

The Scheimpflug Principle—Part III

by Harold M. Merklinger
as published in *Shutterbug*, February 1993.

In Part II we discussed how the focal length of flat-field lenses is not fixed, but rather may vary from ray to ray depending upon the angle the ray makes with the lens axis. As a result we can effectively zoom a lens to some extent simply by tilting it. But as we saw in Part I, tilting the lens also tilts the plane of sharp focus.

In part III we will be addressing some aspects of perspective. Let me say at the outset that I am not an expert on perspective. The rules of perspective are, in my humble view, somewhat arbitrary; they are man-made standards rather than natural physical laws. The accepted conventions for perspective predate photography and even the camera obscura. Furthermore the rules have changed with time. My purpose in discussing perspective here is not primarily to establish what is correct perspective—although I will express opinions. Rather I wish primarily to explain how the relative positions of lens and film influence geometrical distortion and hence apparent perspective.

Tilting the camera lens has no inherent impact upon the apparent perspective of the image. We saw this in Figure 5 of Part II: I tilted the lens by a fairly strong 30°, with no more result upon the image than a 15% enlargement and some effect upon what was in focus and what was not. If one were to stop a lens down to the size of a pin-hole, one would have, in effect, a pin-hole camera. All objects would be roughly equally sharp, and tilting the lens would have essentially no impact upon the image at all. We would not even need to refocus.

The point I am making is this: tilting or swinging the lens of a view

camera has no effect whatever upon the apparent perspective of the image. Swinging and tilting a lens is a focusing operation. For the purpose of analyzing perspective, we can indeed best think of the lens simply as a pin-hole. A corollary to this is that when adjusting a view camera, we can consider the action of focusing and the action of controlling perspective quite independently. These two actions do not affect one another. These actions only become coupled if some other constraint—like the covering power of the lens—limits the degree of adjustment one should properly make.

That being said, the next logical question is: Well, what does control perspective? And, what is it anyway. Perspective in its most general terms means giving visual clues to the viewer of a two-dimensional image, so that he or she can interpret the image in three dimensions. These clues include such nuances as relative size, sharpness, and contrast. Contrasty objects are generally perceived to be closer than objects of low contrast, for example. Fuzzy objects are assumed to be closer than or more distant than a main, clearly delineated subject. The type of perspective I will address here, however, is restricted to geometrical shapes and sizes.

In the context of photography with rectilinear lenses, there are two effects of interest: image magnification and image distortion. One might even argue that there is only one—magnification—but the degree of magnification may be different in different directions. I find it easier to think in terms of relative magnification—how does image size change if I put the object here rather than there—and distortion—do the pro-

portions (relative height and width) of the object change if it is moved from here to there. Both effects depend mostly upon one factor: how far off the bore sight is the image? (You will recall, I hope, that we defined the bore sight of a camera as a line perpendicular to the film plane, passing through the optical center of the lens.)

When a small light ray bundle of circular cross-section strikes the film at an angle, the image on the film will not be circular, but rather elliptical: a circle that has been stretched in one direction. In one direction, perpendicular to a line from image to bore sight, the image has the 'right' dimension. In the other direction, parallel to a line from image to bore sight, the image is too long. This is what I mean by distortion. How much too long depends upon how obliquely the ray bundle strikes the film. On the film plane near the bore sight, objects are imaged with negligible stretch. Far from the bore sight, the stretch can be significant. Figure 1 illustrates. What really matters here is the angle measured at the lens between the bore sight and the rays forming the particular image of interest. Within 25° to 30° of the bore sight, the stretch is almost unnoticeable. Beyond that it becomes increasingly noticeable; the stretch will reach a factor of two for rays 60° off the bore sight.

The implication of this phenomenon for perspective is that the flattening of a circle tends to make us think we are looking at the circle from an oblique angle. If we are looking at a sphere which gets flattened, however, we are confused: the image does not look real.

Figure 2 shows some table-tennis




Figure 1: The image of a spherical object (the moon, say) will be imaged near the bore sight as a circle. At a large angles off the bore sight (more than 30°) the same object at the same distance will be imaged as a somewhat larger ellipse. The image is magnified because the effective focal length of the lens is greater at this angle, and it is elongated because the imaging rays strike the film obliquely.




Figure 2: Here's a photograph of six table-tennis balls and five round white dots taken with a 15 mm lens on a 35 mm camera. The balls in the corners show significant stretch while the ball in the very center is quite round. The dots, on the other hand are round no matter where they are—so long as the film plane is parallel to the object plane—as was the case here.

balls and some flat white dots mounted on a sheet of black cardboard. They are photographed from a distance of one foot with the widest-angle rectilinear lens I have: a Pentax 15 mm lens for 35 mm cameras. In this photograph the dots all look fine. The balls look fine near the center of the image, but rather strange near the corners.

Those with sharp eyes may notice a certain discrepancy between Figures 1 and 2. In Figure 1, the ellipse's smallest dimension is larger than the circle. In Figure 2, the 'width' of the ellipsoidal table-tennis balls is just about the same as diameter of the round table-tennis balls. Why the difference? The answer is that I drew Figure 1 to represent the images of two spherical object each the same distance from the camera lens. Off the bore sight the magnification is greater (as well as the distortion) and so the image is larger. For Figure 2, the balls at the edges and in the corners are farther from the lens than is the ball at the center of the picture. One might expect therefore that their images would be smaller, but the magnification effect exactly cancels the size reduction due to subject distance.

Image magnification depends simply upon lens-to-image distance divided by the lens-to-subject distance. (For images of objects at infinity, the lens-to-film slant distance alone affects relative magnification.) The shortest possible lens-to-image distance—and hence minimum magnification—will occur along the bore sight. 60° off the bore sight, image magnification will be twice what it is at the bore sight, assuming objects are uniformly distant from the lens.

These elongation and magnification effects compound one another: the farther off the bore sight (in angular terms) an object lies, the greater is its magnification, and, the more stretch it undergoes. Even worse, for large objects subtending a significant angle at the lens, the magnification and stretch can be significantly greater at one end of the object than at the other. Figure 3 shows how a large image of a spherical object can become egg-shaped. What has happened here is that a) since the image is off the bore sight, the sphere is imaged as an ellipsoid and b) since the ellipsoid covers a wide range of magnifications, the end of it farthest from the bore sight is magnified significantly more than the end near the bore sight

is.

Users of ordinary cameras equipped with wide angle lenses will be familiar with these "wide-angle distortion" effects. A beach ball looks normal in the center of a picture, but more like a football near the corners of the image. And if the beach ball is close enough to the camera, so that it's image is relatively large, the changing magnification from one end of the 'football' to the other will make it look more egg-shaped than football-shaped.

The 'distortion control' offered by view cameras is really nothing more than manipulation of the 'wide-angle-distortion' described above. What makes the view-camera different from the normal SLR, say, is that in the case of the view camera we see only a portion of the total picture. If the photographer selects that portion near the bore sight, we see a relatively normal image. But he can equally well select the very outside of the coverage circle, where stretch is at a maximum, and where magnification is varying most rapidly. And so he can turn beach balls in the very center of the final image into eggs—or vice versa.




Figure 3: If the image of a spherical object lies off the bore sight and also extends over a large range of off-bore-sight angles, the image will be egg-shaped. The portion of the image farthest from the lens is magnified significantly more than are other parts of the image nearer the lens.




Figure 4: The view camera can produce "wide-angle distortion" even when the lens being used covers only a small angle. The method is simply to swing the camera back. In this case there is no image at all where the bore sight meets the film plane.

Although lenses made for view cameras usually have lots of covering power, even lenses of modest covering power can be made to yield excessive distortion. All the view camera photographer needs to do is to swing the camera back. Figure 4 shows how. In this case there is no image at all where the bore sight meets the film plane. But the film is illuminated at a steep angle, and hence the distortion effects are significant. In this case the photographer might have a focusing problem, but then again, for the right subject, this geometry might be exactly what he needs. I note in passing that there is an exposure factor to consider also when the imaging rays fall obliquely on the film. The image intensity is reduced because the same bundle of light is spread over a larger area of film. If the light rays fall at a 45° angle to the film, one must open the lens by one stop (or double the exposure time) to compensate. At 60° off the bore sight, the compensation required is two stops.

Figure 5 shows exactly the situation just described. Figure 5a shows a

quite normal photo of a globe. For this photograph I used only back rise and front fall to obtain a downward viewing angle. The globe is very slightly taller than it is wide as a result, but one can hardly notice it. Figure 5b shows what happens when I swing the back by 40° . It was also necessary in this case to swing the front by 20° to maintain reasonable focus. I would have used greater back swing to exaggerate the effect even more, but the camera bellows would not allow that without cut-off of the image. Figure 5b was also taken from slightly farther away than Figure 5a, just to let the globe fit comfortably in the frame.

When reproducing images of flat subjects such as drawings or paintings, it is important that the film plane and the object planes are parallel. This will ensure that any distortion of the image by the camera will exactly compensate the true perspective distortion that truly exists. If we were to look objectively at the object from the position of the lens, we would see that the extreme corners looked small and foreshortened—that

is, squashed in one direction. But the intentional distortion introduced by a rectilinear lens exactly compensates and ensures that the final image looks 'right'. This is why the flat white dots in Fig 2 still look circular.

Perhaps the most common examples of a distorted but 'correct' images are pictures of buildings. By convention, tall buildings are represented in art and photography with vertical lines that do not converge, and in some cases even with horizontal lines that do not converge. Preventing converging vertical lines is simple: keep the bore sight horizontal. This means keeping the film plane vertical. Keeping horizontal lines from converging is also simple: keep the bore sight perpendicular to the horizontal lines in the image. If we are photographing a building in such a way that we can see only one face of the building, we usually want to keep both horizontal and vertical lines from converging. This is done by holding the bore sight perpendicular to the visible face of the building. In other words, we keep the film plane parallel to the face of the




Figure 5: a) show a relatively normal photograph of a globe. b) shows the same object from essentially the same camera position, but with the camera back swung by 40° . The world no longer appears spherical, and the books seem to be lying on a slope.

building. We can position the lens anywhere, but keep the film parallel to the building.

If, in another example, one can see two faces of a building (the front and one side, for example), the convention states that we maintain only the vertical lines parallel; we let the horizontals converge. We can still play with the orientation of the bore sight in the horizontal plane. The most natural image usually results when the bore sight is pointing at about the center of the building. There is some freedom, however, to exaggerate the apparent perspective. In general, the building will tend to look as though it is being viewed in a direction parallel to the bore sight—even when that is not the case.

Figure 6 illustrates two images of a building from the same camera position. Figure 6 a) shows a building taken using front rise only. (Well, actually, I also tilted the lens down a bit to sharpen the foreground.) Vertical lines are vertical and horizontal lines converge: for the left face they converge to the left and for the ‘front’ face they converge to the right. For Figure 6 b) I moved the back by about 38° (counterclockwise if one is looking down on the camera). The lens also needed to be swung to re-establish good focus. This somewhat

extreme back swing causes the horizontal lines for the ‘front’ of the building to converge to the left instead of to the right. The result tends to make the photograph look as though it was taken from a position further to the photographer’s right. This photograph is not ‘correct’: the viewer is confused by mixed visual clues.

I have seen examples in books where it is stated that the proper view camera solution to photographing a wall obliquely will involve both lens swing—to maintain focus—and back swing—to restore the proper perspective. I do not believe this is correct. I suggest that what has happened is that when the photographer adjusted the lens for proper focus, he discovered that the horizontal lines no longer converged as much as he expected. The true cause was an effective increase in focal length, but the photographer interpreted the result as a change in apparent perspective. To achieve greater convergence of the horizontal lines in the image, he adjusted the back. This is not a serious error; to a great extent, it is all a matter of art and taste.

So, the lesson for this month is that, for the purposes of perspective, lenses act like a simple pin hole. All that matters is where the image is in

relation to the bore sight—a line perpendicular to the film plane and passing through the pin hole. Near the bore sight, images are reproduced more or less naturally. For images formed at large angles from the bore sight, the image is stretched in one direction and magnified. The view camera is able to reproduce this ‘wide angle distortion’ even with lenses which are not wide-angle. As a vague guide, images often tend to look as though the camera’s line-of-sight was along the bore sight of the camera—even when such is not actually the case. This has not been a complete guide to perspective; the major points are that lens orientation (swing and/or tilt) does not affect perspective, rather, it is the angle of incidence of the image rays on the film that determines the effect.

In part IV we’ll look at another rule similar to the Scheimpflug rule. This second rule is perhaps even more valuable than the Scheimpflug rule, for it is capable of telling us precisely how much we need to swing or tilt the lens. But both rules are needed for a complete solution to the optical problem.

© Harold M. Merklinger, 1992.



Figure 6: a) shows a ‘correct’ image of a building. Front rise only was used for this photo. In b) the back was swung by almost 40° to give a rather odd and ‘incorrect’ image. This second photo looks in some ways as though it was taken from a camera position to the photographer’s right. But other clues indicate this is not the case: we can see the left side of the building, for example, and we know that most buildings have 90° corners. Swinging the back also stretches the image horizontally.

Author’s Note: Since this article was published, it has been pointed out that some of the terminology used here is not that normally used in describing drawing perspective. In particular, what I called the “bore sight” is traditionally called “the line of vision”.