

The Language of Three Dimensions



GETTING
STARTED
WITH 3D

6

LIVING IN A 3D WORLD

In the real world we are surrounded by solid objects that are three-dimensional objects (rocks, trees, mountains, bears, people), which we must be able to recognize and locate in space. Our senses equip us to interpret and interact with the three-dimensional world we inhabit. With our sense of touch we experience the solidity and texture of everyday objects. Our hearing helps us to immediately recognize and locate the source of threatening or interesting sounds. And our highly developed visual system (our eyes working together with our brain) helps us to identify objects and determine whether they are close to us or far away.

Our senses are vital to our survival in the world, but they also provide us with profound aesthetic experiences. "Beauty before me, beauty behind me, beauty all around me" goes a Native American poem, expressing the especially great pleasure we take in the visual perception of nature. This pleasure was one factor that led human beings to begin making graphic art: both sculptural and two-dimensional. And though hunting rituals may have inspired our ancestors' cave paintings of

animals, those elegant and realistic images also show a keen appreciation of the beauty of natural forms. In the thousands of years since these early artists first depicted three-dimensional objects on a flat surface, people have developed and honed their "drawing" skills in order to imitate in pictures the way things look in real life.

This striving to represent the real world (as well as to create new worlds) is a driving force behind the systems of perspective and modeling that were developed during the Renaissance and lead directly to the creation of 3D models and images using computers. Today's computer programs for 3D modeling and rendering are simply the latest tools in a long line of mechanical aids to drawing. All of these methods, old and new, have as their aim the depiction of solid objects on a flat surface in a way that mimics the eye's view of the real world.


THE HUMAN VISUAL SYSTEM

In order to create two-dimensional images that look similar to real scenes, artists take advantage of the way our visual system works, imitating with paint the various depth cues we use to interpret what we see around us.



Two Eyes = Two Images = Depth

Because we have two eyes spaced slightly apart and facing forward, our brain receives two slightly different views and merges them to create a perception of depth. The sense of depth provided by our binocular vision cannot be duplicated on a two-dimensional surface, such as a painting or a computer screen, but various monocular (one-eye) depth cues can be imitated. Techniques for rendering depth were developed during the Renaissance by many architects and painters, including Leonardo da Vinci. The eyes shown above are a detail from his painting *Lady with the Ermine*, painted around 1485. It is in the collection of the Czartoryski Museum, Cracow.



Stereoscopic images

Sir Charles Wheatstone (1802–1875) was the first to develop a process for producing two offset views of a scene, to simulate the differing views seen by each eye. When presented to each eye separately in a device called a *stereoscope*, these two views create a perception of depth. Wheatstone's device was introduced in 1838, shortly before the invention of photography. During the 1850s several inventors created stereoscopic cameras, with lenses placed about 2.5 inches apart to imitate the spacing of the eyes. Stereoscopic photography became very popular around the turn of the century. The image at left, entitled "I Just Knew It Would Be a Boy," was produced by T. W. Ingersoll in 1898. Although the pictures appear identical at first glance, notice how the objects in the right-hand image are slightly to the right in the frame, compared with the objects in the left-hand picture.

To view the images without a stereoscope is tricky, but try placing a piece of paper upright between them and touch your nose to the paper's edge so that each eye sees only one image. Your brain may (or may not) merge the two flat images into one image with depth.

Learning to draw is really a matter of learning to see...

—Kimon Nicolaïdes, *The Natural Way to Draw*

The same techniques used by traditional artists are also incorporated in the programming of desktop 3D software, in order to create computer graphics that look strikingly similar to real objects.


BINOCULAR VISION

How do we perceive depth? The most important aspect of our visual system is the fact that we have two eyes rather than one (*binocular vision*). And in human beings, along with other predators, these eyes are located a small distance apart on the front of our heads (as opposed to prey animals, which have their eyes located on the sides of their heads). This provides us with a large binocular visual field, in which each eye gets a slightly different view (*retinal disparity*). These two views are combined by the brain in a way that gives us the perception of depth. The disparity is sometimes called *parallax* and the combining of the two images is called *stereopsis* (“solid sight”), or *stereo vision*.

Stereo vision helps predators find and track their prey, especially over long distances. But it doesn't do much to help artists depict solid-looking shapes, since two-dimensional art is, by definition, *flat*. So to create a sense of depth, artists must use their paint or pixels to imitate the depth cues that can be perceived with only one eye.

Creating three dimensions with two

Leonardo da Vinci (1452–1519) painted this *Annunciation* when he was only 21 years old. The painting employs all the monocular depth cues we will describe in this chapter, including occlusion, size differences, linear perspective, shading, texture gradients and aerial perspective. Can you find examples of each of these in the painting? For an analysis of the perspective in this painting, see page 10.



THE WORLD
OF THREE
DIMENSIONS

7



Art imitates life

This photograph of coastal palm trees receding into the distance includes all of the monocular depth cues used in the painting above. In the 3D rendering next to it, created in Strata StudioPro, we've tried to imitate these depth cues in an invented landscape based on the photo. Images in any medium that include all six monocular depth cues look the most convincing in terms of the imitation of reality.

We will continue to use examples from painting, photography and 3D imaging to illustrate each of the monocular depth cues on the following pages. (Photo by Janet Ashford)



MONOCULAR DEPTH CUES

In addition to the depth cues provided by binocular vision, there are many cues that function with only one eye, or viewpoint. These are called *monocular depth cues* and include the following:

OCLUSION

Objects that are close to us partly block our view of objects that are “behind” them or further away.

SIZE DIFFERENCES

An object that is closer to us looks larger than a similar-sized object that is further away.

LINEAR PERSPECTIVE

Lines that are parallel with our line of sight (such as railroad rails when we’re standing in the middle of the track) appear to converge together in the distance.

TEXTURE GRADIENTS

Uniform textures (such as a checkered table cloth) look denser when they are farther away.

SHADING

Light shining on an object wraps around it in a way that creates gradations of light and dark. Parts of an object that are further away from the light look darker.

ATMOSPHERIC PERSPECTIVE

Objects that are further away look fuzzier, less detailed and often bluer than near objects because particles of dust in the atmosphere cause light to be scattered and also because different temperatures of air masses between object and viewer cause refraction of light.

While both binocular and monocular depth cues help us avoid bumping into tables, only monocular depth cues can help us interpret an image on a flat surface—such as a painting, a photograph or a 3D rendering. Over the years, artists have learned how to imitate these monocular depth cues in order to create images that look real.

Keeping in mind that the final product of a 3D graphics model is a flat image, programmers have designed their software to render 3D models in a way that imitates the conventions already used by painters. These techniques include occlusion through the creation of objects with opaque surfaces; size differences based on relative positions in 3D “space”; the use of linear perspective to show the convergence of parallel lines on a vanishing point; the use of light and dark areas to show modeling of forms; and the use of *aerial perspective* or atmospheric effects to create a sense of distance. These traditional techniques all have analogs within computer 3D programs.

OCLUSION

This depth cue is one of the most apparent and easiest to understand. Objects which are close to us partially block our view of objects that are further away or “behind.” In addition, when we view a single opaque object, the surfaces that are facing us hide the surfaces that are in back. For example, when looking at a cube, we can see only three of its six surfaces at a time. This effect of occlusion is readily apparent in a photograph or in a painting done by an artist with conventional training. However, occlusion presented a problem in early computer 3D programming.

HIDDEN LINE REMOVAL


Early 3D graphics programs were able to draw only the contours or edges of objects. This was done by a plotting of points very similar to the techniques used by Renaissance artists (see “The Science of Drawing” on page 13). A line image like this is called a wireframe and is the simplest way to display a 3D object. However, a wireframe drawing is often visually confusing because the computer traces and presents to the viewer every edge,



Blocking the view

Solid objects that are close to us partially block our view of objects that are behind them, providing immediate information about their relative positions in space. This is called *occlusion*. In the child’s painting above, occlusion (the tree blocks our view of the house) is the only depth cue that has been provided. The painting is expressive and decorative, but without further depth cues it looks quite flat. In the photo and 3D image below, the trunks of the close palm trees block our view of the trees behind them. These two images also include size differences and some shading, cues which enhance our perception of depth. But with only three monocular depth cues provided in these detail views, the images look almost as flat as the child’s drawing.





Hiding the "hidden lines"

This demonstration of "hidden line removal" was created by Melvin L. Prueitt for his book, *Computer Graphics*, published by Dover Publications in 1975. The 3D objects are visually confusing when all the edges calculated by the computer are visible (top). But when the "occult" or hidden lines are removed (bottom), the objects are easier to interpret and take on the appearance of depth. Algorithms that remove hidden lines and surfaces make it possible to create opaque-looking objects on the computer and duplicate the monocular depth cue of occlusion.

including those that should be hidden behind surfaces. There is no illusion of depth to help the viewer interpret the image. So beginning around 1963, algorithms that enable "hidden-line removal" and "hidden-surface removal" were developed to solve this problem. These processes are also sometimes called *occultation*, not because they involve magic, but because they remove *occulted* or hidden lines. (For information on wireframe views and shading see "Rendering" on page 34.)


SIZE DIFFERENCES

Objects that are the same size look larger when they are close to us and smaller when they are farther away. This occurs because the image made on the retina of the eye by a distant object is smaller than the image made by a similar-size object that is close.

In naive art, the relative size of objects in a picture often has more to do with their emotional importance or status. For example, a child may draw herself as the largest figure in a family group, towering over her mother and father. Obviously, the artistic purpose of the child is not to imitate reality but to express ideas and feelings, so accurate size relationships are not important.

Why do far objects look smaller?

Far objects look smaller because the light rays traveling from their outer edges to our eye converge more than the rays traveling from near objects and "draw" a smaller outline on the retina (and on a "picture plane" or pane of translucent material placed between the eye and the object). The height, width and depth of objects are all affected in this way, as is the space in which the objects exist.



But as systems of linear perspective developed (see next page), it became possible to accurately determine what size distant objects should be in order to produce images that look natural.

TEXTURE GRADIENTS

Regular patterns, either human made (such as a tiled floor or cobbled street) or natural (such as wildflowers in a field or waves on the ocean), appear denser as they recede into the distance. This effect is related, of course, to the fact that objects of the same size look smaller when they are farther away. So each paving brick in the street looks a little smaller than the one in front of it.

Artists often take advantage of this depth cue by using a tiled floor in their interior compositions. The diminishing size of the tiles as they move away from the viewer is reinforced by the receding parallel lines as they move toward the vanishing point.

When a regular pattern, such as a checkerboard, is applied to a 3D object as a surface texture, the program will calculate the rendering so that the pattern appears smaller and denser on those parts of the object that are further from the "camera."



Textures are denser at a distance

The cobblestones paving this street in Zurich, Switzerland look smaller as they move further from the viewer. This decrease in the size of a regular pattern is called a texture gradient and adds to the sense of depth in an image. (Photo by Janet Ashford)



The Art of Illusion



GETTING
STARTED
WITH 3D

10

LINEAR PERSPECTIVE

In Western art before the Renaissance, in most Oriental art and also in naive folk and children's art from around the world, accurate linear perspective is absent. The relation of objects in space is shown in a formal or stylized way, following conventions that are not mathematical; for example, objects which are further away are sometimes shown as being above objects which are closer, even though they may not be drawn smaller.

This type of nonscientific representation works well when accurate depiction is not the goal. For example, in Medieval Europe, artists were more concerned with depicting religious ideas than with duplicating reality. Paintings tended to be flat, decorative surfaces, with little sense of depth. No one expected a painting to look like a real scene. The same aesthetic still prevails in much Eastern art. But during the Renaissance, European artists, as well as philosophers and scientists, turned more toward nature and humanistic values and developed systems both for understanding and celebrating natural phenomena. One result of this change in emphasis was the development of systems of *linear perspective* to aid artists in depicting reality. Many Renaissance artists, including Leonardo Da Vinci and Albrecht Dürer, developed systems and wrote treatises on the relations between geometry and painting and worked out methods of translating 3D forms into 2D images.

THE BASIC CONCEPTS OF PERSPECTIVE

Perspective is defined in the *Oxford English Dictionary (OED)* as "the art of delineating solid objects upon a plane surface so that the drawing produces the same impression of apparent relative positions and magnitudes, or of distance, as do the actual objects when viewed from a particular point." The *art* of perspective is based on two related facts about the way we perceive visually: (1) that parallel lines appear to converge at a vanishing point located on the horizon and (2) that objects which are further away look smaller than those which are closer. Desktop 3D programs are all based on classical systems of perspective and employ the same key words and concepts.


HORIZON LINE

The *horizon* (from a Greek word for "the bounding circle") is the place where the sky meets the earth. If we were on the masthead of a ship far at sea, we could turn in a complete circle and view the entire horizon, which would describe a circle around us. But on land, our view of the horizon is usually obscured by trees, buildings or other objects. Even though the horizon is round, in graphic images it is usually drawn as a straight, horizontal line, since in a landscape view we usually see only a short segment of the horizon. In an image, this is called the *horizon line*.

Note that in real life, the horizon line is always at your own eye level. The place where sky meets earth will be at the level of your eyes whether you are standing, lying down, or at the top of a tall building. In a graphic image, the horizon line may be placed high, low or near the middle depending upon the height of the "eyes" of the viewer whose view is being presented by the artist.

The basic assumption of all perspective systems is that parallel lines never meet, but that they appear to do so.

—P. and L. Murray, *Dictionary of Art and Artists*, 1959



From theory to art

In a detail of Leonardo's *Annunciation* we can see that the parallel lines of the architectural elements on the right recede almost perfectly to a single vanishing point located on the distant horizon. The invisible lines that converge on the vanishing point (they're shown in red above) are called perspective lines, visual rays, or orthogonal lines. This way of drawing in perspective lets us know that the parallel lines of the building blocks and table are parallel with our line of sight. This is an example of 1-point perspective, in which all the lines that are parallel to the viewer's line of sight converge on a single vanishing point.

Viewing the horizon

The horizon line is always at the viewer's eye level, no matter how high or low the viewer may be. When you are standing, the horizon line falls toward the middle of the view. When you are lying down, the horizon line is also low and more of the sky is seen. When you are elevated, the horizon line is raised and more of the ground plane is seen. To illustrate this we constructed a model in Ray Dream Designer using clip art of a city block and various creatures. The top image shows a high horizon (flying duck's-eye view), the center image shows a more centered horizon (standing velociraptor's-eye view) and the bottom image shows a low horizon (ladybug's-eye view).

VANISHING POINT

The *vanishing point* (or *principal vanishing point*) is a point on the horizon toward which we are looking and toward which all the parallel lines in the scene that are parallel to our line of sight will appear to converge. Other receding lines in a scene that are not parallel to the line of sight will appear to converge on other vanishing points. These points may be to the left or the right of the principal vanishing point and may also be above or below the horizon line. These other vanishing points may actually lie outside the boundaries of the image. (See the section on 1-, 2- and 3-point perspective on pages 14–15.)

VIEWPOINT





The *viewpoint* or *station* point is the point at which the viewer of a scene stands. In a 3D model, this point is often called the "camera."

LINE OF SIGHT

The *line of sight* is the line between the viewpoint and the principal vanishing point.


PICTURE PLANE (OR PROJECTION PLANE)

The picture plane is an imaginary plane that is perpendicular to the line of sight and that stands between the viewer or viewpoint and the scene being viewed.



Drawing on the picture plane

The "picture plane" is a hypothetical plane placed in between the observer (viewpoint) and the object (in this case a house) and oriented perpendicular to the line of sight. The situation is similar to that of viewing an object through a window. In theory, the object's image could be drawn by tracing its outlines on the glass of the window. To illustrate this, we constructed a model in Ray Dream Designer using clip art. The picture plane is modeled by a flat rectangle with a clear glass texture map. The visual rays marked above are drawn from the viewpoint to the four corners of the front of the house. As they pass through the picture plane they



describe a rectangle (shown in red lines) that represents the viewer's view of the house front. A rendering made from the same model but with the camera placed in back of the viewer, gives us a view similar to what the viewer sees on or "through" the picture plane.







VISUAL RAYS

Perspective depends upon the concept that when we look at an object, every point on its surface sends a visual ray in a straight line to our eyes. These rays converge on us, so to speak, in the same way that lines drawn along parallel edges of objects converge on a vanishing point. There are an infinite number of such hypothetical visual rays, but important ones can be traced by simply drawing lines from the viewpoint to the object (as in the figure at the right) or by drawing lines from objects to the vanishing point (as in the illustration on page 10).

Visual rays are important in computer 3D because the rendering of 3D models into images uses a process called “ray tracing,” based on the existence of these hypothetical rays. In theory, the “ray” traveling from the model to each pixel in the computer image (which is analogous to the picture plane) determines what the color and intensity of that pixel will be. To assign the correct values, the ray tracing algorithm works backward, tracing the path of the ray from the pixel back to the model, taking into account whether it came directly from a light source, or was reflected by a surface, or was refracted (bent) through a material like glass.

FORESHORTENING

Foreshortening is a distortion that occurs when we look at an object that recedes away from us, because the closer parts of it look larger than the parts that are further away. For example, when we view a person standing up, most of their body is equally distant from us, so their hands look about as large as their feet. However, when we view a person who is lying down, with their feet toward us, their feet will look larger than their hands because the feet are closer. Foreshortening is a challenge to traditional artists, since it requires that we really look at the outlines of an object rather than draw what we expect to see. However, foreshortening is done automatically in a computer 3D program, whenever we move the camera to view an object across its longest dimension.



Foreshortening

When we view a long (or tall) object from a viewpoint in which our line of sight runs perpendicular to the long axis of the object, its length is not distorted. For example, in the diagram above we can see exactly how tall (or long) both mannequins are, because we are looking at them from the side. But when we view the lying down figure from the front (as the mannequin’s friend is doing) so that the line of sight is running parallel to the figure’s long axis, the mannequin’s length is shortened and its shape is distorted. This occurs because the parts of the mannequin that are closer to us (such as the feet) look larger than similar-sized parts (such as the hands) that are further away. If we draw a box grid around the supine mannequin and draw visual rays from five equidistant points on the grid to the viewpoint (the other mannequin’s head), we can see that the distance between the rays becomes unequal as they pass through the picture plane, which is the hypothetical surface that makes up the image we see. The foreshortened image seen by the standing mannequin is shown at left. It was created by placing a camera in the model at the same location as the standing mannequin’s head.



Variations

A photograph of Italian fishing boats illustrates the principle of foreshortening. The length of the boats is condensed into shorter, squatter shapes, creating an interesting image of theme and variations. (Photo by Janet Ashford)

The Science of Drawing

GRIDS AND RAY TRACING IN THE 1500S




Can't draw a straight line, much less a curved one? Try the methods of Albrecht Dürer (1471–1528), a German engraver, painter and draftsman who was one of the greatest artists of his time and a lover of linear perspective. His concern with proper proportion and the application of mathematical systems to art made him the northern counterpart of da Vinci, who also expounded on systems of perspective. Several of Dürer's woodcuts (two are shown here; one is on this chapter's opening page) illustrate his methods for drawing objects accurately, so that the foreshortening of the image on paper is similar to what we see when looking at the object itself.

Dürer's method follows closely the definition given by the *Oxford English Dictionary* for linear perspective: "... an application of projective geometry, in which the drawing is such as would be made upon a transparent vertical plane (*plane of delineation*) interposed in the proper position between the eye and the object, by drawing straight lines from the position of the eye (*point of sight*) to the several points of the object, their intersections with the plane of delineation forming the corresponding points of the drawing."

All of Dürer's woodcuts are reproduced in *The Complete Woodcuts of Albrecht Dürer*, edited by Willi Kurth (Dover Publications, 1963).

From *Underweisung der Messung*, 1525, Albrecht Dürer. In this woodcut from Dürer's treatise on geometry, two men demonstrate an early form of "ray tracing," the method used today in computer 3D to determine how 3D objects should be rendered on a flat, 2D surface.

Here's how this laborious process worked:



Creating a visual ray

A spot on the wall to the right of the table is marked by a small hook which represents the "viewpoint." A string attached to the hook passes through an open wooden frame which is perpendicular to the table. The plane defined by this frame is the "picture plane" and it is also perpendicular to the "line of sight." (The line of sight is defined as a straight line between the viewpoint and the center of the object being viewed, in this case a lute. It can also be called a visual "ray").

Viewing the object

The other end of the string is tied to a pointer, held by the man on the left. This man (let's call him the pointer) touches various points along the contours (edges) of the lute. This is analogous to looking at or viewing the various parts of the lute.

Marking the visual rays

As the pointer touches each point, the other man (we'll call him the plotter) moves two cross-hairs (probably threads made of silk) which are attached to the wooden frame so that they slide. One cross-hair is vertical and slides from left to right, while the other is horizontal and slides from top to bottom. The plotter slides each cross-hair until together they mark the position at which the line-of-sight string or "ray" passes through the frame. Plotting the points


The string is untied for a moment and the plotter swings shut a hinged drawing surface, attached to the frame, on which he has tacked a piece of paper. Using the cross-hairs as a guide, the plotter marks the spot where the visual ray passed through. Then he swings open the drawing surface, the pointer moves his pointer to another spot on the lute and they continue the plotting process until they have produced a foreshortened drawing of the lute delineated by dots.

From *De Symmetria Partium Humanorum Corporum*, 1532, Albrecht Dürer.

In this illustration, Dürer demonstrates a grid method of drawing in which the artist places a grid in a vertical frame between himself and his subject, views the subject through the grid and draws what he sees in each grid "square" on a corresponding grid on paper on the table in front of him. To make sure that his eye view remains constant, the artist always aligns his nose with a vertical guide.



THE WORLD OF THREE DIMENSIONS





1-point perspective

In a 1-point perspective view, objects (such as the cubes shown above, for example) are orthogonal or aligned at right angles to the viewer; that is, the horizontal and vertical lines of the cubes' front and back sides all lie in planes that are perpendicular to our line of sight, while the receding parallel lines of the cubes' sides are parallel with line of sight. All of the receding parallel lines converge on a single vanishing point.

In this type of perspective or "projection," objects that straddle the horizon line show only a front and a side surface, but objects above the horizon line also show a bottom surface and objects below the horizon line show a top surface.



The photo below of a symmetrical hallway in Balboa Park, San Diego, provides a perfect example of 1-point perspective in real life. Because the viewer's line of sight passes through the center of the hallway, all the lines of the building which are parallel with the line of sight converge on a single vanishing point at the end of the hall. (Photo by Janet Ashford)



2-point perspective


When the objects in a scene (such as the cubes above) turn an edge or corner towards the viewer (instead of being "orthogonal" as in the 1-point perspective view at left) then 2-point perspective comes into play. In this situation, some of the receding parallel lines extending from the objects converge on a vanishing point to the left, while others converge on a vanishing point to the right.

Two-point perspective is often encountered in life, as in the photo below of the Point Loma lighthouse in San Diego. Here the line of sight runs from the viewer to a corner of the lighthouse, with the two side walls oriented almost perfectly at 45 degree angles from the line of sight. We can find the two vanishing points by drawing the horizon and then extending receding lines from the walls of the building to the horizon, following the angles of the building. Often in a 2-point perspective view, the vanishing points lie outside the boundaries of the scene. Two-point perspective is a popular view in 3D graphics programs. (Photo by Janet Ashford)



3-point perspective

When objects are viewed from a position that is extremely low (worm's eye view) or extremely high (bird's eye view), then a third vanishing point is created. For example, when looking up at a very tall building, not only do the horizontal rays along the sides of the building recede to two vanishing points on the horizon, as in a 2-point projection, but the parallel vertical rays running up and down the building converge on a third vanishing point in the sky. A photo looking up at a tall bell tower at Balboa Park in San Diego shows this effect clearly. Lines drawn along the main vertical edges of the tower converge on a vanishing point somewhere just off the top of this page. (Photo by Janet Ashford)



DIFFERENT TYPES OF PERSPECTIVE

Different types of perspective occur when objects are viewed from different viewpoints and angles. Schemes of 1-, 2- and 3-point perspective have been devised to accommodate these differences and are used extensively in drawing and painting, as well as in computer 3D. In real life, however, since our eyes are constantly roaming over the scenes around us, several types of perspective may occur simultaneously.

1-POINT PERSPECTIVE

In a 1-point perspective view there is only one vanishing point and all lines that are parallel to the line of sight converge on this point.

2-POINT PERSPECTIVE

In this view, sometimes called an “off-axis” view, the primary receding lines of the objects are not parallel with the line of sight and they recede to two different vanishing points, located at the left and right of the picture.

3-POINT PERSPECTIVE

When objects are viewed from below or above, sometimes a third vanishing point appears, towards which the vertical lines of the objects recede.

ORTHOGRAPHIC VIEWS

In addition to views with linear perspective, most 3D programs also include *orthographic* views. *Ortho* means “straight” or “at right angles to” something else; so an orthographic view is one in which the direction of the projection is perpendicular (at right angles) to the picture plane. In 3D modeling these include the views

called front, back, left, right, top and bottom. In an orthographic view, the parallel lines in the model do not converge on a vanishing point but remain parallel; also, objects that are the same size look the same size even when they are further away. Many 3D program manuals recommend the use of an orthographic view during the initial phase of modeling, since it is easier to measure and align objects.


ISOMETRIC VIEW

In some 3D programs, there is also an *isometric* view, which is a type of orthographic view in which the principal axes (x , y , z) of the object are equally foreshortened so that the same measure scale can be applied to each axis. An isometric projection looks similar to a projection using linear perspective, so it gives a sense of depth, but it is easier to measure and manipulate. This type of projection is often used in mechanical drawing.

For more information on orthographic and isometric views see the section on “Viewing the 3D World” on pages 26–29.

Following the converging lines

A perspective drawing from a 19th-century engraving shows a carefully drawn grid of converging lines receding toward a vanishing point on the horizon. The building was drawn in perspective by aligning its structures to the grid. We used Strata StudioPro to construct a model similar to the one in the engraving. Viewing it in a 1-point perspective projection, we can see that the computer’s view differs a little from the hand engraving, but the parallel beams of the 3D structure do appear to converge on a vanishing point. However, viewing the same model in an isometric projection, there is no converging of parallel elements.



The first objective of the painter is to make a flat plane appear as a body in relief and projecting from that plane.

—Leonardo da Vinci



GETTING
STARTED
WITH 3D

16

LIGHT AND DARK

In addition to perspective, another essential technique for creating the illusion of depth in a two-dimensional image is *chiaroscuro*, or the modeling of forms with areas of light and dark. This method of shading was devised during the Renaissance and again, Leonardo da Vinci was one of its greatest exemplars.

With the technique of *chiaroscuro*, objects do not have clear outlines (as in a coloring book) but appear to emerge from darkness, as though being revealed by the light hitting them. Often a highlight appears where the light strikes an object and reflects brightly. In a black-and-white drawing the highlight may be pure white, or a pale shade of gray, depending upon the reflectiveness of the object being drawn. Darker shades of gray blend gradually across the object as the light decreases, until they reach the edge where the lighted part of the object meets the unlit part. Here there is a more rapid transition to dark shades of gray, though the unlit side of an object rarely becomes completely black because of the presence of reflected or ambient light in most scenes.

The shadows cast by lighted objects are also rendered in varying shades of gray. The part of the shadow that is closest to the object is usually the darkest. It becomes lighter as it recedes away from the object, as it is affected by ambient light.



Leonardo and *chiaroscuro*

Writing in his *History of Art* (Prentice-Hall and Harry N. Abrams, 1962), H. W. Janson observes the technique of da Vinci: "Leonardo ... thinks not of outlines, but of three-dimensional bodies made visible, in varying degrees, by the incidence of light. In the shadows, these shapes remain incomplete, their contours are merely implied. In this method of modeling (called *chiaroscuro*, "light-and-dark") the forms no longer stand abruptly side by side but partake of a new pictorial unity, the barriers between them having been partially broken down."

In this drawing from his notebooks, we can see the process of *chiaroscuro* at work, as Leonardo uses a light pencil sketch to rough in the outlines of the woman's face, head and shoulders and then works into the face with shadow to model the forms of mouth, eyes and nose so that they appear to almost rise from the surface of the paper.

THE SYMBOLISM OF LIGHT AND DARK

Chiaroscuro shading of this type can be very subtle and expressive of emotional states. In fact, light and shadow are seen as symbols for more elevated dualities by one critic speaking of Rembrandt's work:

The richness and variety of Rembrandt's achievements are inexhaustible, but close to their centre is a curious duality of light and shadow, manifestation and concealment, matter and spirit.

The light and shade of his mature art is most striking; it fluctuates unaccountably, it hides what it might reveal. As light and shade alternate across the form of the painting, it is hard at first glance to perceive whether the shifts are caused by changing planes, varying textures, or some unregarded object casting its shadow. Yet the light and shade are expressions of a more profound duality.

—(J. M. Nash, *The Age of Rembrandt and Vermeer* Phaidon Press, 1972)



Exploring light in photography

The photo at left, *Tea Still Life*, was taken by Heinrich Kuehn around 1908. It shows the soft focus and lighting that were being explored by photographers of this period. The use of chiaroscuro (light and dark areas) to delineate the forms is especially striking. The photo is part of the Alfred Steiglitz Collection at the Metropolitan Museum of Art in New York.

The Metropolitan Museum of Art, The Alfred Steiglitz Collection, 1933. (33.43.273)
All rights reserved, The Metropolitan Museum of Art

CHIAROSCURO ON THE COMPUTER


Shading with values of light and dark is also done in computer 3D, using algorithms that calculate the effects of lights placed in a 3D “scene,” taking into account the color, intensity and direction of the lights as well as the surface properties and reflectiveness of the objects in the scene. In the same way as traditional chiaroscuro, the calculations made by the computer “reveal” the objects in a model by shading them.

GOING BEYOND ACCURACY

The moodiness of a chiaroscuro treatment can sometimes be captured in photography, particularly when a subject is lit by a single light. Around the turn of the century, in fact, the members of the Photo Secessionist movement sought to lessen the “wealth of trivial detail” that photographs could capture and instead used different printing techniques to produce photographs that looked more like paintings than photos.

In the new art of computer 3D imaging, the limits of memory and programming capability mean that light and shadow and the modeling they create are less sophisticated than in painting or photography, with tones and the transitions between them sometimes appearing more harsh. However, the crisp, hard-edged look that often results in computer 3D shares in the aesthetic of the Super Realist school of painting that developed during the 1960s. This way of painting exaggerates trivial detail to produce images that are even more “realistic” than the real world. (For an example of a 3D rendering in super realist style see the motorcycle on page 98.)

At the same time, various techniques can be used to soften the effects of light and dark in a 3D image—both



From plain to fancy

To explore the use of chiaroscuro in computer 3D we used clip art in Ray Dream Designer to assemble a still life similar to the photo above. Using a single point light source, we made a rendering that clearly delineates the forms but lacks charm or emotional content (A). 3D renderings often tend to be rather utilitarian (similar to product photos in a catalog, for example) rather than aesthetically pleasing. However, renderings can be modified in an image editing to add a more artistic touch. To create a more dramatic “photo realist” treatment we increased the contrast and saturation and applied Photoshop’s Watercolor filter (B). To create a softer, more glowing version, similar to the Kuehn photograph above, we increased the contrast, applied a Gaussian blur and applied the Paint Daubs filter (C). Some 3D programs, including Designer, also make it possible to add filtered effects during the rendering process.

before and after rendering—so that it looks less like a digital image and more like a painting of the old school. One can use atmospheric effects in the model or edit the rendered image afterward with filters. (We describe more techniques for achieving a “painterly” look with 3D in “Applying filter Effects” on page 158.) In either case, the world of feeling and style in 3D imagery is just beginning to be explored.



Looking Through a Black Box

THE CAMERA OBSCURA AND THE CAMERA LUCIDA AS EARLY DRAWING AIDS

The camera is one of the greatest aids to drawing ever invented. The camera does such an excellent job of translating real world objects into flat images on paper, that after the invention of film-based cameras in the 19th century, artists largely abandoned the task of representational art to develop the more subjective styles of Impressionism and Expressionism.

But early cameras, without film, had been used by artists since the Renaissance as mechanical aids to drawing. One of the early precursors of the modern camera is the *camera obscura*, from the Latin words for “dark chamber.” This device is a darkened box or sometimes a darkened room with a small opening in one wall containing a double convex lens. Images outside the box are projected onto the interior wall opposite the lens.

A related device, called a *camera lucida* (“light chamber”) is an instrument in which the rays of light from an object or scene pass through a special prism and are reflected downward onto a piece of paper placed under the device. In this way, an artist can trace the contours of the objects being reflected. Many historians believe that the Dutch artist Jan Vermeer (1632–1675) used a *camera obscura* or *lucida* to project scenes in his house onto a canvas for tracing. Vermeer’s paintings display not only very accurate perspective, but also include highlights or blobs of light and also unfocused areas, both of which appear to be artifacts of the lenses he used.

Computer 3D images can be used in the same way, not only as images in themselves, but as guides for painting or drawing. For more information see “Using 3D Images as References” on page 156.

20th Century Vermeer

Vermeer excelled at realistic scenes of everyday life, especially interiors in which everyday people and things were transformed by light into extraordinary objects. *The Music Lesson* (or *Lady and Gentleman at the Virginals*) shows a room in Vermeer’s own house and is one of several paintings made in that room that have fascinated art historians, who have speculated that Vermeer used a *camera lucida* to aid his rendering of perspective. Certainly his paintings have a modern, photographic quality.


Researchers at Cornell University’s Program of Computer Graphics created a 3D graphics model of Vermeer’s room to illustrate a scientific paper (“A Two Pass Solution to the Rendering Equation: a Synthesis of Ray Tracing and Radiosity Methods” by John R. Wallace, Michael F. Cohen and Donald P. Greenberg, 1987).

In addition, Professor Phillip Steadman of The Open University in England built a physical model of the room and furnished it with paintings, objects and mannequins in period dress to produce a photograph that closely resembles Vermeer’s painting. An article by Steadman describing his work (“In the Studio of Vermeer”) appears in *The Artful Eye* (Gregory, Harris, Heard and Rose, editors) published by Oxford University Press in 1995. All the images shown on this page were taken from World Wide Web sites maintained by Cornell University, Professor Steadman and aficionados of Vermeer.

Physical model of *The Music Lesson* constructed by Phillip Steadman of the Open University, England



The Music Lesson by Vermeer, 1662–65
(in the collection of Buckingham Palace, London)



Computer rendering of *The Music Lesson* made at Cornell University



GETTING
STARTED
WITH 3D

ATMOSPHERIC PERSPECTIVE

Another key to depicting three-dimensional reality is the fact that our eyes perceive distant objects as less clear and distinct than those in the foreground. Near objects look sharper and more focused than those that are far away. Also, objects in the distance look bluer and more monochromatic than foreground objects, which have brighter, richer colors. This blueness is often seen in landscapes, both real and painted, in which distant hills, for example, are rendered in soft blues, in contrast with the brighter greens and browns of foreground features. Atmospheric perspective was first described and named by Leonardo da Vinci, who called it *aerial perspective*.

WHY DISTANT OBJECTS ARE BLUE

Distant objects look blue for the same reason the sky looks blue—because blue light is scattered by the atmosphere more easily than light of redder colors. The white light that comes to us from the sun contains a whole rainbow of colors, which we can see when light is passed through a prism. But the light from the blue end of this rainbow spectrum has shorter wavelengths than does light from the red end. So blue light is more likely to be scattered or deflected by the particles of oxygen, water and dust in the atmosphere. In a way, the whole atmosphere acts like a prism, but because of the match between the short wavelengths of blue light and the small size of the atmospheric particles, the blue light gets deflected the most. This scattering of blue light occurs throughout the atmosphere, so that when we look up at the sky we see blue light coming to us from all directions.

In addition, when we look at distant objects we are looking through more atmosphere than when we look at near objects. So we are receiving more blue light (from

Getting bluer and bluer

This photo of the coastline of northern Italy shows five hillsides jutting out into the Mediterranean. The hill closest to us is full of rich color, with sharp focus and detail. But though the plants and rocks on each hillside are essentially the same, each distant one appears less saturated, less focused, less detailed and bluer as the hills recede away from the viewer. (Photo by Janet Ashford)



the atmosphere), along with the reflected light showing the color of the objects themselves. So distant objects tend to look bluer. The further away they are the bluer they look.

WHY DISTANT OBJECTS LOOK FUZZY

In addition to creating a blue tone, the scattering of the blue light rays coming to us from distant objects means that they look less focused than near objects, whose light rays travel through less atmosphere on their way to our eyes.

Most computer 3D programs include atmospheric controls (sometimes called fog) that mimic the effects of aerial perspective. A blue haze may be applied to the more distant objects in a model (controls allow the user to specify the starting and ending position of the fog) or may also be applied to the foreground as well, to imitate the technique of *sfumato*.



Creating a poetic vision

In addition to acting as an important depth cue, mist and soft focus can also be used by artists to create emotional effects. Writing of Leonardo da Vinci in his *History of Art* (Prentice-Hall and Harry N. Abrams, 1962), H. W. Janson describes his use of these techniques in the painting *The Virgin of the Rocks*: “Here the figures emerge from the semidarkness of the grotto, enveloped in a moisture-laden

atmosphere that delicately veils their forms. This fine haze (called *sfumato*) ... lends a peculiar warmth and intimacy to the scene. It also creates a remote, dreamlike quality and makes the picture seem a poetic vision rather than an image of reality pure and simple.” In this case, the artist has applied a haze even to objects that are closer to the viewer.


A similar effect can be seen in another painting of the Virgin by da Vinci, shown above. In a detail of the landscape, we can see clearly the artist’s use of blue paint and soft focus to indicate the distance of the far mountains.





“FICTIVE SPACE”

Through the use of rational systems of perspective along with the emotional values created by the depiction of light and atmosphere, Renaissance artists were able to create a *fictive space* in their landscapes and backgrounds that not only provides a convincing imitation of reality, but that evokes feelings of awe, wonder and pleasure in the viewer. Desktop 3D programs employ all three traditional techniques (perspective, chiaroscuro and atmospheric depth cues) as well as other monocular depth cues (occlusion, size differences and texture gradients) to create renderings that imitate reality, and in skilled hands 3D graphics can also be used to create images that touch the emotions.



Breaking the rules

One of Gauguin's paintings of Tahiti manages to create a convincing fictive space while violating many of the usual cues. For example, while he does make use of naturalistic size differences and occlusion, Gauguin painted both foreground and background objects with equal detail and saturation, choosing to dispense with the visual cues of atmospheric perspective. As a result, the painting is both a realistic landscape and a flat, decorative surface. 3D graphics images can be modified with filters to achieve similar effects.



Creating a sense of place

The Impressionists painted over 300 years after the height of the Renaissance, yet they employed the representational techniques developed by the earlier masters (or else played off tradition by breaking the rules). *Paris Street* (left) was painted by Caillebotte in the late 1800s and provides a beautiful example of depth and a sense of place using all the monocular depth cues described in this chapter. The artist's use of a texture gradient on the street is especially striking, as is the dramatic use of size differences (among the pedestrians with umbrellas) and atmospheric perspective.

The eerie towers in our 3D rendering (below) show occlusion, size differences and atmospheric perspective, even though the model itself is quite simple. The techniques of representing depth that artists have used over the centuries, including the logic of perspective, are built into today's 3D programs. But we can still learn from the masters how to apply those techniques imaginatively to create a unique sense of place in computer generated images.

