Chapter 5

Implementation and Applications

5.1 Implementation.

We have implemented the transformations described in Section 4.3.5 and Section 4.3.6. The implementation of our projections is quite straightforward. The Π_{plane} projection of Theorem 3.1 practically coincides with the standard perspective/parallel projection. There is, however, an important implementation detail that is absent in some rendering systems. As we mentioned before, our center of projection need not coincide with the position of the camera or the eye. It is chosen according to perceptual requirements. For instance, it can happen that the most appropriate center of projection for an office scene is outside the room (Figure 5.1c.) In these cases it is necessary to have a mechanism for making parts of the model invisible. (These parts of the model should participate in lighting calculations but should be ignored by the viewing transformation.) This can be done using clipping planes.

The 2D part of the viewing transformation, T_{plane}, can be implemented as a separate postprocessing stage. In this way, the transformation can be applied to any perspective image, computer-generated or photographic. The only additional information required is the position of the center of projection relative to the image. The basic structure of the implementation is very simple:

```
for all output pixels (i,j) do
  r := sqrt(i^2 + j^2)
  setpixel( i,j,
     interpolated_color( ρ^{-1}(r)i/r, ρ^{-1}(r)j/r))
end
```

For the geometric interpolation (4.21) the inverse function ρ^{-1} can be calculated directly using the formula

$$\rho^{-1}(r) = \tan{(2-\lambda)}\arctan\left(\tan\left(\frac{\arctan R}{2-\lambda}\right)\frac{r}{R}\right) \tag{5.1}$$

Equation 4.21 can be used in the cases of constant and variable λ.

For linear interpolation (4.20), the inverse function ρ^{-1} can be computed numerically with any of the standard root-finding methods, such as those found in *Numerical Recipes* [Press et al., 1988] As we have mentioned in Section 4.3.5, for moderately wide angles the difference between any transformation from the family that we are considering and the identity transformation is quite small. That means that the initial approximation of identity for any numerical root-finding procedure is good, and even the dichotomy method works quite well. There is no particular advantage in using the linear interpolation (at least in our implementation), but it was the first one that we have implemented, and some of our images were generated using this transformation.

The interpolated_color(x, y) function computes the color for any point (x, y) with real coordinates in the original image by interpolating the colors of the integer pixels. This interpolation can be done using convolution with any reasonable filter. Using a wider filter slows down the computation, but virtually guarantees the absence of aliasing.

Implementations of (4.23) and (4.22) are practically identical to the implementation described above. There is a significant difference in the case of (4.22): the values of ρ^{-1} cannot be precomputed, because they depend on ϕ. Using (4.23) is preferable, because the inverse function can be computed explicitly.

We also need a mechanism that allows user to specify desired values of λ for various directions in the picture. One way to implement this is to allow user to specify the values of λ at a number of points, compute the direction from the center of the picture to each point, and use a closed spline curve to interpolate the values of λ for all directions (see, for example, [Hill, 1990].)

The algorithm is linear in n and should be about as fast as any other image transformation (change of the size, filtering or other geometric transformations.) Our non-optimized program (in which no attempt was made to eliminate extra function calls in inner loops, etc.) took about 30 sec on an HP9000/700 to transform a 600 by 400 image with a filter of width 4.

5.2 Deep and Shallow Scenes: When Do We Need a Wide-Angle Projection?

As we have discussed in detail in Chapter 2, narrow-angle perspective projections are likely to produce perceptually acceptable images, while wide-angle projections are likely to produce distortions and are practically unusable for extremely wide angles (Figure 5.5a.) As we have shown in Chapter 3 and Chapter 4, it is in general impossible to find a projection which would produce undistorted images for any scene, given a fixed angle of projection.

Before we try to use compromises suggested in Chapter 4, we have to answer the following question: when is it possible for a given scene to use a narrow-angle projection instead of a wide-angle one?

A general direction from the scene to the projection point is determined by our choice of the side of the objects that we want to see in the picture (we assume that our preferences are consistent with the fiber structure of the perspective projection, i.e. we don't want to see around the corners or oppposite sides of one object at the same time.) We are only free to choose the distance between the projection point and the scene.

Suppose we want an object of size s is located at the distance l from the center of projection to have approximately the same size as an object of size S located at the distance $l + L$, $S > s$. As the ratio of the sizes of their images for a perspective projection is $Sl/(l + L)s$, the distance l, determining the projection angle, can be estimated from $Sl/(l + L)s \approx 1$. The projection angle α can be found from the expression $\tan\alpha \approx S/(l + L)$. The resulting estimate fo r l is $S - s/L$. Of course, S need not be a size of an object; it can be a distance between objects and in general any characteristic length in the direction perpendicular to the principal ray. We can consider S to be the "width" of the scene in the background, and s to be the "width" of the scene in the foreground. L is just the distance between the closest and most remote objects in the scene ("depth" of the scene).

Therefore, the projection angle has to be large if the difference between the foreground width of the scene and the background width of the scene is large compared to the depth of the scene.

We shall call a scene *shallow*, if the ratio $S - s/L$ is small; if this ratio is large (> 1), then we shall call the scene *deep*.

Deep scenes require wide-angle projections, while narrow-angle projections can be used for shallow scenes.

Almost any indoor scene is inherently shallow, because the objects in such scenes typically have comparable sizes. Therefore, in most cases for closed spaces of relatively small size we can get away with using the perspective projection. It is important to note that the choice of the projection point for a scene doesn't have to reflect an actual physical position of the observer in the scene. For example, any real-life photo showing a whole wall of an office would require a very wide projection angle. However, in computer graphics, nothing prevents us from choosing a projection point in an unreachable position in or outside the room. Our choice can be based purely on the perceptual quality of the picture. Minimization of shape distortion requires narrow angles of projection. At the same time, extremely small angles of projection result in a lack of parallel foreshortening and excessive perpendicular foreshortening for some objects. Therefore, a moderate projection angle is most preferable. In Figure 5.1 we can see pictures of the same part of the office from three projection points, corresponding to the projection angles 90°, 36°and 3°. Most people prefer the projection in the middle. In fact, the number 36° was determined by showing to people a set of pictures with the projection angle varying from 3 to 90°, and asking them to choose the nicest picture. The variance of the results was small: nobody chose any image outside the limits 32-40°.

The effect of the change of the projection angle on the image of a deep scene is radically different. For a deep scene the general layout and relative sizes of the objects are unchanged only for a small range of angles. Consider Figure 5.2. The 90° image has the desired composition, but distortions are clearly visible: the men close to the edges of the picture appear to be "bulkier" than the men in the center of the picture, although their shape was identical. The shape of the head appears to be distorted in some cases, and the men close to the edge of the picture appear to be asymmetrical.

The two 60° images are radically different: in the first case, when all the objects are still present in the picture, the men in the picture almost disappear. In the second case, the pyramids in the background are absent, and relative sizes of the images of the men are changed.

a

b

c

Figure 5.1: Example of a shallow scene: a. 90° projection angle; b. 3° projection angle; c. 36° projection angle. Position of the camera is shown on the right for each picture. Note that for 36° image the camera is outside the room. □

For deep scenes, the projection angle is practically determined by the scene and composition of the picture, and cannot be varied. In this case transformations of the type described in Section 5.1 can be applied to decrease the distortions in the perspective image. The results of applying the transformation (4.20) to the 90° picture is shown in Figure 5.2b. Note the reduction in the distortion of shape of the men. As a tradeoff, the barrel distortion is introduced, but due to the absence of straight lines, its only manifestation is a slant of some figures.

Striking distortions occur in wide-angle motion pictures, especially when the camera rotates around a fixed vertical axis (pan.) They are primarily associated with two factors:

– when an object moves to from the edge to the center of the picture, the shape of its image changes from distorted to non-distorted. This change is noticeable and produces an impression of non-rigidity.

– the vertical coordinate of objects that are located at a large angle from the horizon changes considerably (cf. Figure 2.22). This effect appears unnatural.

These effects are reduced if transformations like (4.20) or (4.21) are applied to the frames.

Summarizing the discussion above, we suggest the following recipe for the choice of an appropriate projection for an image:

1. Determine if the scene is shallow or deep.

2. If the scene is shallow, choose a perceptually optimal projection angle which is likely to be in the range 15-45°.

3. If the scene is deep, determine the minimal projection angle that is acceptable for the scene. Then choose a transformation from (4.21) or (4.23) which produces the best results.

5.3 Applications of Correcting Transformations

In the previous section we described the conditions which require a wide-angle projection under the assumption that we are restricted in our choice of projection only by the composition of the scene. In the case of photography there are several additional restrictions: there should be no obstacles between the camera and the objects that must appear in the image, and the size of the camera cannot be made arbitrarily large. For example, it is impossible to make a picture of a room from

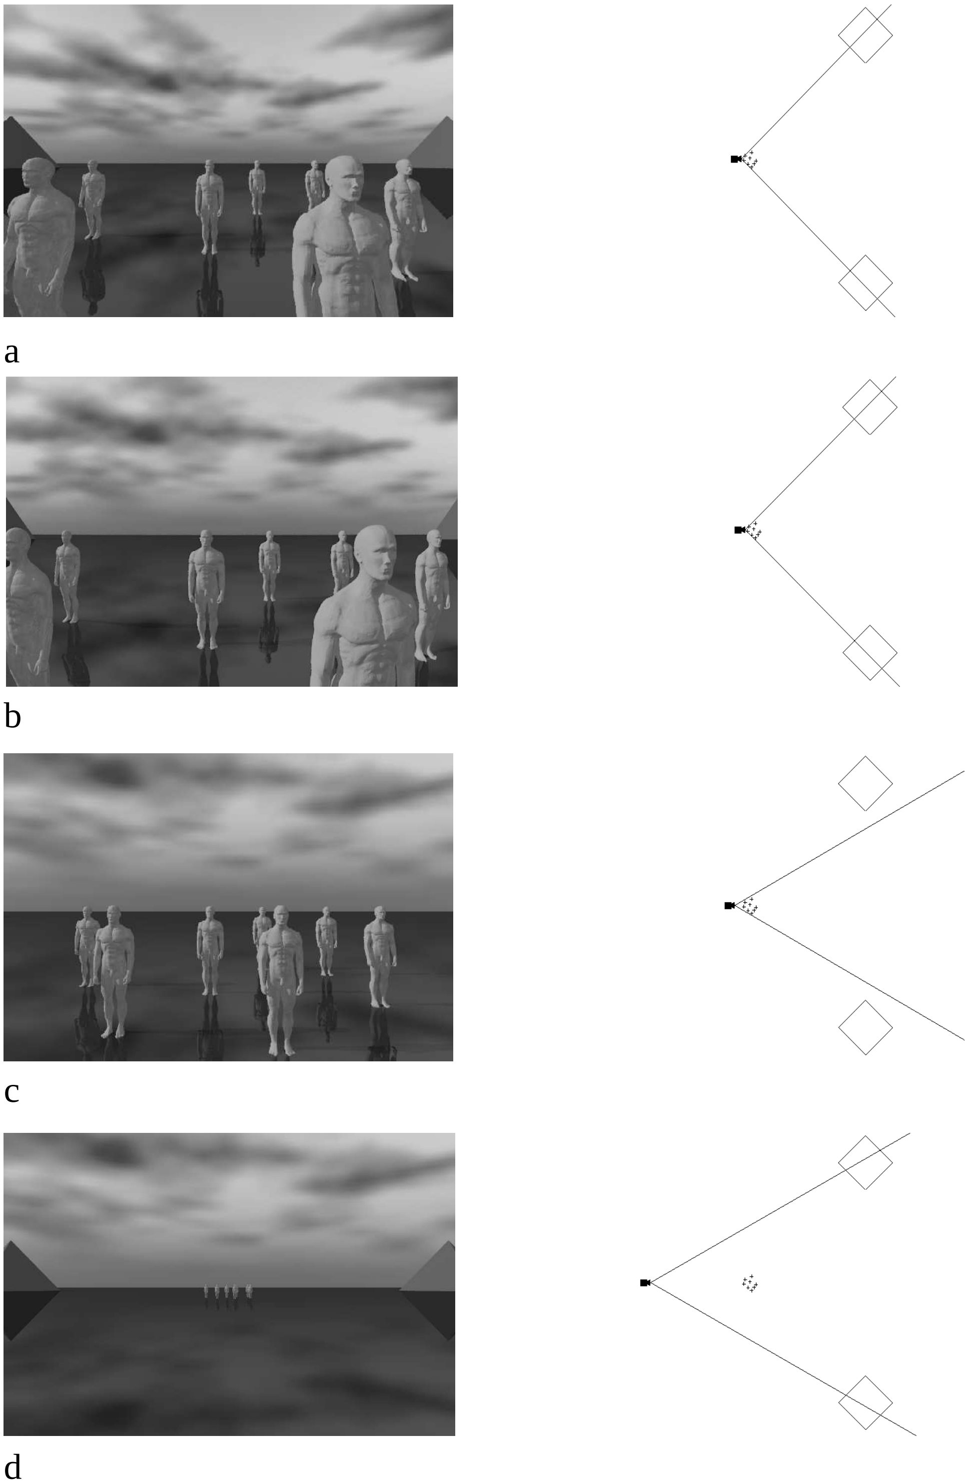

Figure 5.2: Example of a deep scene: a. 90° projection angle; b. 90° projection angle, correcting transformation with $\lambda = 0$ applied. c. 60° projection angle, keeping the figures of men approximately the right size; d. 60° projection angle, keeping the pyramids in the field of view; □

a projection point behind a wall. As a result, the projection angle can be chosen only in a limited range, and some pictures must use a wide-angle projection.

In many cases the best camera positions are physically inaccessible or there is no time to make a good choice of the position.

All these reasons explain why wide-angle images are inevitable in many cases even for shallow scenes. Transformations 4.21, 4.23, 4.21, 4.23 can be used to correct for distortions of shape in such images.

The position of the center of projection is usually known for computer-generated images, but is more difficult to obtain for photographs. For photographs it can be calculated if we know the size of the film and the focal distance of the lens used in the camera. Alternatively, it can be computed directly from the image if there is a rectangular object of known aspect ratio present in the picture (Appendix B.)

If our task is correcting a distorted photographic image, it is important to understand what type of distortion is present.

For example, distortions related to verticality (Section 2.4.3) can be corrected by a linear transformation, as it was known to photographers for a long time ([Feininger, 1953], Figure 5.3.)

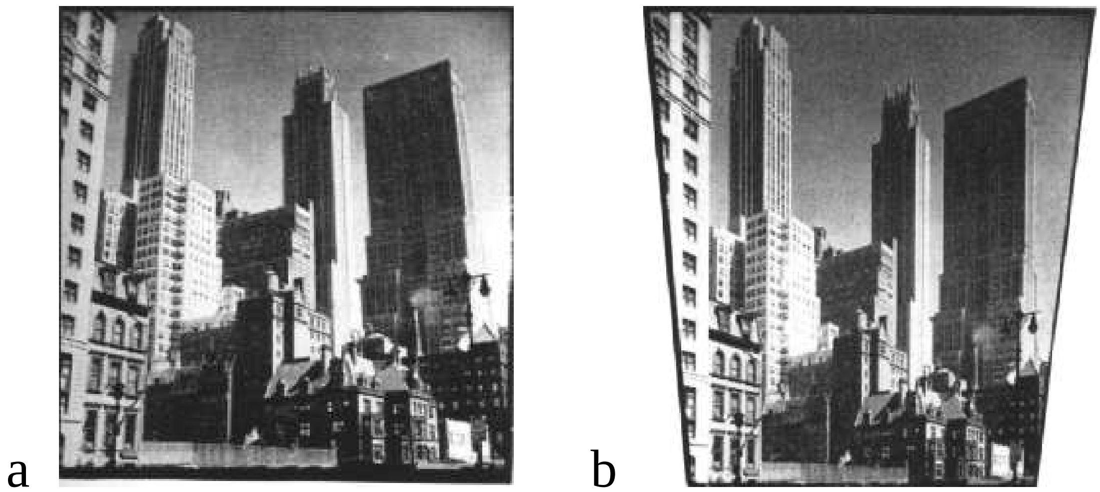

Figure 5.3: Photos from [Feininger, 1953]: a. The original photo with convergent verticals; b. Same photo printed with slanted easel, the slant is chosen so that the verticals become parallel. □

Distortion of shape can be corrected using our transformations. In some cases these transformations also decrease distortions of relative size due to the excessive perpendicular or parallel foreshortening.

Let's consider in greater detail an example of application of transformations 4.21 and 4.23.

The head of the author in Figure 5.4a is considerably distorted, and his right shoulder appears to be longer than the left. He also appears to have much more muscle than in real life. The shape of small objects on the table is distorted. The rectangular corner indicated in the left part of Figure 5.4a also looks distorted (cf. Section 2.3.1.)

After application of the transformation (4.21) for $\lambda = 0$, the author looks much more like himself and the rectangular corner looks better. The straight edges in the left part of the picture are curved.

If we use the transformation (4.23) with λ gradually increasing from 0 on the left to 1 on the right, we can eliminate the curvature of edges, while leaving the corrections in the left part o the picture intact. Unfortunately, the rectangular corner doesn't look right again - we cannot have both zero-curvature and correct shape at the same place.

Again, we give a short recipe for correction of wide-angle photos:

1. If vertical objects in the picture appear to be falling, apply a linear transformation which makes convergent verticals parallel.

2. If distortions of shape are present, apply (4.21) or (4.23). Transformation (4.21) is sufficient when all long straight lines in the picture pass close to the center of the picture; otherwise, choose $\lambda(\phi)$ close to 1 in 4.21 for the directions where there are straight lines and no objects of simple or familiar shape, and close to 0 in the directions where there are such objects and no long straight edges.

5.4 Other Applications

There are two more applications of our transformations that might be useful.

Zooming. If it is necessary to extract a part of a wide-angle image, the objects in the sub-image are likely to look quite distorted. In the context of the whole image this distortion appears to be more acceptable, but when a part of the picture is clipped, the distortions become more objectionable. If the clipped part is small, no long straight edges will appear in the picture, and transformation (4.21) with λ close to 0 can be applied before clipping. The results are demonstrated in Figure 5.4d and e.

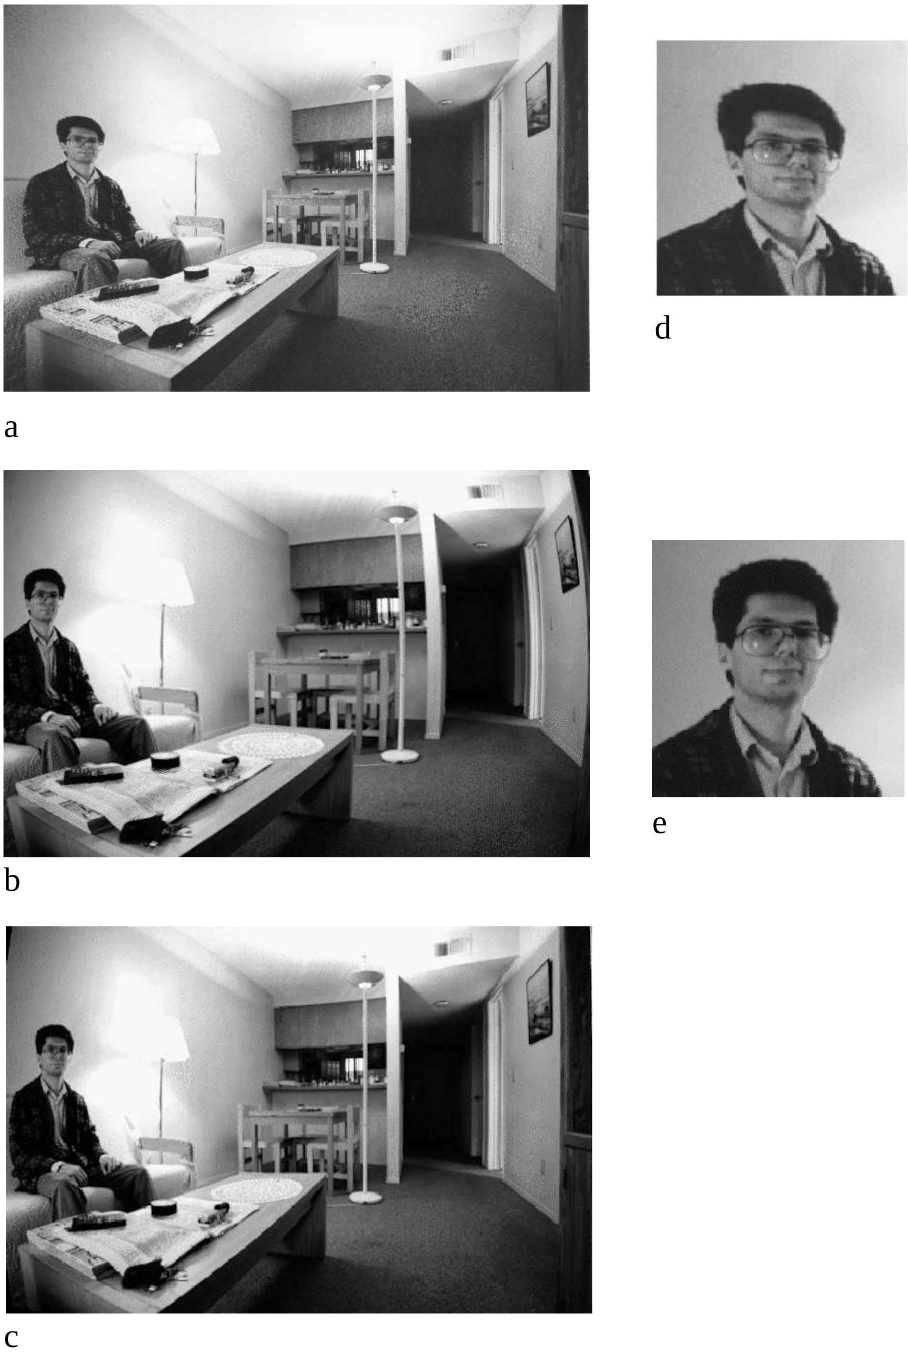

Figure 5.4: An example of application of (4.21), (4.23) to a wide-angle photo: a. Original picture, 180°; note the distortion of the shape of the head and of the marked rectangular corner in the right part of the picture. b. Transformation 4.21 with $\lambda = 0$ applied; the straight edges in the right part of the picture are curved. c. Transformation 4.23 applied; λ changes from 0 to 1. d. Zoom of the head in the original photo. e. Zoom of the head in the transformed photo (b). □

Extremely wide angle images. As we have seen, (Figure 5.5a), linear perspective is unsuitable for extremely wide angles. Rather than using traditional fisheye views, which "overcompensate" for the distortions of the perspective projection, we can use one of our transformations. Visible distortions are practically inevitable, but for our transformations they are likely to be less than for the perspective projection and fisheye views.

5.5 Summary

Here is a summary of possible applications of the transformations that we propose:

- Creation of wide-angle computer-generated pictures with minimal distortions.
- Reduction of distortions in photographic images.
- Creation of wide-angle motion pictures with reduced distortion.
- A better alternative to fisheye views.
- Zooming of parts of a wide-angle picture.

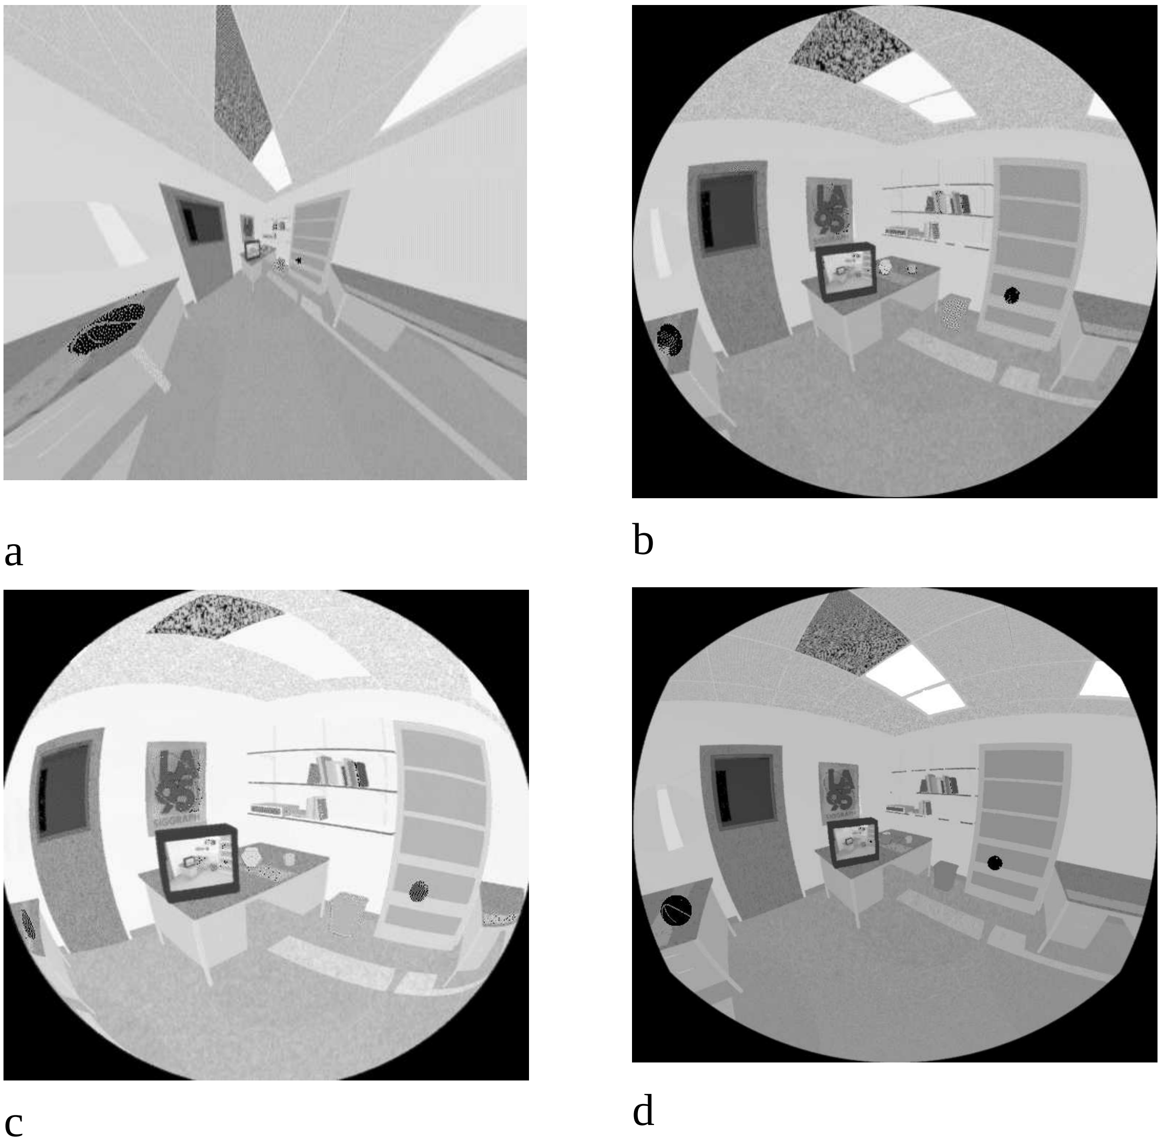

Figure 5.5: Extremely wide angle images (160°): a. Linear perspective. b. Angular fisheye (equidistant projection). c. Hemispherical fisheye (orthographic projection). d. Our direct projection. Images (b) and (d) are quite similar, but in (c) we can observe excessive compression of the objects in the margins –the tables appear to be narrower and the ball doesn't look spherical. □