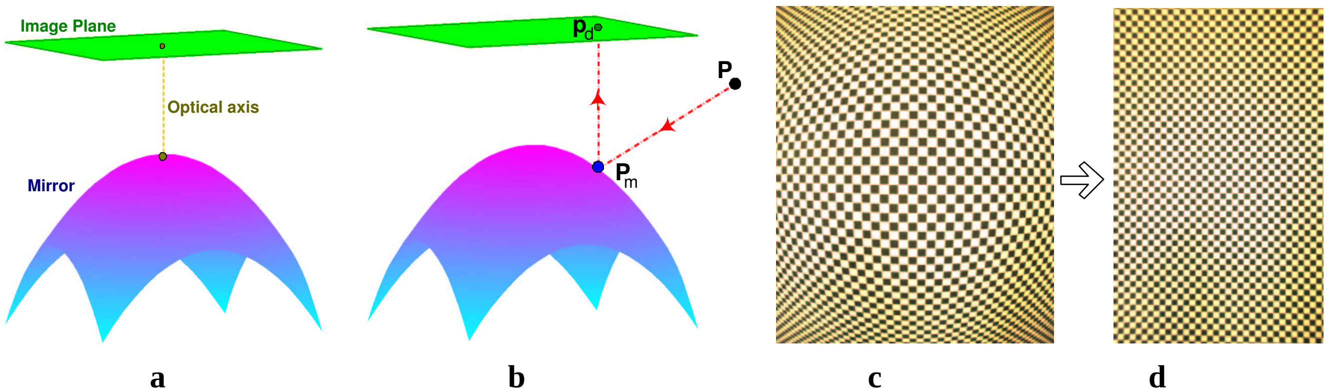

Figure 3: Catadioptric imaging model: (a) The mirror-panel pair is modelled as a catadioptric acquisition system, composed by a curved mirror and an acquiring camera. **(b)** A ray emanating from a point $\mathbf{P}$ in the scene is reflected by the mirror onto the corresponding image (panel) point $\mathbf{p}_d$. The ray is reflected off the mirror surface in the point $\mathbf{P}_m$. **(c)** A square grid reflected by a convex mirror appears barrel-distorted. **(d)** A goal of this paper 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.

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 one.
3. Measuring the discrepancy between the rectified image and the "perspectivally correct" image to assess the accuracy of the painted mirror.

As it will be clearer in the remainder of the paper, these basic geometric tools will help us cast a new light on the analysed 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 A1**: The mirror has a parabolic shape,
- **Assumption A2**: The mirror has a spherical shape,
- **Assumption A3**: 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 which can "remove" the distortion in the reflected image and rectify it into a perspective image.

It must be noticed, though, that from a historical point of view the parabolic can be discarded since 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 analyses that follow.

2.1. *Assumption A1*: Parabolic mirror

The parabolic shape assumption is convenient since it allows us to re-use considerable amount of techniques developed in the field of catadioptric imaging [1, 10]. 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-centre acquisition device[2] and the corresponding rectification equations are straightforward.

The diagram in fig. 4 describes a plan-view cross-section of the basic parabolic mirror-camera system and simple algebra leads to 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 formulae below:

$$u_c = \frac{\Delta}{w} u_d \qquad \text{with} \qquad w = \frac{h^2 - r^2}{2h}, \tag{1}$$

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 $\mathbf{P}_m$.

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

2.2. *Assumption A2*: Spherical mirror

Unlike parabolic mirrors, spherical mirrors do not present a single point of projection but a whole locus of viewpoints which is called a *caustic surface* [14]. The rectification process becomes more complicated here, since it is now necessary to know the three-dimensional 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 $\mathbf{P}$ in fig. 3b) lie at infinity; or more realistically, far away from the mirror itself. From a historical point of view this latter assumption is, obviously, more plausible and effectively operates in our two chosen examples.

From an analysis of the diagrams in fig. 5 and some simple algebra we can derive the following rectification equations:

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

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

2.3. *Assumption A3*: The convex mirror induces radial distortion on the image plane

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

[2] The interested reader may refer to [9, 13] for other single-viewpoint catadioptric systems

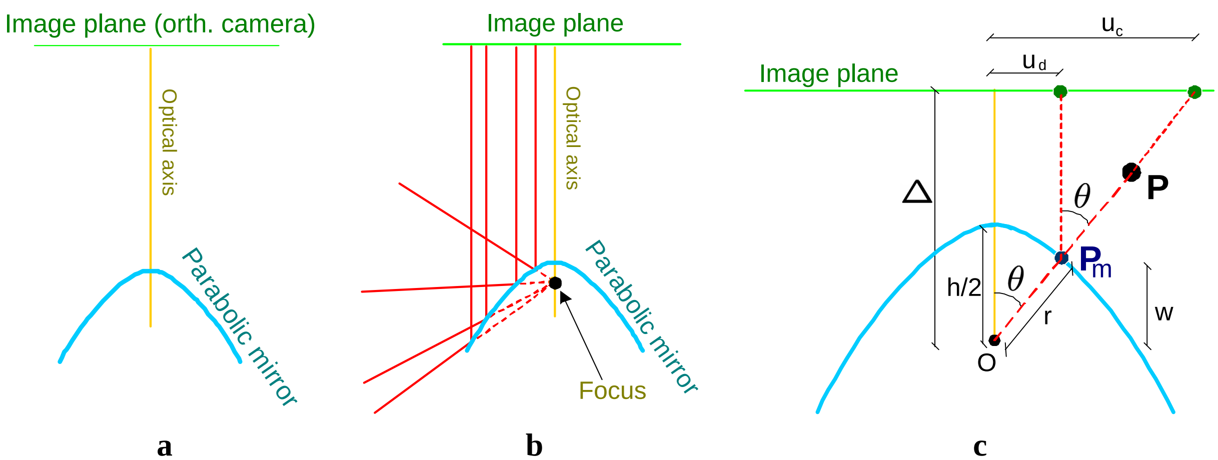

Figure 4: The geometry of parabolic mirrors: (a) Plan-view cross-section of a parabolic-mirror and the image (panel) plane. **(b)** All the rays coming from the orthographic camera are reflected by the mirror as rays which intersect in a single point, namely the *focus* of the mirror. The existence of a single focus makes parabolic mirrors particularly convenient, see text for details. **(c)** A point $\mathbf{P}$ in the three-dimensional scene is reflected by the mirror into its corresponding image point at u_d distance from the centre of the image. In other words, u_d is the observed (distorted) coordinate. The corresponding, undistorted position is labelled u_c. The centre of projection (focus) is denoted $\mathbf{O}$ and the parabolic parameter h.

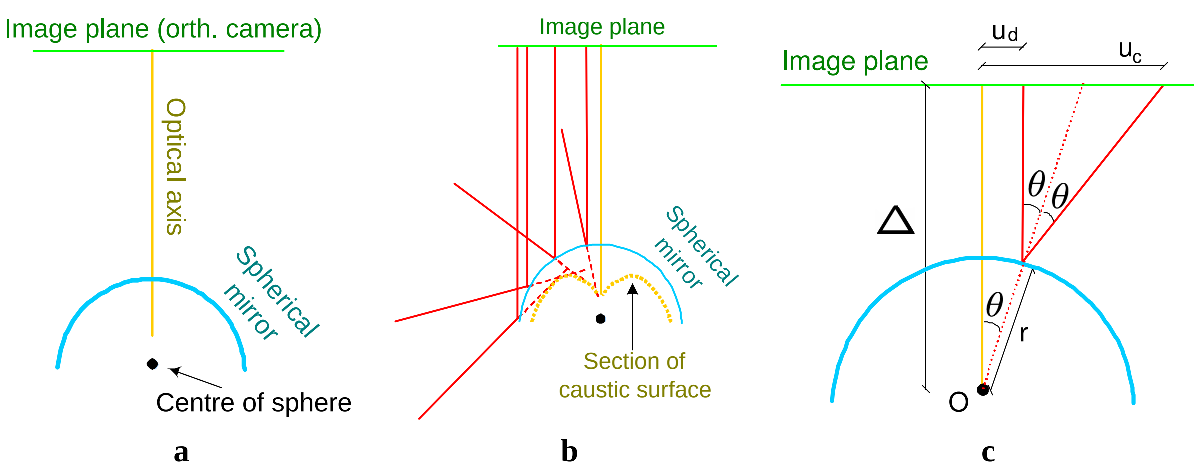

Figure 5: The geometry of spherical mirrors: (a) Plan-view cross-section of a spherical-mirror and the image (panel) plane. **(b)** All the rays coming from the orthographic camera are reflected by the mirror as rays which do *not* intersect in a single point, but in a *caustic surface*. This characteristic makes spherical mirrors harder to use, unless the 3D point $\mathbf{P}$ is assumed to be at infinite distance (or very far) from the centre of the mirror. **(c)** u_d and u_c are the observed and corrected coordinates, respectively.

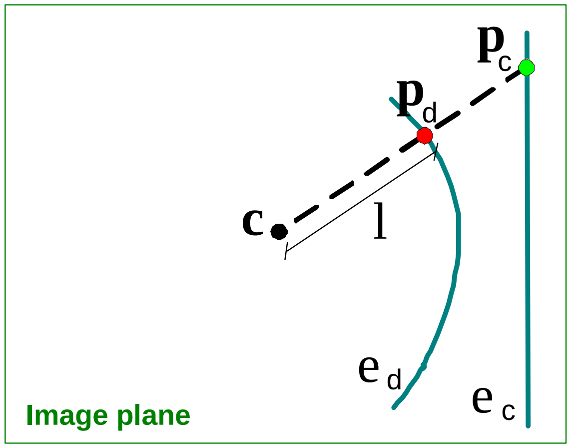

Figure 6: Mirror rectification via a generic radial model: A radial transformation maps a point $\mathbf{p_d}$ in the original image into the corresponding point $\mathbf{p_c}$ in the rectified (corrected) image by applying a radial transformation centered in the image centre $\mathbf{c}$, see text for details. The complete rectified image is obtained by applying this point-based tranformation to all the points in the original (distorted) image. Notice that the radial transformation straightens curved edges, *e.g.* 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.

As shown in figs. 3c,d, in the image reflected by a convex mirror straight scene edges appear curved. Here we make the plausible assumption that the distorting transformation can be modelled as a standard radial distortion model. Radial distortion typically occurs in cameras with very wide-angle lenses [4]. Cheap web-cameras are a good example.

A generic radial model can be described by the following equations (*cf.* fig. 6):

$$\mathbf{p}_c = \mathbf{c} + f(l)(\mathbf{p}_d - \mathbf{c}), \tag{3}$$

where $\mathbf{p}_d$ is a 2D point on the plane of the original image and $\mathbf{p}_c$ is its corresponding, corrected point. The quantity $l = d(\mathbf{p}_d, \mathbf{c})$ is the distance in the image plane of the point $\mathbf{p}_d$ from the centre of the image $\mathbf{c}$ and the function $f(l)$ is defined as

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

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 perspective image. It can be shown that a generic radial model subsumes both the parabolic and spherical models. This is done by expanding the right hand sides of (1) and (2).

We can choose the complexity of the radial model by selecting how many and which of the k_i parameters to use. In the results below, in order to keep the model simple, we have chosen to use only the k_1 and k_2 parameters.

2.4. Image rectification

Estimating the parameters. In all the three cases described above, in order 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 1.

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

Table 1: Mirror shapes and parameters to be computed.

As mentioned previously, we expect the rectification algorithm to straighten the curved images of straight scene edges such as window and door frames. Therefore, the best set of rectification parameters may be computed as the set of values which best straightens some carefully selected edges in the original (distorted) images. Our semi-automatic rectification algorithm proceeds as follows:

1. A Canny Edge Detection [2] algorithm automatically detects sub-pixel accurate image edges;
2. A user manually selects some curved edges which correspond to straight 3D scene lines;
3. A numerical optimization routine such as Levenberg-Marquardt [11] automatically estimates the parameters values that best straighten the selected edges.

Rectification results. Results of the rectification of Campin and van Eyck's mirrors are illustrated in fig. 7 and fig. 8, respectively.

One first observation is that the three different mirror shape assumptions lead to very similar (but not identical) rectifications. This can be explained by considering the fact that the tips of the parabolic and spherical mirrors are very similar in shape. In fact, in Optics, parabolic assumptions are made for lenses and mirrors when small portions of a spherical surface are being considered. If the visible parts of the mirrors are just the tips, as in these two cases, it is not surprising that the rectification results are very similar. Furthermore, the fact that the spherical assumption and the parabolic one yield very similar rectification results validates the assumption of substantial distances between the 3D scene points and the mirror, which needs to be made in the spherical mirror case. This is an indirect way of assessing the validity of the rectification approach we have undertaken.

Notice that horizontal flipping of the rectified images in figs. 7b,c,d and figs. 8b,c,d would produce the perspective images that an observer would see if standing 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 three-dimensional computer vision techniques [5, 6] to reconstruct complete virtual models of the depicted environment. Alternatively, single-view reconstruction [3] may be applied twice: to the image from the front (the painting itself) and to the image from the back (the rectified mirror reflection) and the two resulting three-dimensional reconstructions merged together to create a single and complete shoe-box-like virtual model of the scene.

Figure 7: Rectification of the convex mirror in Campin's *St. John Panel*: (a) Original (distorted) mirror from fig 2b. **(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 radial model fitting; Notice that in all the rectified images the edges of the window and door have been straightened.

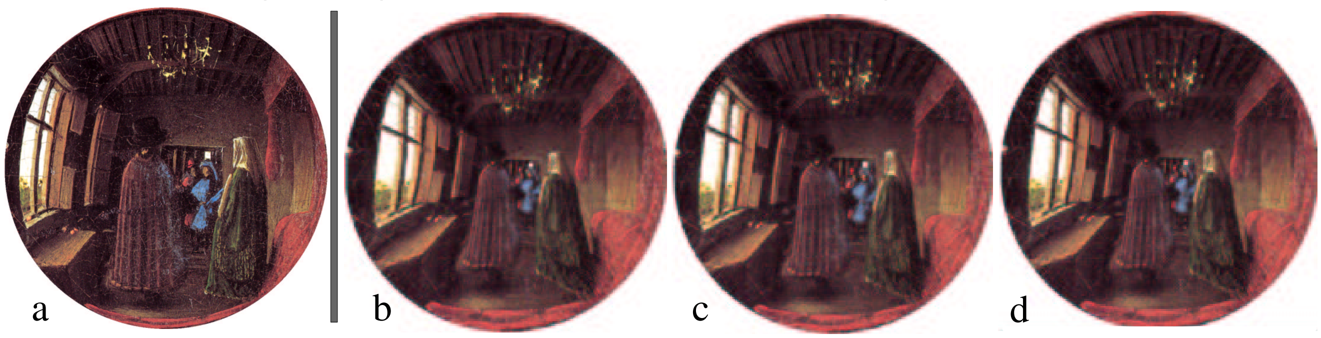

Figure 8: Rectification of the convex mirror in van Eyck's *Portrait*: (a) Original (distorted) mirror from fig 2d. **(b,c,d)** Rectified images. **(b)** Result of rectification via parabolic assumption; **(c)** Result of rectification via spherical assumption; **(d)** Result of rectification via generic radial model fitting;

2.5. Comparing mirror protrusions

As mentioned above, in the case of the spherical mirror assumption, in order to rectify the input, reflection images it is necessary to compute the radius r of the spherical mirror itself (Table 1).

Since 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 only as relative measurements. For instance, we can measure the ratio between the radius r and the radius of the disk in each painting (we call *disk* the visible part of a spherical mirror on the image plane, *cf.* fig. 9a).

Figure 9a shows a plan view of a spherical mirror sticking out a wall with the disk radius clearly marked. Figure 9b shows a comparison between two spherical mirrors with the second mirror (namely mirror 2) bulgier than the first one. We now define a measure of the *protrusion* (bulge) of a spherical mirror as the ratio $\mathcal{P}$ between the disk radius and the radius of the sphere.

$$\mathcal{P} = \frac{r_{disk}}{r} \qquad (4)$$

By inspection of fig. 9b, since $r_2 < r_1$ and the disk radius r_{disk} is the same, then $\mathcal{P}_2 > \mathcal{P}_1$.

From the protrusion measures reported in Table 2.5 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 $\mathcal{P}$) but the

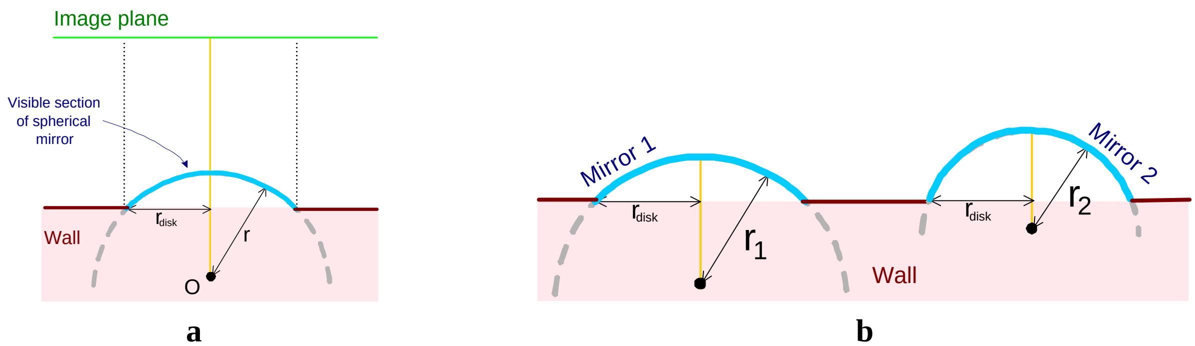

Figure 9: **(a)** We call the visible part of a spherical mirror a *disk*. The quantity r_{disk} is the disk radius. **(b)** A plan view of two spherical mirrors sticking out a wall. Since $r_1 > r_2$ the second mirror sticks out more and thus is bulgier than the first mirror.

Painting	Protrusion factor, $\mathcal{P} = r_{disk}/r$
Campin's *St. John with a donor*	$\mathcal{P} = 0.83$
van Eyck's *Arnolfini Wedding*	$\mathcal{P} = 0.78$

Table 2: "Bulge" estimates for the two paintings.

former is slightly bulgier (greater value of $\mathcal{P}$) than the latter.

This result should be taken with some care. In fact, in order 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 while 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 2.5 could also be used to compute the focal lengths of the two mirrors. In fact, elementary Optics teaches us that the radius of curvature r of a mirror is twice its focal length. Therefore, from equation (4), if we knew r_{disk} we could compute r and therefore $f = r/2$ would be the focal length. Notice, though, that an absolute metric value for r_{disk} can be estimated only after making further assumptions on shapes and sizes of objects that appear in the scene. This kind of assumptions would make sense if, like in the case of a few paintings by Jan Vermeer, some of the depicted objects still exist and can be measured [8, 12]. But in the case of the two paintings analysed in this paper, we feel that assumptions not supported by strong physical evidence may conduct to misleading results on the value of the mirrors focal lengths.

2.6. Accuracy analysis

The previous sections have 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. This section, instead, analyses the rectified images to assess their geometric consistency.

Figure 10: Accuracy of Campin's mirror: **(a)** Rectified image of mirror from spherical assumption; **(b)** Straight lines have been superimposed to (a) to demonstrate the accuracy of the rectification process. **(c)** The accuracy of the rectified image is sufficient for extracting metric information out of the painting. By applying single view metrology techniques the ratio between the heights of the two monks is computed to be $\frac{h_1}{h_2} = 1.38$.

2.6.1 Accuracy of Campin's *St. John Panel*

Figure 10a (identical to fig. 7d) is the rectification of Campin's mirror obtained from the assumption of spherical mirror.

Straightening of curved edges. Figure 10b shows that images of straight scene lines now have become consistently straight. 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). This demonstrates 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, painting a convex 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 [3] techniques to measure the ratio between the heights of the two monks. The height-estimation algorithms in [3] applied to the image in fig.10a produce the ratio $\frac{h_1}{h_2} = 1.38$ between the heights of the two monks; thus proving that the difference in the monks' imaged height is not due just to perspective effects (reduction of farther objects) but 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 suggests that Campin's otherwise meticulous scaling of objects in the mirror image has gone awry in this small detail

2.6.2 Accuracy of van Eyck's *Arnolfini Portrait*

Inconsistency of mirror geometry. As it can be observed in figs. 8b,c,d, whatever the parameters used, it does not seem possible to simultaneously straighten all the edges in Jan van Eyck's mirror (see also fig. 11). The most striking error can be observed in the leftmost edge of the large window in the rectified images. In fact, while 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 left-most window edge

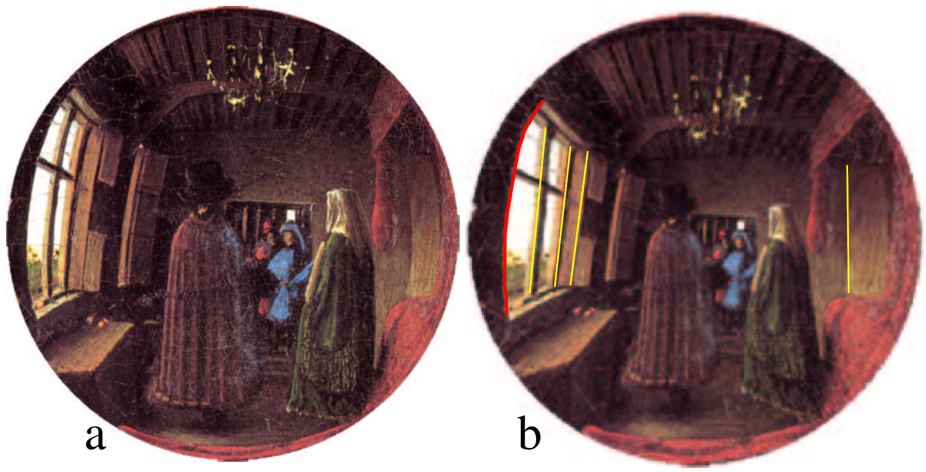

Figure 11: Accuracy of van Eyck's mirror: **(a)** Original image; **(b)** The best possible rectification with spherical mirror assumption. Most edges have been correctly straightened (marked in yellow), but the left-most edge of the window (marked in red) has remained noticeably curved. This is indication of inconsistency in the geometry of the painted mirror.

remains curved. Figure 11 illustrates this concept more clearly. While all the edges marked in yellow in the rectified image in fig. 11b are straight, the edge marked in red is still clearly curved. This can be interpreted as lack in geometric consistency; *i.e.* it is not possible for a physical parabolic or spherical mirror to produce the kind of image observed in the original painting (fig 2).

Fig. 12 shows the effect of rectifying the reflected image by the spherical model with different values of the sphere radius (the mirror radius r). As it can be seen, none of these values can straighten up 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 (*i.e.* zooming effect, see fig. 12).

Correcting inaccuracies in van Eyck's *Portrait*. Further inconsistencies characterise the geometry of Jan van Eyck's mirror.

In fig. 11b, 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 horizontal in the original painting (fig. 11a). However, this fact is physically impossible: these edges, due to 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 fig. 13. The steps taken were as follow:

1. Rectify the original mirror (fig. 13a) using the process described in the previous section and spherical assumption. The resulting rectified image (fig. 13b) 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. 13c) is perspectivally and physically "correct" in that projected straight 3D edges are now straight and the bench and gown edges (circled with a dotted line) are horizontal.

Figure 12: The effect of different radii: Rectifying van Eyck's mirror by means of spherical assumption and varying the radius of the sphere. Despite all the efforts there is no single value of the mirror radius that can straighten up all the edges at the same time.

3. Warp the image in fig. 13c using the inverse of the spherical transformation employed in the first step.

The resulting image in fig. 13d illustrates what the artist "should have painted". When compared with the original painting, the image in fig. 13d 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 fig. 13d, the bottom edge of the bench more realistically follows a sloped curve and so does the bottom of the woman's green gown. Furthermore, the curvature of the left-most edge of the large window is much less pronounced than in the original painting. Overall the image in fig. 13d looks more consistent.

The above 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 do so?

Plausible answers may be: i) the artist accentuated the curvature of the left-most edge of the window to convey a stronger sense of the bulge of the mirror (in fact, our analysis shows that Campin's mirror is bulgier than that of van Eyck); ii) he may have painted the edges of the bench and gown as straight to make them look less strange, *i.e.* accentuate reality.

In the interest of clarity, fig 14 shows enlarged versions of the images in fig. 13a and fig. 13d, respectively, with the edges of interest marked with different colours.

3. 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 Wedding*. 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 an historical point of view. The computer vision algorithms we have employed have allowed us to analyse the shapes of the mirrors, compare them and transform the reflected images into ones that correspond with orthodox perspective. We are thus provided

Figure 13: Correcting Jan van Eyck's *Portrait*: (a) The original painted mirror; **(b)** The rectified mirror, see text for details. **(c)** The left-most window edge is manually straightened. The bottom edges of the woden 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.

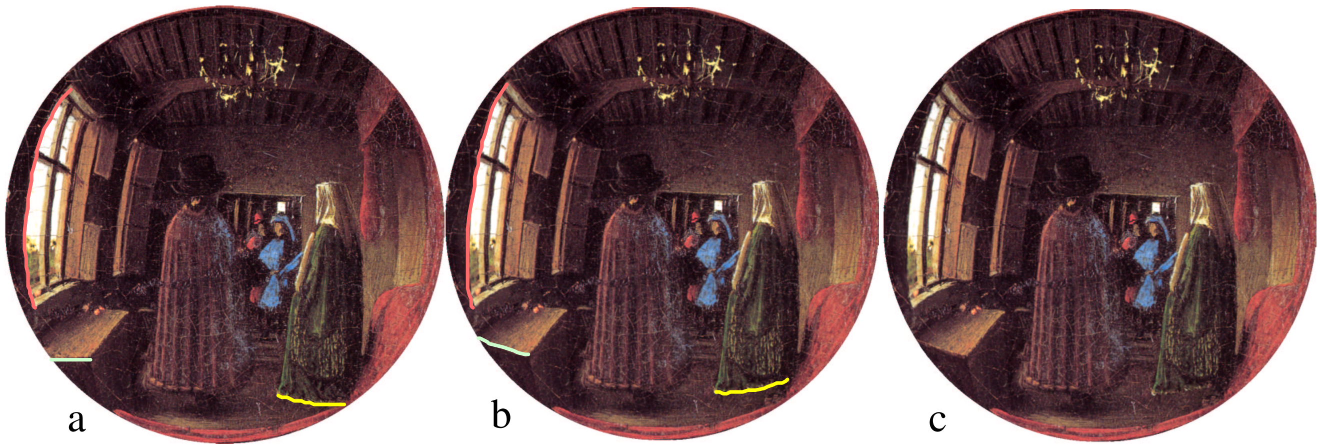

Figure 14: Correcting Jan van Eyck's *Portrait*: **(a)** The original painted mirror; **(b)** Our corrected version, see text for details. The left-most edge of the large window is less pronounced in the corrected version and the bench and gown bottom edges are correctly curved, consistently with the laws of Optics. **(c)** Enlarged version of fig. 13d, to show the "corrected" edges. To be compared to the original painting in fig. 13a.

with new views from the back of the depicted rooms. A rigorous comparison between the two analysed paintings leads to the following conclusions:

- 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's has purposely modified the geometry of the image in certain parts of the painted mirror.

A series of notable art-historical consequences flow from these findings, most particularly with respect to the *St. John* panel.

3.1. 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. 15a). The height and distance of the viewpoint 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[3]. However, the most prominent verticals in the rectified image (marked in yellow in fig. 15b) undergo upwards convergence towards a second, single vanishing point, which suggests 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

[3]The horizontal edges of the window (marked in green in fig. 15b) 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.

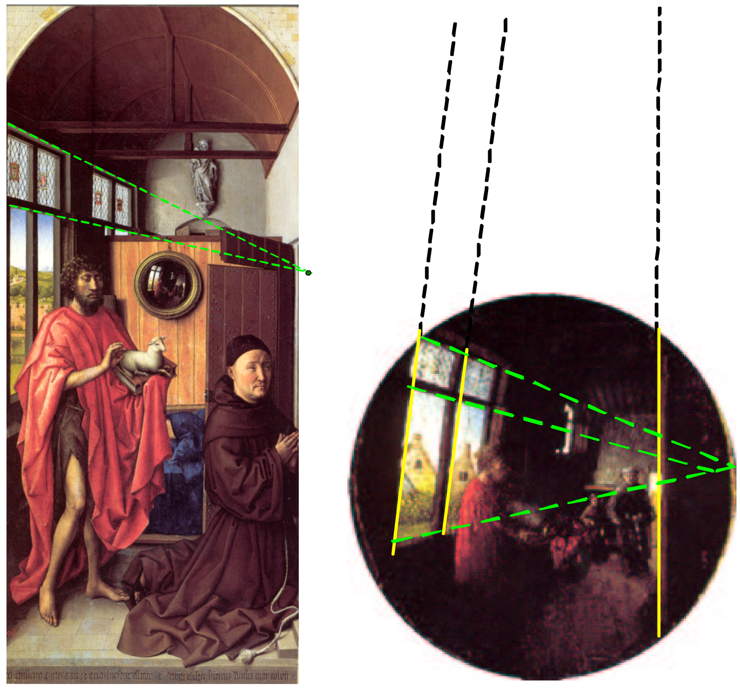

a **b**

Figure 15: 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 (in green) converge roughly to a point on the right boundary of the mirror, while the verticals (in yellow) converge upwards. The upwards convergence suggests a tilting of the "camera/panel".

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.

These 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 set-up as portrayed as a 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 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 mid-line, *i.e.* effectively tilting the picture plane or "camera" in such a way as to cause the upwards converge.

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. While it may be possible to envisage in a general way what an imaginary view would look like in a spherical mirror, it is inconceivable that such consistently accurate optical effects could have been achieved by simply thinking about it. We are drawn to what seems to be the inescapable conclusion that the artist has directly observed and recorded the effects visible when actual figures and objects are located in a specific interior. This means that at some point, models must have been posed in exactly the positions occupied in the painting and dressed in appropriate costumes - with the exception of the two Franciscan witnesses, who are in their normal habits. If this particular picture was achieved through such a stage-managed set-up in real spaces, 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 be decided. 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 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, this would facilitate the process. 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, he 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 and van Eyck seems entirely unjustified.

3.2. The Arnolfini Wedding

Judged on its own merits, the accuracy of the image in Jan'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 set-up 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 Jan 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 centre of the mirror, which suggests that he did not adopt the expedient we have suggested for Campin, *i.e.* moving closer to the mirror and slightly to the side in order to portray the reflected scene more readily.

3.3. 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 less 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 stop short of providing incontrovertible support. The main implications of our work lies 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 precisely the same time. Italian perspectival pictures, as pioneered by Masaccio, were achieved through the synthetic construction of geometrical spaces according to pre-conceived 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 as working as literally as they could with what they could see. They are true champions of what Gombrich has described as "making and matching". Such literalism of matching, much to the taste of 19th-century critics like Ruskin, has not found favour with recent art historians, who prefer more complicated explanations. It is almost as if we are precociously entering the realm of Courbet, the 19th-century realist, who declared, "show me an angel and I will paint one". Campin, five centuries earlier, seems to be saying, "to paint a St. John, I need to see one".

References

[1] R. Benosman and S. B. Kang, editors. *Panoramic Vision: Sensors, Theory, and Applications*. Springer-Verlag, 2001.

[2] J. F. Canny. A computational approach to edge detection. *Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.

[3] A. Criminisi. *Accurate Visual Metrology from Single and Multiple Uncalibrated Images*. Distinguished Dissertation Series. Springer-Verlag London Ltd., September 2001. ISBN: 1852334681.

[4] F. Devernay and O. D. Faugeras. Automatic calibration and removal of distortion from scenes of structured environments. In *Proc. of SPIE*, volume 2567, San Diego, CA, July 1995.

[5] O. D. Faugeras. *Three-Dimensional Computer Vision: a Geometric Viewpoint*. MIT Press, 1993.

[6] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.

[7] D. Hockney. *Secret Knowledge: Rediscovering the Lost Techniques of the Old Masters*. Viking Press, October 2001.

[8] M. Kemp. *The Science of Art*. Yale University Press, New Haven and London, 1989. ISBN: 0-300-05241-3.

[9] S. S. Lin and R. Bajcsy. True single view point cone mirror omni-directional catadioptric system. In *Proc. Int'l Conf. on Computer Vision*, volume II, pages 102–107, Vancouver, British Columbia, Canada, July 2001.

[10] S. Nayar. Catadioptric omnidirectional camera. In *CVPR*, pages 482–488, Puerto Rico, June 1997.

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

[12] P. Steadman. *Vermeer's Camera: Uncovering the Truth Behind the Masterpieces*. Oxford University Press, Oxford, 2001. ISBN: 0-19-215967-4.

[13] T. Svoboda, T. Pajdla, and V. Hlavac. Epipolar geometry for panoramic cameras. In *ECCV*, pages 218–231, Freiburg, Germany, June 1998.

[14] R. Swaminathan, M. D. Grossberg, and S. K. Nayar. Non-single viewpoint catadioptric cameras: Geometry and analysis. Technical report, Columbia University, 2001.