THE ECOLOGICAL APPROACH TO VISUAL PERCEPTION

Classic Edition

James J. Gibson

Psychology Press
Taylor & Francis Group
NEW YORK AND LONDON

Classic Edition published 2015
by Psychology Press
711 Third Avenue, New York, NY 10017

and by Psychology Press
27 Church Road, Hove, East Sussex BN3 2FA

Psychology Press is an imprint of the Taylor & Francis Group, an informa business

First published by Taylor & Francis Group, 1986

Library of Congress Cataloging in Publication Data
A catalog record for this book has been requested

ISBN: 978-1-84872-577-5 (hbk)
ISBN: 978-1-84872-578-2 (pbk)
ISBN: 978-1-315-74021-8 (ebk)

Typeset in Bembo
by RefineCatch Limited, Bungay, Suffolk, UK

TABLE OF CONTENTS

5

THE AMBIENT OPTIC ARRAY

The central concept of ecological optics is the ambient optic array at a point of observation. To be an *array* means to have an arrangement, and to be *ambient at a point* means to surround a position in the environment that could be occupied by an observer. The position may or may not be occupied; for the present, let us treat it as if it were not.

What is implied more specifically by an *arrangement?* So far I have suggested only that it has *structure,* which is not very explicit. The *absence of structure* is easier to describe. This would be a homogeneous field with no differences of intensity in different parts. An array cannot be homogeneous; it must be heterogeneous. That is, it cannot be undifferentiated, it must be differentiated; it cannot be empty, it must be filled; it cannot be formless, it must be formed. These contrasting terms are still unsatisfactory, however. It is difficult to define the notion of structure. In the effort to clarify it, a radical proposal will be made having to do with *invariant* structure.

What is implied by *ambient at a point?* The answer to this question is not so difficult. To be ambient, an array must surround the point completely. It must be environing. The field must be closed, in the geometrical sense of that term, the sense in which the surface of a sphere returns upon itself. More precisely, the field is unbounded. Note that the field provided by a picture on a plane surface does not satisfy this criterion. No picture can be ambient, and even a picture said to be panoramic is never a completely closed sphere. Note also that the temporary field of view of an observer does not satisfy the criterion, for it also has boundaries. This fact is obviously of the greatest importance, and we shall return to it in Chapter 7 and again in Chapter 12.

Finally, what is implied by the term *point* in the phrase *point of observation?* Instead of a geometrical point in abstract space, I mean a position in ecological space, in a medium instead of in a void. It is a place where an observer *might* be and from which an act of observation *could* be made. Whereas abstract space consists of points, ecological space consists of places—locations or positions.

A sharp distinction will be made between the ambient array at an unoccupied point of observation and the array at a point that is occupied by an observer, human or other. When the position becomes occupied, something very interesting happens to the ambient array: it contains information about the body of the observer. This modification of the array will be given due consideration later.

The point of observation in ecological optics might seem to be the equivalent of the station point in perspective geometry, the kind of perspective used in the making of a representative painting. The station point is the point of projection for the picture plane on which the scene is projected. But the terms are not at all equivalent and should not be confused, as we shall see. A station point has to be stationary. It cannot move relative to the world, and it must not move relative to the picture plane. But a point of observation is never stationary, except as a limiting case. Observers move about in the environment, and observation is typically made from a moving position.

How is Ambient Light Structured? Preliminary Considerations

If we reject the assumption that the environment consists of atoms in space and that, hence, the light coming to a point in space consists of rays from these atoms, what do we accept? It is tempting to assume that the environment consists of *objects* in space and that, hence, the ambient array consists of *closed-contour forms* in an otherwise empty field, or "figures on a ground." For each object in space, there would correspond a form in the optic array. But this assumption is not close to being good enough and must also be rejected. A form in the array could not correspond to each object in space, because some objects are hidden behind others. And in any case, to put it radically, the environment does not consist of objects. The environment consists of the earth and the sky with objects *on* the earth and *in* the sky, of mountains and clouds, fires and sunsets, pebbles and stars. Not all of these are segregated objects, and some of them are nested within one another, and some move, and some are animate. But the environment is all these various things—places, surfaces, layouts, motions, events, animals, people, and artifacts that structure the light at points of observation. The array at a point does not consist of forms in a field. The figure-ground phenomenon does not apply to the world in general. The notion of a closed contour, an outline, comes from the art of drawing an object, and the phenomenon comes from the experiment of presenting an observer with a drawing to find out what she perceives. But this is not the only way, or even the best way, to investigate perception.

We obtain a better notion of the structure of ambient light when we think of it as divided and subdivided into component parts. For the terrestrial environment, the sky-earth contrast divides the unbounded spherical field into two hemispheres, the upper being brighter than the lower. Then both are further subdivided, the lower much more elaborately than the upper and in quite a different way. The components of the earth, as I suggested in Chapter 1, are nested at different levels of size—for example, mountains, canyons, trees, leaves, and cells. The components of the *array* from the earth also fall into a hierarchy of subordinate levels of size, but the components of the array are quite different, of course, from the components of the earth. The components of the array are the *visual angles* from the mountains, canyons, trees, and leaves (actually, what are called *solid angles* in geometry), and they are conventionally measured in degrees, minutes, and seconds instead of kilometers, meters, and millimeters. They are *intercept angles,* as we shall see. All these optical components of the array, whatever their size, become vanishingly small at the margin between earth and sky, the horizon; moreover, they change in size whenever the point of observation moves. The substantial components of the earth, on the other hand, do not change in size.

There are several advantages in conceiving the optic array in this way, as a nested hierarchy of solid angles all having a common apex instead of as a set of rays intersecting at a point. Every solid angle, no matter how small, has form in the sense that its cross-section has a form, and a solid angle is quite unlike a ray in this respect. Each solid angle is unique, whereas a ray is not unique and can only be identified arbitrarily, by a pair of coordinates. Solid angles can fill up a sphere in the way that sectors fill up a circle, but it must be remembered that there are angles within angles, so that their sum does not *add* up to a sphere. The surface of the sphere whose center is the common apex of all the solid angles can be thought of as a kind of transparent film or shell, but it should not be thought of as a picture.

The structure of an optic array, so conceived, is without gaps. It does not consist of points or spots that are discrete. It is completely filled. Every component is found to consist of smaller components. Within the boundaries of any form, however small, there are always other forms. This means that the array is more like a hierarchy than like a matrix and that it should not be analyzed into a set of spots of light, each with a locus and each with a determinate intensity and frequency. In an ambient hierarchical structure, loci are not defined by pairs of coordinates, for the relation of location is not given by degrees of azimuth and elevation (for example) but by the relation of inclusion.

The difference between the relation of *metric location* and the relation of *inclusion* can be illustrated by the following fact. The stars in the sky can be located conveniently by degrees to the right of north and degrees up from the horizon. But each star can also be located by its inclusion in one of the

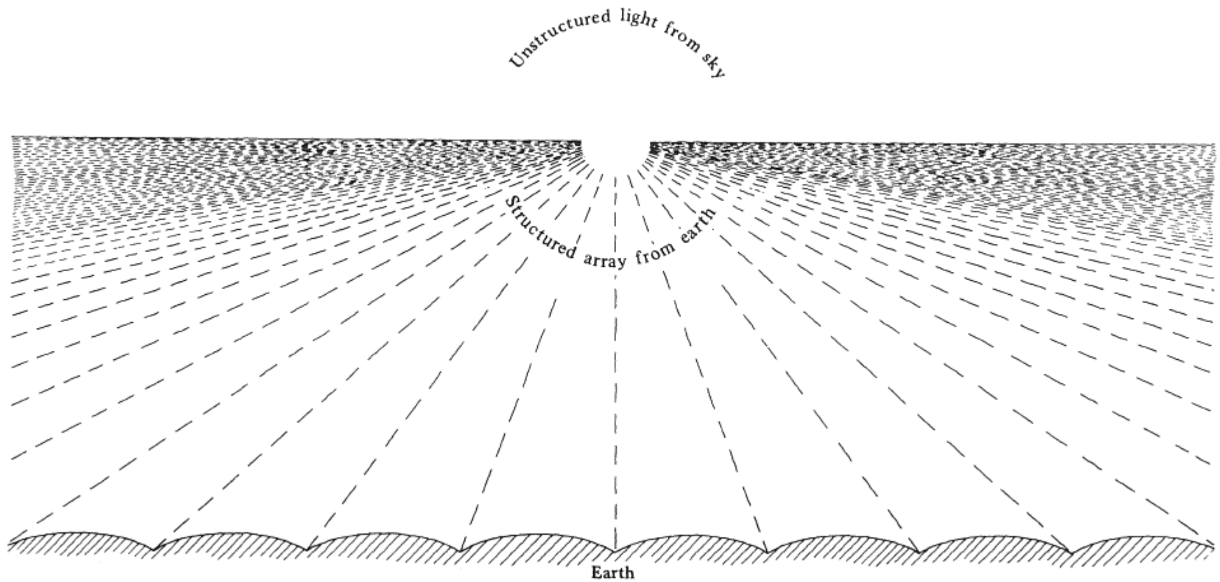

FIGURE 5.1 The ambient optic array from a wrinkled earth outdoors under the sky.

In this illustration it is assumed that illumination has reached a steady state. The earth is shown as wrinkled or humped, but not as cluttered. The dashed lines in this drawing depict the envelopes of visual solid angles, not rays of light. The nesting of these solid angles has not been shown. The contrasts in this diagram are caused by differential illumination of the humps of the earth. Compare this with the photograph of hills and valleys in Figure 5.9. This is an optic array at a single fixed point of observation. It illustrates the main invariants of natural perspective: the separation of the two hemispheres of the ambient array at the horizon, and the increasing density of the optical texture toward its maximum at the horizon. These are invariant even when the array flows, as it does when the point of observation moves.

constellations and by the superordinate pattern of the whole sky. Similarly, the optical structures that correspond to the leaves and trees and hills of the earth are each included in the next larger structure. The texture of the earth, of course, is dense compared to the constellations of discrete stars and thus even less dependent than they are on a coordinate system. If this is so, the perception of the direction of some particular item on the earth, its direction-from-here, is not a problem in its own right. The perceiving of the environment does not consist of perceptions of the differing directions of the items of the environment.

The Laws of Natural Perspective: The Intercept Angle

The notion of a visual angle with its apex at the eye and its base at an object in the world is very old. It goes back to Euclid who postulated what he called a "visual cone" for each object in space. The term is not exact, for the object need not be circular and the figure does not have to be a cone. Ptolemy spoke of the "visual pyramid," which implied that the object was rectangular. Actually, we should refer to the *face* of an object, which can have any shape whatever, and to a corresponding *solid angle,* having an envelope. A cross-section of this envelope is what we call the *outline of the object.* We can now note that the solid angle shrinks as the distance of the object from the apex increases, and it is laterally squeezed as the face of the object is slanted or turned. These are the two main laws of perspective for objects. Euclid and Ptolemy and their successors for

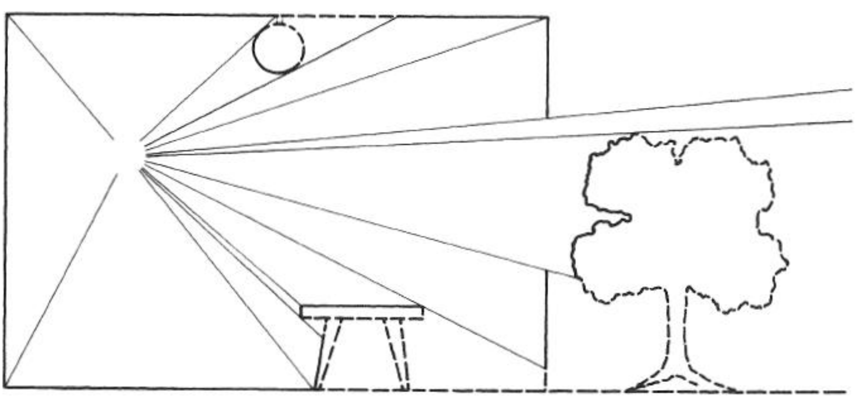

FIGURE 5.2 The ambient optic array from a room with a window.

This drawing shows a cluttered environment where some surfaces are projected at the point of observation and the remainder are not, that is, where some are unhidden and the others are hidden. The hidden surfaces are indicated by dotted lines. Only the faces of the layout of surfaces are shown, not the facets of their surfaces, that is, their textures.

many centuries never doubted that objects were seen by means of these solid angles, whether conical, pyramidal, or otherwise. They were the basis of ancient optics. Nothing was then known of inverted retinal images, and the comparison of the eye with a camera would not be made for a thousand years. The ancients did not understand the eye, they were puzzled by light, they had no conception of the modem doctrine that nothing gets into the eye but light, but they were clear about visual angles.

The conception of the ambient optic array as a set of solid angles corresponding to objects is thus a continuation of ancient and medieval optics. Instead of only freestanding objects present to an eye, however, I postulate an environment of illuminated surfaces. And instead of a group of solid angles, I postulate a nested complex of them. The large solid angles in the array come from the *faces* of this layout, from the facades of detached objects, and from the interspaces or holes that we call background or sky (which Euclid and Ptolemy seem never to have thought of). The small solid angles in the array come from what might be called the *facets* of the layout as distinguished from the faces, the textures of the surfaces as distinguished from their forms. As already has been emphasized, however, the distinction between these size-levels is arbitrary.

Natural perspective, as I conceive it, is the study of an ambient array of solid angles that correspond to certain distinct geometrical parts of a terrestrial environment, those that are separated by edges and corners. There are elegant trigonometric relations between the angles and the environmental parts. There are gradients of size and density of the angles along meridians of the lower half of the array, the earth, with sizes vanishing and density becoming infinite at the horizon. These relations contain a great amount of information about the parts of the earth. No one who understood them would think of questioning their validity. It is a perfectly clear and straightforward discipline, although neglected and undeveloped. But the environment does not *wholly* consist of sharply differentiated geometrical parts or forms. Natural perspective does not apply to shadows with penumbras and patches of light. It does not apply to sunlit surfaces with varying degrees of illumination. It geometrizes the environment and thus oversimplifies it. The most serious limitation, however, is that natural perspective omits motion from consideration. The ambient optic array is treated as if its structure were frozen in time and as if the point of observation were motionless.

Although I have called this discipline *natural perspective,* the ancients called it *perspectiva,* the Latin word for what we now call *optics.* In modem times, the term *perspective* has come to mean a technique—the technique of picture-making. A picture is a surface, whether it be painted by hand or processed by photography, and perspective is the art of "representing" the geometrical relationships of natural objects on that surface. When the Renaissance painters discovered the procedures for perspective representation, they very properly called the method *artificial perspective.* They understood that this had to be distinguished from the natural perspective that governed the ordinary perception of the environment. Since that time we have become so picture-minded, so dominated by pictorial thinking, that we have ceased to make the distinction. But to confuse pictorial perspective with natural perspective is to misconceive the problem of visual perception at the outset. The so-called cues for depth in a picture are not at all the same as the information for surface layout in a frozen ambient array, although pictorial thinking about perception tempts us to assume that they are the same. Pictures are artificial displays of information frozen in time, and this fact will be evident when the special kind of visual perception that is mediated by such displays is treated in detail in Part IV.

Natural perspective, as well as artificial perspective, is restricted in scope, being concerned only with a frozen optical structure. This restriction will be removed in what follows.

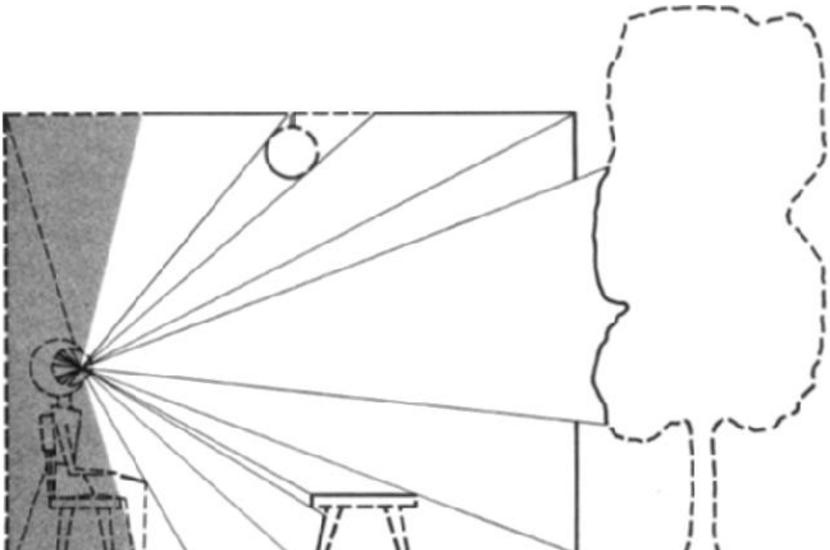

FIGURE 5.3 The same ambient array with the point of observation occupied by a person.

When an observer is present at a point of observation, the visual system begins to function.

Optical Structure with a Moving Point of Observation

A point of observation at rest is only the limiting case of a point of observation in motion, the null case. Observation implies movement, that is, locomotion with reference to the rigid environment, because all observers are animals and all animals are mobile. Plants do not observe but animals do, and plants do not move about but animals do. Hence, the structure of an optic array at a stationary point of observation is only a special case of the structure of an optic array at a moving point of observation. The point of observation normally proceeds along a path of locomotion, and the "forms" of the array change as locomotion proceeds. More particularly, every solid angle included within the array, large or small, is enlarged or reduced or compressed or, in some cases, wiped out. It is wiped out, of course, when its surface goes out of sight.

The optic array *changes,* of course, as the point of observation moves. But it also does *not* change, not completely. Some features of the array do not persist

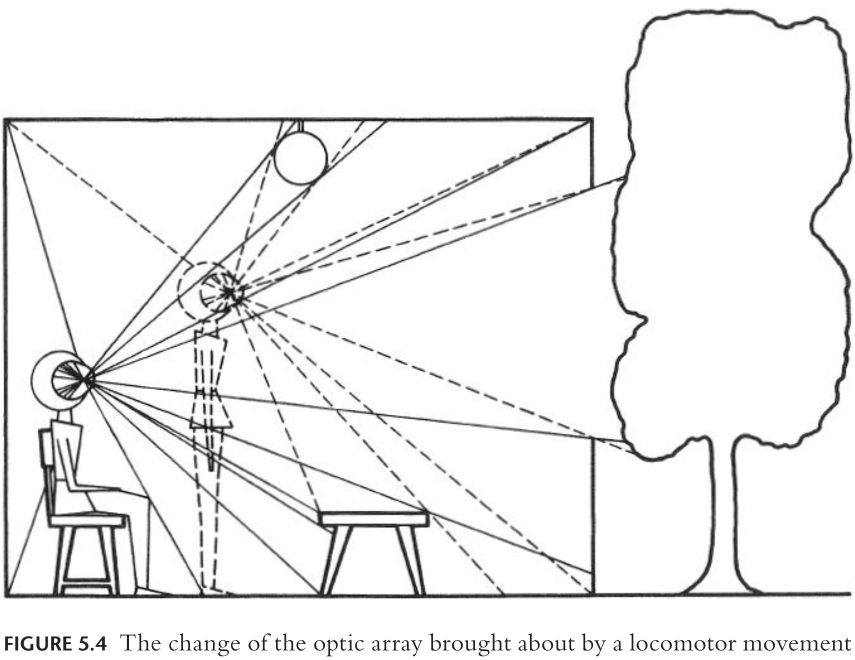

FIGURE 5.4 The change of the optic array brought about by a locomotor movement of the observer.

The thin solid lines indicate the ambient optic array for the seated observer, and the thin dashed lines the altered optic array after standing up and moving forward. The difference between the two arrays is specific to the difference between the points of observation, that is, to the path of locomotion. Note that the whole ambient array is changed, including the portion behind the head. And note that what was previously hidden becomes unhidden.

and some do. The changes come from the locomotion, and the nonchanges come from the rigid layout of the environmental surfaces. Hence, the nonchanges specify the layout and count as information about it; the changes specify locomotion and count as another kind of information, about the locomotion itself. We have to distinguish between two kinds of structure in a normal ambient array, and I shall call them the *perspective structure* and the *invariant structure.*

Perspective Structure and Invariant Structure

The term *structure* is vague, as we have seen. Let us suppose that a kind of essential structure underlies the superficial structure of an array when the point of observation moves. This essential structure consists of what is invariant despite the change. What is invariant does not emerge unequivocally except with a flux. The essentials become evident in the context of changing nonessentials.

Consider the paradox in the following piece of folk wisdom: "The more it changes, the more it is the same thing." Wherein is it true and wherein false? If *change* means *to become different but not to be converted into something else,* the assertion is true, and the saying emphasizes the fact that whatever is invariant is more evident with change than it would be without change. If *change* means *to become different by being converted into something else,* the assertion is self-contradictory, and the paradox arises. But this is not what the word ordinarily means. And assuredly it is not what change in the ambient array means. One arrangement does not become a wholly different arrangement by a displacement of viewpoint. There is no jump from one to another, only a variation of structure that serves to reveal the nonvariation of structure. The pattern of the array does not ordinarily scintillate; the forms of the array do not go from triangular to quadrangular, for example.

There are many invariants of structure, and some of them persist for long paths of locomotion while some persist only for short paths. But what I am calling the *perspective structure* changes with every displacement of the point of observation—the shorter the displacement the smaller the change, and the longer the displacement the greater the change. Assuming that the environment is never reduplicated from place to place, the arrested perspective is unique at each stationary point of observation, that is, for each point of observation there is one and only one arrested perspective. On the other hand, invariants of structure are common to all points of observation—some for all points in the whole terrestrial environment, some only for points within the boundaries of certain locales, and some only for points of observation within (say) a single room. But to repeat, the invariant structure separates off best when the frozen perspective structure begins to flow.

Consider, for example, the age-old question of how a rectangular surface like a tabletop can be given to sight when presumably all that an eye can see is a large number of forms that are trapezoids and only one form that is rectangular, that one being seen only when the eye is positioned on a line perpendicular to the center of the surface. The question has never been answered, but it can be reformulated to ask, What are the invariants underlying the transforming perspectives in the array from the tabletop? What specifies the shape of this rigid surface as projected to a moving point of observation? Although the changing angles and proportions of the set of trapezoidal projections are a fact, the unchanging relations among the four angles and the invariant proportions over the set are another fact, equally important, and they uniquely specify the rectangular surface. There will be experimental evidence about optical transformations as information in Chapter 9.

We tend to think of each member of the set of trapezoidal projections from a rectangular object as being a form in space. A change is then a transition from one form to another, a transformation. But this habit of thought is misleading. Optical change is not a transition from one form to another but a reversible process. The superficial form becomes different, but the underlying form remains the same. The structure changes in some respects and does not change in others. More exactly, it is variant in some respects and invariant in others.

The geometrical habit of separating space from time and imagining sets of frozen forms in space is very strong. One can think of each point of observation in the medium as stationary and distinct. To each such point there would correspond a unique optic array. The set of all points is the space of the medium, and the corresponding set of all optic arrays is the whole of the available information about layout. The set of all line segments in the space specifies all the possible displacements of points of observation in the medium, and the corresponding set of transformation families gives the information that specifies all the possible paths. This is an elegant and abstract way of thinking, modeled on projective geometry. But it does not allow for the complexities of optical change and does not do justice to the fact that the optic array *flows in time* instead of going from one structure to another. What we need for the formulation of ecological optics are not the traditional notions of space and time but the concepts of variance and invariance considered as reciprocal to one another. The notion of a *set* of stationary points of observation in the medium is appropriate for the problem of a whole crowd of observers standing in different positions, each of them perceiving the environment from his own point of view. But even so, the fact that all observers can perceive the same environment depends on the fact that each point of view can move to any other point of view.

REDUPLICATION

It is easy to make copies or duplicates of a picture but the world is never exactly the same in one place as it is in another. Nor is one organism ever exactly the same as another. One cubic yard of empty abstract space is exactly the same as another, but that is a different matter.

The Significance of Changing Perspective in the Ambient Array

When the moving point of observation is understood as the general case, the stationary point of observation is more intelligible. It no longer is conceived as a single geometrical point in space but as a pause in locomotion, as a temporarily fixed position relative to the environment. Accordingly, an arrested perspective structure in the ambient array specifies to an observer such a fixed position, that is, rest; and a flowing perspective structure specifies an unfixed position, that is, locomotion. The optical information for distinguishing locomotion from nonlocomotion is available, and this is extremely valuable for all observers, human or animal. In physics the motion of an observer in space is "relative," inasmuch, as what we call motion with reference to one chosen frame of reference may be nonmotion with reference to another frame of reference. In ecology this does not hold, and the locomotion of an observer in the environment is absolute. The environment is simply that with respect to which either locomotion or a state of rest occurs, and the problem of relativity does not arise.

Locomotion and rest go with flowing and frozen perspective structure in the ambient array; they are what the flow and the nonflow *mean.* They contain information about the potential observer, not information about the environment, as the invariants do. But note that information about a world that surrounds a point of observation implies information about the point of observation that is surrounded by a world. Each kind of information implies the other. Later, in discussing the occupied point of observation, I shall call the former *exterospecific information* and the latter *propriospecific information.*

Not only does flowing perspective structure specify locomotion, but the particular instance of flow specifies the particular path of locomotion. That is, the difference of perspective between the beginning and the end of the optical change is specific to the difference of position between the beginning and the end of the locomotor displacement. But more than that, the *course* of the optical flow is specific to the *route* the path of locomotion takes through the environment. Between one place and another there are many different routes. The two places are specified by their different arrested perspectives, but the different routes between them are in correspondence with different optical sequences between the two perspectives. There will be more of this later. It is enough now to point out that the visual control of locomotion by an observer, purposive locomotion such as homing, migrating, finding one's way, getting from place to place, and being oriented, depends on just the kind of sequential optical information described.

It is important to realize that the flowing perspective structure and the underlying invariant structure are concurrent. They exist at the same time. Although they specify different things, locomotion through a rigid world in the first instance and the layout of that rigid world in the second instance, they are like the two sides of a coin, for each implies the other. This hypothesis, that optical change can seemingly specify two things at the same time, sounds very strange, as if one cause were having two effects or as if one stimulus were arousing two sensations. But there is nothing illogical about the idea of concurrent specification of two reciprocal things. Such an idea is much needed in psychology.

The Change between Hidden and Unhidden Surfaces: Covering Edges

We are now prepared to face a fact that has seemed deeply puzzling, a fact that poses the greatest difficulty for all theories of visual perception based on sensations. The layout of the environment includes unprojected (hidden) surfaces at a point of observation as well as projected surfaces, but observers perceive the layout, not just the projected surfaces. Things are seen *in the round* and one thing is seen *in front* of another. How can this be? Information must be available for the whole layout, not just for its facades, for the covered surfaces as well as the covering surfaces. What is this information? Presumably it becomes evident over time, with changes of the array. I will argue that the information is implicit in the *edges that separate* the surfaces or, rather, in the optical specification of these edges. I am suggesting that if covering edges are specified, both the covered and the covering surfaces are also specified.

To suggest that an observer can see surfaces that are unseen is, of course, a paradox. I do not mean that. I am not saying that one can see the unseen, and I am suspicious of visionaries who claim that they can. A vast amount of mystification in the history of human thought has arisen from this paradox. The suggestion is that one can perceive surfaces that are temporarily *out of sight,* and what it is to be out of sight will be carefully defined. The important fact is that they come into sight and go out of sight as the observer moves, first in one direction and then in the opposite direction. If locomotion is reversible, as it is, whatever goes out of sight as the observer travels comes into sight as the observer returns and conversely. The generality of this principle has never been realized; it applies to the shortest locomotions, in centimeters, as well as to the longest, in kilometers. But it has not been elaborated. I will call it the *principle of reversible occlusion.* The theory of the cues for depth perception includes one cue called "movement parallax" and another called "superposition," both related to the above principle, but these terms are vague and do not even begin to explain what needs to be explained. What we see is not depth as such but *one thing behind another.* The new principle can be made explicit. I will attempt to do so, at some length.

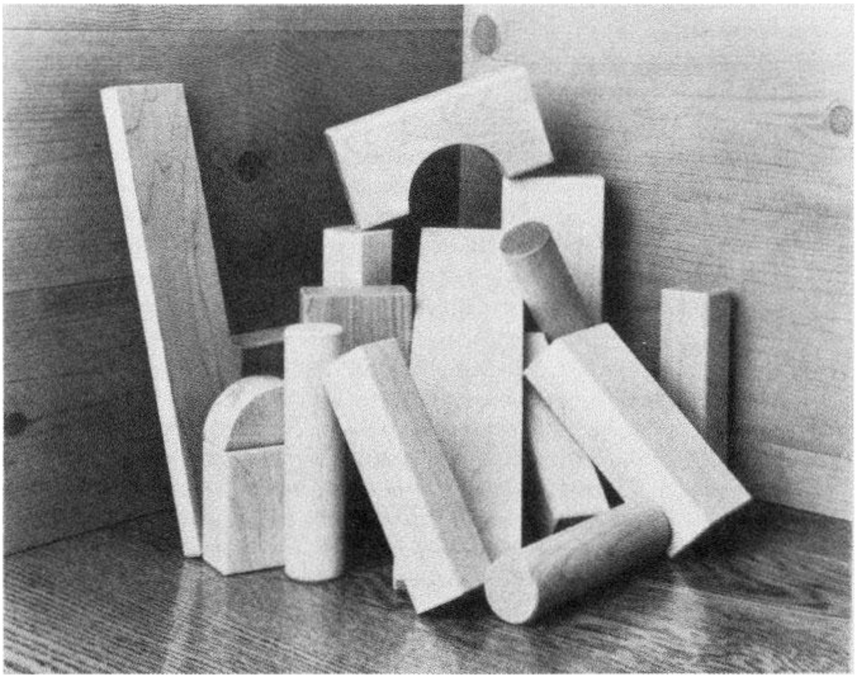

FIGURE 5.5 Objects seen in the round and behind other objects.

Do you perceive covered surfaces as well as covering surfaces in this photograph? (Photo by Jim Scherer.)

Projected and Unprojected Surfaces

There are many commonsense words that refer to the fact of covered and uncovered things. Objects and surfaces are said to be hidden or unhidden, screened or unscreened, concealed or revealed, undisclosed or disclosed. We might borrow a technical word in astronomy, *occultation,* but it means primarily the shutting off of the light from a celestial source, as in an eclipse. We need a word for the cutting off of a visual solid angle, not of light rays. I have chosen the word *occlusion* for it. An occluded surface is one that is out of sight or hidden from view. An occluding edge is the edge of an occluding surface. The term was first introduced in a paper by J. J. Gibson, G. A. Kaplan, H. N. Reynolds, and K. Wheeler (1969) on the various ways in which a thing can pass between the state of being visible and the state of being invisible. The experiment will be described in Chapter 11.

Occlusion arises because of two facts about the environment, both described in Chapter 2. First, surfaces are generally opaque; and second, the basic environment, the earth, is generally cluttered. As to the first, if surfaces were as transparent as air, they would not reflect light at all and there would be no use for vision. Most substances are nontransmitting (they reflect and absorb instead), and therefore light is reflected back from the surface. A few substances are partially transmitting or "translucent," and hence a sheet of such a substance will transmit part of the radiant light but will not transmit the structure of the ambient array; it will let through photons but not visual solid angles. There can be an obstructing of the *view* without obstructing of the *light,* although an obstructing of the light will of course also obstruct the view. If we add the fact that surfaces are also generally textured, the facts of opaque surfaces as contrasted with the surfaces of semitransparent and translucent substances become intelligible.

The second fact is that the environment is generally cluttered. What I called an open environment is seldom or never realized, although it is the only case in which all surfaces are projected and none are unprojected. An open environment has what we call an unobstructed view. But the flat and level earth receding unbroken to a pure linear horizon in a great circle, with a cloudless sky, would be a desolate environment indeed. Perhaps it would not be quite as lifeless as geometrical space, but almost. The *furniture* of the earth, like the furnishings of a room, is what makes it livable. The earth as such affords only standing and walking; the furniture of the earth affords all the rest of behavior. The main items of the clutter (following the terminology adopted in Chapter 3) are *objects,* both attached and detached, *enclosures, convexities* such as hills, *concavities* such as holes, and *apertures* such as windows. These features of surface layout give rise to occluding surfaces or, more exactly, to the *separation* of occluding and occluded surfaces.

A surface is *projected* at a point of observation if it has a visual solid angle in the ambient optic array; it is *unprojected* if it does not. A projected surface may become unprotected in at least three ways—if its solid angle is diminished to a point, if the solid angle is compressed to a line, or if the solid angle is wiped out. In the first case we say that the surface is too far away, in the second that it is turned so as not to face the point of observation, in the third that the view is obstructed. The second case, that of facing toward or away, is instructive. A wall or a sheet of paper has two "faces" but only one can *face* a fixed point. The relation between the occluding and occluded surfaces is given by the relation of each to the point; the relation is not merely geometrical but also optical. The relation is designated when we distinguish between the *near* side and the *far* side of an object. (It is not, however, well expressed by the terms *front* and *back,* since they are ambiguous. They can refer to such surfaces as the front and the back of a house or the front and the back of a head. Terms can be borrowed from ordinary language only with discretion!)

Going Out of and Coming Into Sight

A point of observation is to be thought of as moving through the medium to and fro, back and forth, often along old paths but sometimes along new ones.

Displacements of this position are reversible and are reversed as its occupier comes and goes, even as she slightly shifts her posture. Any face or facet, any surface of the layout, that is progressively hidden during a displacement is progressively unhidden during its reversal. Going out of sight is the inverse of coming into sight. Hence, occluding and occluded surfaces interchange. The occluding ones *change into* the occluded ones and vice versa, not by changing from one entity to another but by a special transition.

The terms *disappearance* and its opposite, *appearance,* should not be used for this transition. They have slippery meanings, like *visible* and *invisible.* For a surface may disappear by going out of existence as well as by going out of sight, and the two cases are profoundly different. A surface that disappears because it is no longer projected to *any* point of observation, because it has evaporated, for example, should not be confused with a surface that disappears because it is no longer projected to a fixed point of observation. The latter can be seen from another position; the former cannot be seen from any position. Failure to distinguish these meanings of *disappear* is common; it encourages careless observation and vague beliefs in ghosts, or in the reality of the "unseen." To *disappear* can also refer to a surface that continues to exist but is no longer projected to any point of observation because of darkness. Or we might speak of something disappearing "in the distance," referring to a surface barely projected to a point of observation because its visual solid angle has diminished to a limit. These modes of so-called disappearance are quite radically different. The differences between (1) a surface that ceases to exist, (2) a surface that is no longer illuminated, (3) a surface that lies on the horizon, and (4) a surface that is occluded are described in a paper by Gibson, Kaplan, Reynolds, and Wheeler (1969) and are illustrated in a motion picture film (Gibson, 1968*a*). An experimental study of the perception of occlusion using motion picture displays has been reported by Kaplan (1969).

The Loci of Occlusion: Occluding Edges

We must now distinguish an edge that is simply the junction of two surfaces from an edge that causes one surface to hide another, an *occluding edge.* In the proposed terminology of layout in Chapter 3, I defined an *edge* as the apex of a convex dihedral (as distinguished from a *corner,* which is the apex of a concave dihedral). But an occluding edge is a dihedral where only one of the surfaces is projected to the point of observation—an *apical* occluding edge. I also defined a *curved convexity* (as distinguished from a *curved concavity*), and another kind of occluding edge is the *brow* of this convexity, that is, the line of tangency of the envelope of the visual solid angle—a *curved* occluding edge. The apical occluding edge is "sharp," and the curved occluding edge is "rounded." The two are illustrated in Figure 5.6. The latter slides along the surface as the point of observation moves, but the former does not. Note that an occluding edge always requires a convexity of some sort, a protrusion of the substance into the medium.

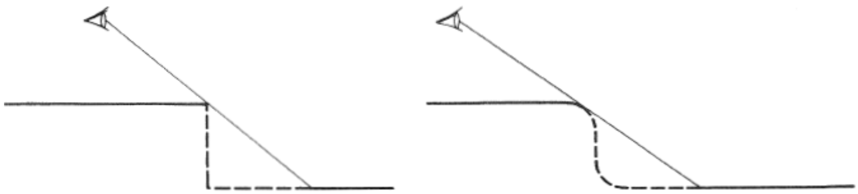

FIGURE 5.6 The sharp occluding edge and the rounded occluding edge at a fixed point of observation.

The hidden portions of the surface layout are indicated by dotted lines.

These two kinds of occluding edges are found in the ells of corridors, the brinks of cliffs, the brows of hills, and the near sides of holes in the ground. One face or facet or part of the layout hides another to which it may be connected and which it may adjoin. This is different from what I called a detached object, by which I mean the movable or moving object having a topologically closed surface with substance inside and medium outside. The detached object produces a visual solid angle in the optic array, as noted by Euclid and Ptolemy, and yields a closed-contour figure in the visual field, as described by Edgar Rubin and celebrated by the gestalt psychologists under the name of the "figure-ground phenomenon." Occluding edges are a special case, because not only does the near side of the object hide the far side but the object covers a sector of the surface behind it, the ground, for example. The occluding edges may be apical, as when the object is a polyhedron, or the locus of the tangent of the envelope of the solid angle to the surface, as when the object is curved. These are illustrated in Figure 5.7, where both the hiding of the far side and the covering of the background are shown. The object is itself rounded or

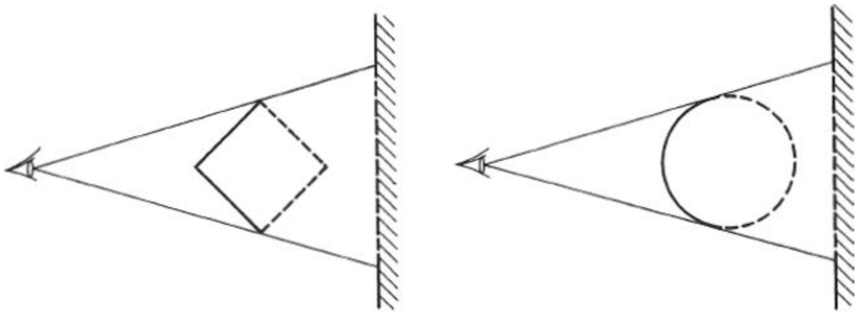

FIGURE 5.7 Both the far side of an object and the background of the object are hidden by its occluding edges.

Two detached objects are shown, one with sharp occluding edges and the other with rounded occluding edges.

solid, and it is superposed on the ground, which is also continuous behind the object. These two kinds of occlusion may be treated separately.

Self-occlusion and Superposition

An object, in the present terminology, is both voluminous and superposed. It exists in volume and it may lie in front of another surface, or another object. In short, an object always occludes itself and generally also occludes something else. The effect of a moving point of observation is different in the two cases.

Projected and unprojected surfaces interchange as the point of observation moves, but the interchange between parts of the object is not like that between parts of the background. There is an interchange between *opposite faces* of the object but an interchange of *adjacent areas* of the surface behind the object. For the object, the near side turns into the far side and vice versa, whereas for the background an uncovered area becomes covered and vice versa. The change of optical structure in the former case is by way of perspective transformation, whereas the disturbance of optical structure in the latter case is more radical, a "kinetic disruption" being involved.

In Figure 5.7, as the point of observation moves each face of the facade of the polyhedron undergoes transformation, for example, from trapezoid to square to trapezoid. Ultimately, when the face is maximally foreshortened, it is what we call "edge on," that is, it becomes an occluding edge. The near face turns into a far face by way of the edge. While this is happening at one edge, the other edge is revealing a previously hidden face. A far face turns into a near face. The two occluding edges in the diagram are perfectly reciprocal; while one is converting near into far, the other is converting far into near. The width of the polyhedron goes into depth, and the depth comes back into width. Width and depth are thus interchangeable.

Similarly, one could describe the transformation of each facet of the textured surface of the curved object. If the object is a sphere, the circular occluding edge (the *outline,* in pictorial terminology) does not transform, but the optical structure within it does. At one edge the texture is progressively turning from projected into unprojected, from near into far, while at the other edge the texture is progressively turning from unprojected into projected, from far into near. The transition occurs at the limit of the slant transformation, the ultimate of perspective foreshortening, but actually the optical texture reaches and goes beyond this purported limit. It has to go beyond it because it comes from beyond that limit at the other occluding edge.

Superposition

Now consider the separated background behind the objects in Figure 5.7, the fact of superposition as distinguished from the fact of solidity. As the point of observation moves, the envelope of the visual solid angle sweeps across the surface. The leading edge progressively covers the texture of the surface, while the trailing edge progressively uncovers it. I have suggested metaphorically that the texture is "wiped out" and "unwiped" at the lateral borders of the figure (Gibson, 1966*b*, pp. 199 ff.). This was inspired by the metaphors used by A. Michotte in describing experiments on what he called the "tunnel effect" (Michotte, Thinès, and Crabbé, 1964). A somewhat more exact description of this optical change will be given below. But note that if the texture that is progressively covered has the same structure as the texture that is progressively uncovered the unity of the surface is well specified.

The metaphor of "wiping" is inexact. A better description of the optical transition was given by Gibson, Kaplan, Reynolds, and Wheeler (1969), and it was also described by Kaplan (1969) as a "kinetic disruption." There is a disturbance of the structure of the array that is not a transformation, not even a transformation that passes through its vanishing limit, but a breaking of its adjacent order. More exactly, there is either a progressive decrementing of components of structure, called *deletion,* or its opposite, a progressive incrementing of components of structure, called *accretion.* An edge that is covering the background deletes from the array; an edge that is uncovering the background accretes to it. There is no such disruption for the surface that is covering or uncovering, only for the surface that is being covered or uncovered. And nondisruption, I suggest, is a kind of invariance.

The Information to Specify the Continuation of Surfaces

A surface always "bends under" an occluding edge, and another surface generally "extends behind" it. These surfaces are connected or continuous. Is there information in a changing optic array to specify the connnectedness or continuity?

Here is a tentative hypothesis for the continuous object surface:

> Whenever a perspective transformation of form or texture in the optic array goes to its limit and when a series of forms or textures are progressively foreshortened to this limit, a continuation of the surface of an object is specified at an occluding edge. This is the formula for going out of sight; the formula is reversed for coming into sight.

Here is a tentative hypothesis for the continuous background surface:

> Whenever there occurs a regular disturbance of the persistence of forms and textures in the optic array such that they are progressively deleted at a contour, the continuation of the surface of a ground is specified at an occluding edge. This is for going out of sight; substituting accretion for deletion gives the formula for coming into sight.

These two hypotheses make no assertions about perception, only about the information that is normally available for perception. They do not refer to *space,* or to the *third dimension,* or to *depth,* or to *distance.* Nothing is said about forms or patterns in two dimensions. But they suggest a radically new basis for explaining the perception of solid superposed objects, a new theory based not on cues or clues or signs but on the direct pickup of solidity and superposition. An object is in fact voluminous; a background is in fact continuous. A picture or an image of an object is irrelevant to the question of how it is perceived.

The assumption for centuries has been that the sensory basis for the perception of an object is the outline form of its image on the retina. Object perception can only be based on form perception. First the silhouette is detected and then the depth is added, presumably because of past experiences with the cues for depth. But the fact is that the progressive foreshortening of the face of an object is perceived as the turning of the object, which is precisely what the transformation specifies, and is never perceived as a change of form, which ought to be seen if the traditional assumption is correct—that the silhouette is detected and then the depth is added.

The two hypotheses stated above depend on a changing optic aray, and so far the only cause of such change that has been considered is the moving point of observation. The reader will have noted that a moving *object* will also bring about the same kinds of disturbance in the structure of the array that have been described above. A moving object in the world is an event, however, not a form of locomotion, and the information for the perception of events will be treated in Chapter 6.

The Case of Very Distant Surfaces

It is interesting to compare the occluding edges of objects and other convexities on the surface of the earth with the *horizon* of the earth, the great circle dividing the ambient array into two hemispheres. It is the limit of perspective minification for terrestrial surfaces, just as the edge-on line is the limit of perspective compression (foreshortening) for a terrestrial surface. Objects such as railroad trains on the Great Plains and ships on the ocean are said to vanish in the distance as they move away from a fixed point of observation. The line of the horizon in the technology of pictorial perspective is said to be the locus of vanishing points for the size of earth-forms and for the convergence of parallel edges on the earth. The railroad train "vanishes" at the same optical point where the railroad tracks "meet" in the distance. The horizon is therefore analogous to an occluding edge *in* being one of the loci at which things go out of and come into sight. But going out of sight in the distance is very different from going out of sight at a sharp or a rounded edge nearby. The horizon of the earth, therefore, is not an occluding edge for any terrestrial object or earth-form. It does not in fact *look* like an occluding edge. It could only be visualized as an occluding edge for the lands

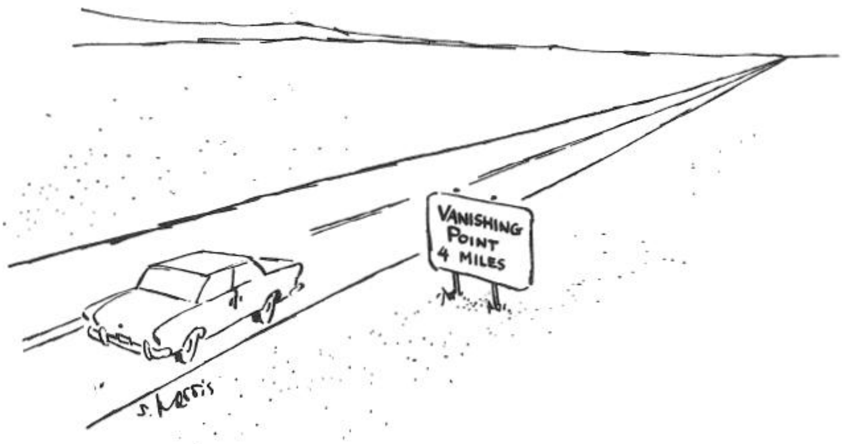

FIGURE 5.8 Cartoon. (Drawing by S. Harris; © 1975 The New Yorker Magazine, Inc.)

and seas beyond the horizon if the seemingly flat earth were conceived as curved and if the environment were thought of as a globe too vast to see.

It has long been a puzzle to human observers, however, that the horizon is in fact visibly an occluding edge for *celestial* objects such as the sun and the moon. Such objects undergo progressive deletion at a contour, as at sunset, and undergo progressive accretion at the same contour, as at moonrise. This is in accordance with the second hypothesis above. The object is obviously beyond the horizon, more distant than the visible limit of earthly distance, and yet there is some information for its being a solid surface. This conflicting information explains, I think, the apparently enormous size of the sun and the moon at the horizon. It also explains many of the ideas of pre-Copernican astronomy about heavenly bodies. We should realize that the terrestrial environment was the only environment that people could be *sure* of before Copernicus—the only environment that could be perceived directly. Terrestrial objects and surfaces had affordances for behavior, but celestial objects did not. More will be said about the perception of objects on earth as distinguished from objects in the sky in Part III.

Summary: The Optics of Occlusion

1. In the ideal case of a terrestrial earth without clutter, all parts of the surface are projected to all points of observation. But such an open environment would hardly afford life.

2. In the case of an earth with furniture, with a layout of opaque surfaces on a substratum, some parts of the layout are projected to any given fixed point of observation and the remaining parts are unprojected to that point.

3. The optically uncovered surface of an object is always *separated* from the optically covered surface at the occluding edge. At the same time, it is always *connected* with the optically covered surface at the occluding edge.

4. The continuation of the far side with the near side is specified by the *reversibility* of occlusion.

5. Any surface of the layout that is hidden at a given fixed point of observation will be unhidden at some other fixed point.

6. Hidden and unhidden surfaces *interchange.* Whatever is revealed by a given movement is concealed by the reverse of that movement. This principle of reversible occlusion holds true for both movements of the point of observation and motions of detached objects.

7. We can now observe that the *separation* between hidden and unhidden surfaces at occluding edges is best specified by the *perspective* structure of an array, whereas the *connection* between hidden and unhidden surfaces at edges is specified by the underlying *invariant* structure. Hence, probably, a pause in locomotion calls attention to the difference between the hidden and the unhidden, whereas locomotion makes evident the continuousness between the hidden and the unhidden.

The seeming paradox of the perceiving or apprehending of hidden surfaces will be treated further in Chapter 11.

How is Ambient Light Structured? A Theory

Let us return to the question of how ambient light is given its invariant structure, the question asked at the beginning of this chapter but not answered except in a preliminary way. Ambient light can only be structured by something that surrounds the point of observation, that is, by an environment. It is not structured by an empty medium of air or by a fog-filled medium. There have to be surfaces—both those that emit light and those that reflect light. Only because ambient light is structured by the substantial environment can it contain information about it.

So far it has been emphasized that ambient light is made to constitute an array by a single feature of these surfaces, their layout. But just *how* does the layout structure the light? The answer is not simple. It involves the puzzling complexities of light and shade. Moreover, the layout of surfaces is not the only cause of the structuring of light; the *conglomeration* of surfaces makes a contribution, that is, the fact that the environment is multicolored. The different surfaces of the layout are made of different substances with different reflectances. Both lighted or shaded surfaces and black or white surfaces make their separate contributions to the invariant structure of ambient light. And how light-or-shade can be perceived separately from black-or-white has long been a puzzling problem for any theory of visual sense perception.

I tried to formulate a theory of the structuring of ambient light in my last book (Gibson, 1966*b*), asserting that three causes existed, the layout of surfaces, the pigmentation of surfaces, and the shadowing of surfaces (pp. 208–216). But the third of these causes is not cognate with the other two, and the interaction between them was not clearly explained. The theory was static. Here, I shall formulate a theory of the sources of *invariant* optical structure in relation to the sources of *variation* in optical structure. What is clear to me now that was not clear before is that structure as such, frozen structure, is a myth, or at least a limiting case. Invariants of structure do not exist except in relation to variants.

The Sources of Invariant Optical Structure

The main invariants of the terrestrial environment, its persisting features, are the layout of its surfaces and the reflectances of these surfaces. The layout tends to persist because most of the substances are sufficiently solid that their surfaces are rigid and resist deformation. The reflectances tend to persist because most of the substances are chemically inert at their interfaces with the air, and their surfaces keep the same composition, that is, the same colors, both achromatic and chromatic. Actually, at the level of microlayout (texture) and microcomposition (conglomeration), layout and reflectances merge. Or, to put it differently, the layout texture and the pigment texture become inseparable.

Note once more that an emphasis on the geometry of surfaces is abstract and oversimplified. The faces of the world are *not* made of some amorphous, colorless, ghostly substance, as geometry would lead us to believe, but are made of mud or sand, wood or metal, fur or feathers, skin or fabric. The faces of the world are colorful as well as geometrical. And what they afford depends on their substance as well as their shape.

The Sources of Variant Optical Structure

There are two regular and recurrent sources of changing structure in the ambient light (apart from local events, which will be considered in the next chapter). First, there are the changes caused by a moving point of observation, and second, there are the changes caused by a moving source of illumination, usually the sun. Many pages have been devoted to the former, and we must now consider the latter. The motion of the sun across the sky from sunrise to sunset has been for countless millions of years a basic regularity of nature. It is a fact of ecological optics and a condition of the evolution of eyes in terrestrial animals. But its importance for the theory of vision has not been fully recognized.

The puzzling complexities of light and shade cannot be understood without taking into account the fact of a *moving* source of illumination. For whenever the source of light moves, the direction of the light falling on the surfaces of the world is altered and the shadows themselves move. The layout and coloration of surfaces persist, but the lightedness and shadedness of these surfaces do not. It is not just that the optic array is different at noon with high illumination from what it is at twilight with low illumination; it is that the optic array has a different structure in the afternoon than it has in the morning.

Variants and Invariants with a Moving Source of Illumination

Just how does pure layout structure the ambient light? It is easy to understand how a mosaic of black and white substances would structure the ambient light but not how a pure layout would do so. For in this case the structuring would have to be achieved wholly by differential illumination, by light and shade. There are two principles of light and shade under natural conditions that seem to be clear: the direction of the prevailing illumination and the progressive weakening of illumination with multiple reflection.

The illumination on a surface comes from the sun, the sky, and other surfaces that face the surface in question. A surface that faces the sun is illuminated "directly," a surface that faces away from the sun but still faces the sky is illuminated less directly, and a surface within a semienclosure that faces only other surfaces is illuminated still less directly. The more the light has reverberated, the more of it is absorbed and the dimmer it becomes. Hence it is that surfaces far from the mouth of a cave are more weakly illuminated than those near the mouth. But within any airspace, any concavity of the terrain or any semienclosure, there is a direction of the *prevailing* illumination, that is, a direction from which more light comes than from any other.

The illumination of any face of the layout relative to adjacent faces depends on its inclination to the prevailing illumination. Crudely speaking, the surface that "faces the light" gets more than its neighbor. More exactly, a surface perpendicular to the prevailing illumination gets the most, a surface inclined to it gets less, a surface parallel to it gets still less, and a surface inclined *away* from it gets the least. The pairs of terms *lighted* and *shadowed* or *in light* and *in shadow* should not be taken as dichotomies, for there are all gradations of relative light and shade. These two principles of the direction and the amount of illumination are an attempt to distill a certain ecological simplicity from the enormous complexities of analytical physical optics and the muddled practice of illumination engineering.

A wrinkled surface of the same substance evidently structures the ambient light by virtue of two facts: there is always a prevailing direction of illumination, and consequently the slopes facing in this direction throw back more energy than the slopes not facing in this direction. A *flat* surface of *different* substances structures the ambient light by virtue of the simple fact that the parts of high reflectance throw back more energy than the parts of low reflectance.

Figure 5.9 shows an array from a wrinkled layout of terrestrial surfaces, actually an aerial photograph of barren hills and valleys. The bare earth of this desert has everywhere the same reflectance. The top of the photograph is to the north of the terrain. The picture was taken in the morning, and the sun is in the east. Some of the slopes face east, and some face west; the former are lighted and the latter shaded. It can be observed that various inclinations of these surfaces to the direction of the prevailing illumination determine various relative intensities in the array; the more a surface departs from the perpendicular to this direction, the darker is the corresponding patch in the optic array.

Now consider what happens as the sun moves across the sky. All those surfaces that were lighted in the morning will be shaded in the afternoon, and all those that were shaded in the morning will be lighted in the afternoon. There is a continual, if slow, process of change from lighted to shaded on certain slopes of the layout and the reverse change on certain other slopes. These slopes are related by orientation. Two faces of any convexity are related in this way, as are two faces of any concavity. A ridge can be said to consist of two opposite slopes, and so can a valley. The reciprocity of light and shade on such surfaces might be described by saying that the lightness and the shadedness *exchange places*. The underlying surfaces do not interchange of course, and their colors, if any, do not interchange. They are persistent, but the illumination is variable in this special reciprocal way.

In the optic array, presumably, there is an underlying invariant structure to specify the edges and corners of the layout and the colors of the surfaces, and at the same time there is a changing structure to specify the temporary direction of the prevailing illumination. Some components of the array never exchange places—that is, they are never permuted—whereas other components of the array do. The former specify a solid surface; the latter specify insubstantial shadows only. The surface and its color are described as opaque; the shadow is described as transparent.

The decreasing of illumination on one slope and the increasing of illumination on an adjacent slope as the sun moves are analogous to the foreshortening of one slope along with the inverse foreshortening of an adjacent slope as the point of observation moves. I suggest that the true relative colors of the adjacent surfaces emerge as the lighting changes, just as the true relative shapes of the adjacent surfaces emerge as the perspective changes. The perspectives of the convexities and concavities of Figure 5.9 are variant with locomotion; the shadows of these convexities and concavities are variant with time of day; the constant properties of these surfaces underlie the changing perspectives and the changing shadows and are specified by invariants in the optic array.

It is true that the travel of the sun across the sky is very slow and that the correlated interchange of the light and the shade on surfaces is a very gradual fluctuation. Neither is as obvious as the motion perspective caused by loco-

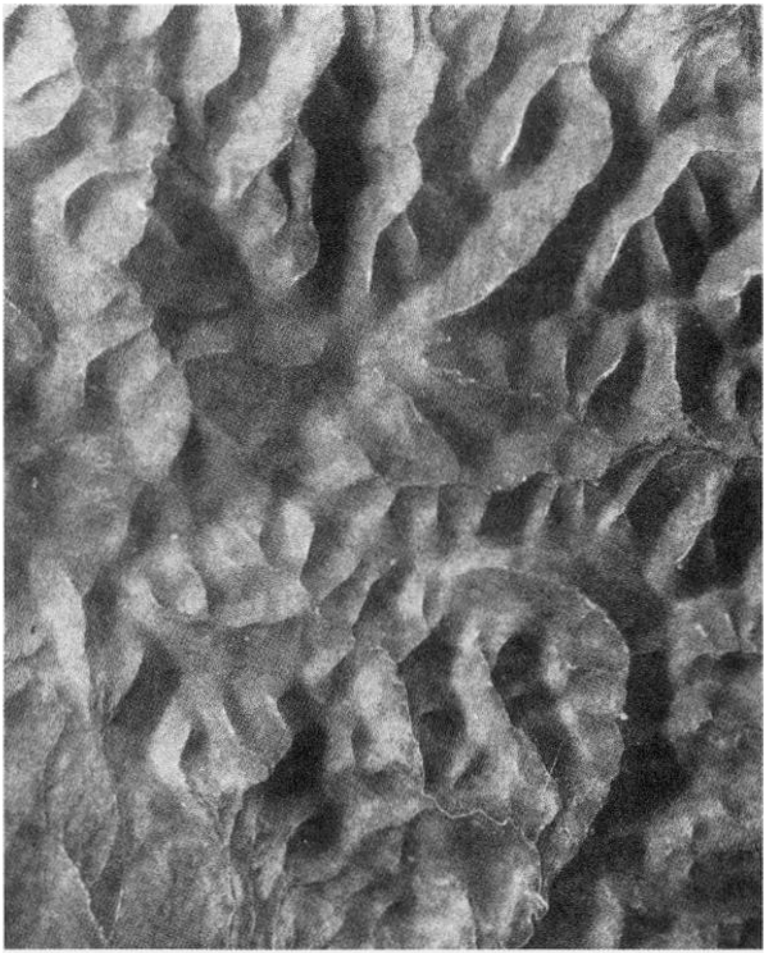

FIGURE 5.9 Hills and valleys on the surface of the barren earth.

The hills in this aerial photograph, the convexities or protuberances, can be compared to the "humps" shown in Figure 5.1

motion. But the fact is that shifting shadows and a moving sun are regularities of ecological optics whether or not they are ever noticed by any animal. They have set the conditions for the perception of the terrain by terrestrial animals since life emerged from the sea. They make certain optical information available. And, although shifting shadows and a moving sun are too slow to be noticed in daylight, a moving source of illumination and the resultant shadows become more obvious at night. One has only to carry a light from place to place in a cluttered environment in order to notice the radical shifts in the pattern of the optic array caused by visibly moving shadows. And yet, of course, the layout of surfaces and their relative coloration is visible underneath the moving shadows.

How the differential colors of surfaces are specified in the optic array separately from the differential illumination of surfaces is, of course, a great puzzle. The difference between black and white is never confused with the difference between lighted and shadowed, at least not in a natural environment as distinguished from a controlled laboratory display. There are many theories of this so-called constancy of colors in perception, but none of them is convincing. A new approach to the problem is suggested by the above considerations.

From an ecological point of view, the color of a surface is relative to the colors of adjacent surfaces; it is not an absolute color. Its reflectance ratio is specified only in relation to other reflectance ratios of the layout. For the natural environment is an *aggregate* of substances. Even a surface is sometimes a *conglomerate* of substances. This means that a range of black, gray, and white surfaces and a range of chromatically colored surfaces will be projected as solid angles in a normal optic array. The colors are not seen separately, as stimuli, but together, as an arrangement. And this range of colors provides an invariant structure that underlies *both* the changing shadow structure with a moving sun and the changing perspective structure with a moving observer. The edges and corners, the convexities and concavities, are thus specified as multicolored surfaces, not as mere slopes; as speckled or grained or piebald or whatever, not as ghostly gray shapes.

The experimental discoveries of E. H. Land (1959) concerning color perception with what he calls a "complete image" as distinguished from color perception with controlled patches of radiation in a laboratory are to be understood in the above way, I believe.

Ripples and Waves on Water: A Special Case

It is interesting and revealing to compare the optical information for a solid wrinkled surface as shown in Figure 5.9 and the information for a liquid wavy surface, which the reader will have to visualize. Both consist of convexities and concavities, but they are motionless on the solid surface and moving on the liquid surface. In both cases the convexities are lighted on one slope and shadowed on the other. In both cases the surface is all of the same color or reflectance. The difference between the two arrays is to be found chiefly in the two *forms of fluctuation* of light and shade. In the terrestrial array, light and shade exchange places slowly in one direction; they do not oscillate. In the aquatic array, light and shade interchange rapidly in both directions; they oscillate. In fact, when the sun is out and the ripples act as mirrors, the reflection of the sun can be said to flicker or flash on and off. This specific form of fluctuation is characteristic of a water surface.

Summary

When ambient light at a point of observation is structured it is an ambient optic array. The point of observation may be stationary or moving, relative to the persisting environment. The point of observation may be unoccupied or occupied by an observer.

The structure of an ambient array can be described in terms of visual solid angles with a common apex at the point of observation. They are angles of intercept, that is, they are determined by the persisting environment. And they are nested, like the components of the environment itself.

The concept of the visual solid angle comes from natural perspective, which is the same as ancient optics. No two such visual angles are identical. The solid angles of an array change as the point of observation moves, that is, the perspective structure changes. Underlying the perspective structure, however, is an invariant structure that does *not* change. Similarly, the solid angles of an array change as the sun in the sky moves, that is, the shadow structure changes. But there are also invariants that underlie the changing shadows.

The moving observer and the moving sun are conditions under which terrestrial vision has evolved for millions of years. But the invariant principle of reversible occlusion holds for the moving observer, and a similar principle of reversible illumination holds for the moving sun. Whatever goes out of sight will come into sight, and whatever is lighted will be shaded.

11

THE DISCOVERY OF THE OCCLUDING EDGE AND ITS IMPLICATIONS FOR PERCEPTION

The facts of occlusion have been described in Chapter 5. They are part of ecological optics. But they were not recognized as facts until observations and experiments made them compelling. The experiments described in the last two chapters about surfaces, layout, change, and kinesthesis were radical enough, but they culminated in the most radical of all, in what I can only call the *discovery of the occluding edge.* This discovery is radical for the following reason. If it is true that there are places where opaque surfaces are seen one behind another, if it is true that one can perceive a *hidden* surface, a paradox arises. For we are not now allowed to say that a hidden surface is *perceived;* we can only say that it is *remembered.* To be perceived, a thing must be "present to the senses"; it must be stimulating receptors. If it is not, it can only be experienced by means of an *image;* it can be recalled, imagined, conceived, or perhaps known, but not perceived. Such is the accepted doctrine, the theory of sensation-based perception. If an occluded surface is perceived, the doctrine is upset.

Kaplan's Experiment

The crucial experiment, which was performed by G. A. Kaplan (1969), involved kinetic, not static, displays of information. Each display was a motion picture shot of a random texture filling the screen, with a progressive deletion (or accretion) of the optical structure on one side of a contour and preservation of the structure on the other side. Photographs of a randomly textured paper were taken frame by frame, and successive frames were modified by careful paper-cutting. No contour was ever visible on any single frame, but progressive decrements of the texture were produced on one side of the invisible line by cutting off thin slices of paper in succession. Progressive increments of the texture could be obtained by reversing the film. This particular kind of decrementing or incrementing of structure had not previously been achieved in a visual display.

In effect, a reversible disturbance of structure in a sample of the optic array had been isolated and controlled, a reversible transition. It is called a *transition,* not a *transformation,* since elements of structure were lost or gained and one-to-one correspondence was not preserved. What was perceived?

All observers, without exception, saw one surface *going behind* another (or *coming from* behind another) that was always concealing (or revealing) the first. Deletion always caused the perception of covering, and accretion always caused the perception of uncovering. The surface going out of sight was never seen to go out of existence, and the surface coming into sight was never seen to come into existence. In short, one surface was seen in a legitimate sense *behind* another *at an occluding edge.*

When the array was arrested by stopping the film, the edge perception ceased and a wholly continuous surface replaced it; when the optical transition was resumed, the edge perception began. The "motion" of the display as such, however, had nothing to do with the occluding edge; what counted was accretion or deletion and whether it was on one side or the other.

These results were striking. There were no uncertainties of judgment, no guessing as in the usual psychophysical experiment. What the observers saw was an edge, a *cut* edge, the edge of a *sheet,* and another surface behind it. But this depended on an array changing in time.

The surface that was being covered was seen to persist after being concealed, and the surface that was being uncovered was seen to pre-exist before being revealed. The hidden surface could not be described as remembered in one case or expected in the other. A better description would be that it was perceived retrospectively and prospectively. It is certainly reasonable to describe perception as extending into the past and the future, but note that to do so violates the accepted doctrine that perception is *confined* to the present.

The crucial paper by Kaplan (1969) was published along with a motion picture film called *The Change from Visible to Invisible: A Study of Optical Transitions* (Gibson, 1968) and an article having the same title by Gibson, Kaplan, Reynolds, and Wheeler (1969). A sharp distinction was made between *going out of sight* and *going out of existence,* and it was proposed that there is information to specify the two cases. I have described the information in Chapters 5 and 6. The former is a *reversing* transition, but the latter is not.

Anticipations of the Occluding Edge

The important result of Kaplan's experiment was not the perceiving of depth at the occluding edge but the perceiving of the persistence of the occluded surface. Depth perception requires no departure from traditional theories, but persistence perception is radically inconsistent with them. Only in the experimental work of

Michotte had anything like persistence perception ever been hinted at (Michotte, Thinès, and Crabbé, 1964). He discovered what he called the "tunnel phenomenon" or the "tunnel effect," the perception of a moving object during the interval between going into a tunnel and coming out of it. He ascribed it, however, not to progressive deletion and accretion of structure for going in and coming out but to a tendency for perception to be completed across a gap, in the style of gestalt theorizing. He did not realize how universal occlusion is during locomotion of the observer. But he was very much aware of the paradox of asserting that an object could be seen during an interval when there was no sensory basis for seeing it. The "screening" or "covering" of an object, he realized, was a fact of visual perception. But he could only suppose that the perception of an object must somehow persist after the sensory input ends; he did not entertain the more radical hypothesis that the persistence of the object is perceived as a fact in its own right. There is a vast difference between the persistence of a percept and the perception of persistence.

It had long been recognized that in pictures, or other displays with a frozen array, the appearance of *superposition* could be obtained. Likewise, Rubin's discovery that a closed contour or *figure* in a display involved the appearance of a *ground* that seemed to extend without interruption behind the figure was well known. But these demonstrations were concerned with the seeing of contours and lines and the perceiving of forms, not with the perceiving of the occluding edges of surfaces in a cluttered terrestrial environment. They showed that what might be called depth-by-superposition could be induced by a picture but not that an occluded surface is seen to persist.

The occluding edge seems to have escaped notice in both physics and psychology. In truth, it is not a fact of physics or a fact of psychology as these disciplines have been taught. It depends on the combined facts of a surface layout and a point of observation.

The Theory of Reversible Occlusion

The theory of reversible occlusion was formulated in Chapter 5 in terms of what I called projected and unprojected surfaces for an ambient optic array at a given time. Reversible occlusion was said to be a consequence of the reversibility of locomotions and motions in the medium, and this was contrasted in Chapter 6 with the unreversibility of changes such as disintegration, dissolution, and the change from a solid to a liquid or a gas. These changes, I said, were not such that the waning of a surface was the temporal inverse of waxing, not such that if a film of one event were *run backward* it would represent the opposite event (Gibson and Kaushall, 1973).

Then, in Chapter 7 on the self, the principle of reversible occlusion was extended to the head turning of the observer, and the margins of the field of view were compared to the occluding edges of a window. The principle is widely applicable. It would be useful to bring together all this theorizing and to summarize it in a list of propositions.

Terminology

The reader should be reminded again that many pairs of terms can be used to denote what I have called *occlusion.* In what follows, the words *hidden* and *unhidden* are chosen to have a general meaning (although they have the unwanted flavor of buried treasure!). *Unprojected* and *projected*, the terms used in Chapter 5, are all right except for the implication of throwing an image on a screen, which gives precisely the wrong emphasis. *Covered* and *uncovered* are possible terms, or *screened* and *unscreened,* and these were employed by Michotte. Other possibilities are *concealed* and *revealed,* or *undisclosed* and *disclosed.* All these terms refer to various kinds of occlusion. The most general terms are *out of sight* and *in sight,* which contrast with *out of existence* and *in existence.* It should be kept in mind that all these terms refer to reversible transitions, that is, to *becoming* hidden or unhidden, to *going* out of sight or *coming* into sight. Terms that should *not* be employed are *disappear* and *appear.* Although in common use, these words are ambiguous and promote sloppy thinking about the psychology of perception. The same is true of the words *visible* and *invisible.*

There seem to be a number of different ways of going out of sight, some not by occlusion and some by occlusion. The latter always involves an occluding edge with progressive deletion on one side of a contour, but the former does not. I can think of three kinds of going out of sight *not* by occlusion: first, going into the distance by minification of the solid angle to a so-called vanishing point in the sky or on the horizon; second, going out of sight in "the dark" by reduction of illumination; and third, going out of sight by closure or covering of the eyes. Perhaps going out of sight in fog or mist is another kind, but it is similar to loss of structure by darkness (Chapter 4). I can also think of three kinds of occlusion other than self-occlusion (Chapter 5): first, at the edge of an opaque covering surface; second, at the edge of the field of view of an observer; and third, for celestial bodies, at the horizon of the earth. As for the going out of *existence* of a surface, there seem to be many kinds of destruction, so many that only a list of examples could be given in Chapter 6 on ecological events.

Locomotion in a Cluttered Environment

The following seven statements about reversible occlusion are taken from Chapters 1 to 5.

1. The substances of the environment differ in the degree to which they persist, some resisting dissolution, disintegration, or vaporization more than others.

2. The surfaces of the environment, similarly, differ in the degree to which they persist, some being transitory and others being relatively permanent. A surface goes out of existence when its substance dissolves, disintegrates, or evaporates.

3. Given an illuminated medium, a surface is unhidden at a fixed point of observation if it has a visual solid angle in the ambient optic array at that point. If it does not (but has at another point of observation), it is hidden.

4. For any fixed point of observation, the persisting layout of the environment is divided into hidden and unhidden surfaces. Conversely, for every persisting surface, the possible points of observation are divided into those at which it is hidden and those at which it is not.

5. A surface that has no visual solid angle at any point of observation is neither hidden nor unhidden. It is out of existence, not out of sight.

6. Any movement of a point of observation that hides previously unhidden surfaces has an opposite movement that reveals them. Thus, the hidden and the unhidden interchange. This is the *law of reversible occlusion* for locomotion in a cluttered habitat. It implies that after a sufficient sequence of reversible locomotions *all* surfaces will have been both hidden and unhidden.

7. The loci of occlusion are those places at which the hidden and unhidden surfaces into which a layout is temporarily divided are separated at occluding edges, there being two sorts, apical and curved. They are also the places where the hidden and unhidden surfaces are *joined* at occluding edges. Thus, to perceive an occluding edge of an object, even a fixed occluding edge at a fixed point of observation, is to perceive both the separation and the junction of its far and near surfaces.

The Motions of Detached Objects

Three more statements about reversible occlusion follow; they are taken from Chapter 5.

8. For any opaque object, the near surface, the temporary "front," hides the far surface, the temporary "back," at a fixed point of observation. The two interchange, however, when the object is rotated. The near surface also hides the *background* of the object, if present, but when the object is displaced the parts that go behind at one edge come from behind at the other. These facts can be observed in the film entitled *The Change from Visible to Invisible:* A *Study of Optical Transitions* (Gibson, 1968).

9. For both solidity and superposition, any motion of an object that conceals a surface has a reverse motion that reveals it.

10. To the extent that the objects of the environment have moved or been moved, the near and far sides of every object will have interchanged many times. This holds true over and above the extent to which the observer has moved around.

Head Turning

Following is the theorem about reversible occlusion when the observer looks around by turning her head. It is now assumed that the point of observation is occupied (Chapter 7).

11. For any fixed posture of the head, surfaces of the surrounding layout are divided into those inside the boundaries of the field of view and those outside the boundaries of the field. But with every turn of the head surfaces come into sight at the leading edge of the field of view and go out of sight at the trailing edge. The observer who looks around can thus see undivided surroundings and see herself in the middle of them.

Nonpersisting Surfaces

The next theorem is about the unreversing destruction and creation of surfaces and the unreversing optical transitions that accompany them (Chapter 6).

12. The going out of existence of a surface is not the reverse of its coming into existence, nor is the disturbance of optical structure that specifies one the reverse of the disturbance of structure that specifies the other. Hence, the disappearance of a surface by, say, dissolution can be distinguished from its disappearance by occlusion if the observer has learned to see the difference between the optical transitions. Such evidence as there is suggests that the two kinds of disappearance are usually distinguished (Gibson, Kaplan, Reynolds, and Wheeler, 1969). This is not to say that infants notice the difference, or even that adults always notice the difference. The difference may sometimes be hard to notice, as when a conjurer is playing tricks with one's perception. It is only to say that anyone can learn to see the difference.

The occlusion of a surface can be nullified, whereas the destruction of a surface cannot. Occlusion can be canceled by a movement of the body, head, or limbs in the opposite direction. Destruction, although it can sometimes be remedied, cannot simply be canceled by an opposite movement. It seems to me that young children must notice the optical transitions that can be thus nullified and those that cannot. How could they fail to pay attention to them? They play peek-a-boo, turn their heads, and watch their hands, all cases of reversible occlusion, and they also spill the milk, break the glass, and knock down the tower of blocks, things that cannot be reversed. But this hypothesis has not been tested with babies, because the only experiments carried out are in the spirit of rationalism promoted by J. Piaget, which asserts that children must form a concept of persistence or permanence and emphasizes what the children believe instead of what they see (for example, Bower, 1974, Ch. 7).

What is Seen at this Moment from this Position does not Comprise What is Seen

The old approach to perception took the central problem to be how one could see into the distance and never asked how one could see into the past and the future. These were not problems for perception. The past was remembered, and the future was imagined. Perception was of the present. But this theory has never worked. No one could decide how long the present lasted, or what distinguished memory from imagination, or when percepts began to be stored, or which got stored, or any other question to which this doctrine led. The new approach to perception, admitting the coperception of the self to equal status with the perception of the environment, suggests that the latter is timeless and that present-past-future distinctions are relevant only to the awareness of the self.

The environment seen-at-this-moment does not constitute the environment that is seen. Neither does the environment seen-from-this-point constitute the environment that is seen. The seen-now and the seen-from-here specify the self, not the environment. Consider them separately.

What is seen now is a very restricted sample of the surfaces of the world, limited to those that are inside the boundaries of the field of view at this head-posture. It is even limited to that surface being fixated at this eye-posture, if by *seen* one means *clearly seen.* This is at most less than half of the world and perhaps only a detail of that.

What is seen from here is at most the optically uncovered surfaces of the world at this point of observation, that is, the near sides of objects, the unhidden portions of the ground, the walls, and the bits that project through windows and doors.

The fact is that, although one can become aware of the seen-now and the seen-from-here if one takes the attitude of introspection, what one perceives is an environment that surrounds one, that is everywhere equally clear, that is in-the-round or solid, and that is all-of-a-piece. This is the experience of what I once called the visual world (Gibson, 1950*b*, Ch. 3). It has vistas that are connected and places that adjoin, with a continuous ground beneath everything, below the clutter, receding into the distance, out to the horizon.

The surface being fixated now at this momentary eye-posture is not a depthless patch of color, and the surfaces inside the field of view seen now at this head-posture are not a depthless patchwork of colors, for they have the quality that I called *slant* in the last chapter. The seen-at-this-moment is not the same, therefore, as the supposedly flat visual *field* analogous to the colors laid on a canvas by a painter that the old theory of color sensations asserted. I once believed that you could with training come to see the world as a picture, or almost do so, but I now have doubts about *it.* That comes close to saying that you can almost see your retinal image, which is a ridiculous assertion.

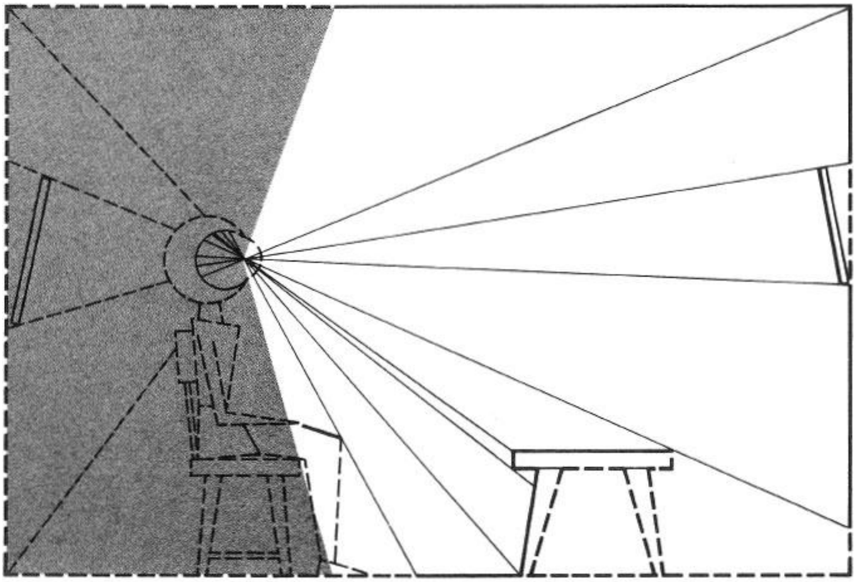

FIGURE 11.1 The surfaces viewed now from here by an observer seated in a room. At this temporary eye posture and this temporary head posture, the surfaces projected into the retinal image are indicated by solid lines and the remaining surfaces by dashed lines. The awareness of the here-and-now surfaces might be called *viewing* the room as distinguished from *seeing* the room. This is a vertical section of the observer and his monocular field of view.

The seen-from-here, from this stationary point of observation, is also not the supposedly flat visual field of tradition, for it is ambient. But it might justly be called *viewing the world in perspective,* or *noticing the perspectives of things.* This means the natural perspective of ancient optics, not the artificial perspective of the Renaissance; it refers to the set of surfaces that create visual solid angles in a frozen ambient optic array. This is a very small sample of the whole world, however, and what we perceive is the world.

Perception Over Time from Paths of Observation

It is obvious that a motionless observer can see the world from a single fixed point of observation and can thus notice the perspectives of things. It is not so obvious but it is true that an observer who is moving about sees the world at *no* point of observation and thus, strictly speaking, *cannot* notice the perspectives of things. The implications are radical. Seeing the world at a traveling point of observation, over a long enough time for a sufficiently extended set of paths, begins to be perceiving the world at *all* points of observation, as if one could be everywhere at once. To be everywhere at once with nothing hidden is to be all-seeing, like God. Each object is seen from all sides, and each place is seen as connected to its neighbor. The world is *not* viewed in perspective. The underlying invariant structure has emerged from the changing perspective structure, as I put it in Chapter 5.

Animals and people do in fact see the environment during locomotion, not just in the pauses between movements. They probably see better when moving than when stationary. The arrested image is only necessary for a photographic camera. An observer who is getting around in the course of daily life sees from what I will call a *path* of observation. A path does not have to be treated as an infinite set of adjacent points at an infinite set of successive instants; it can be thought of as a unitary movement, an excursion, a trip, or a voyage. A path of observation is the normal case, short paths for short periods of observation and long paths for hours, days, and years of observation. The medium can be thought of as composed not so much of points as of paths.

It sounds very strange to say that one can perceive an object or a whole habitat at no fixed point of observation, for it contradicts the picture theory of perception and the retinal image doctrine on which it is based. But it has to be true *if it is acknowledged that one can perceive the environment during locomotion.* The perception of the environment is understood to accompany the visual proprioception of the locomotion, of course, and the hypothesis of invariant structure underlying the changing perspective structure is required for this to be intelligible. These are unfamiliar notions. But the notion of ambulatory vision is not more difficult, surely, than the notion of successive snapshots of the flowing optic array taken by the eye and shown in the dark projection room of the skull.

The Problem of Orientation

Animals and humans are capable of being oriented to the habitat. This state is the opposite of being *disoriented* or "lost." The rat who can find its way directly to the goal box of a maze is said to be oriented to the goal. If there are many paths to the goal, the animal is capable of taking the shortest path. A person, similarly, can learn the way to work, to the post office, to the grocery store, and back home again through the passageways of his town. When he can do so in an unfamiliar town, he has become oriented in the new habitat. Both animals and humans are capable of homing. More generally, they are capable of way-finding. Or, in still other terms, they can do place-learning. Observers can go to the places in their environment that have affordances for them. If they are human observers, moreover, they may be able to *point* to these places, that is, to indicate their direction from here through the walls or other surfaces that hide them.

Two current explanations of how animals learn to find their way to hidden places are the theory of response chains and the theory of cognitive maps.

Neither is adequate. Way-finding is surely not a sequence of turning responses conditioned to stimuli. But neither is it the consulting of an internal map of the maze, for who is the internal perceiver to look at the map? The theory of reversible occlusion can provide a better explanation.

An alley in a maze, a room in a house, a street in a town, and a valley in a countryside each constitutes a place, and a place often constitutes a *vista* (Gibson, 1966*b*, p. 206), a semienclosure, a set of unhidden surfaces. A vista is what is seen from here, with the proviso that "here" is not a point but an extended region. Vistas are serially connected since at the end of an alley the next alley opens up; at the edge of the doorway the next room opens up; at the corner of the street the next street opens up; at the brow of the hill the next valley opens up. To go from one place to another involves the opening up of the vista ahead and closing in of the vista behind. A maze or a cluttered environment provides a choice of vistas. And thus, to find the way to a hidden place, one needs to see which vista has to be opened up next, or which occluding edge hides the goal. One vista leads to another in a continuous set of reversible transitions. Note that in a terrestrial environment of semienclosed places each vista is unique, unlike the featureless passageways of a maze. Each vista is thus its own "landmark" inasmuch as the habitat never duplicates itself.

When the vistas have been put in order by exploratory locomotion, the invariant structure of the house, the town, or the whole habitat will be apprehended. The hidden and the unhidden become one environment. One can then perceive the ground below the clutter out to the horizon, and at the same time perceive the clutter. One is oriented to the environment. It is not so much having a bird's-eye view of the terrain as it is being everywhere at once. The getting of a bird's-eye view is helpful in becoming oriented, and the explorer will look down from a high place if possible. Homing pigeons are better at orientation than we are. But orientation to goals behind the walls, beyond the trees, and over the hill is not just a looking-down-on, and it is certainly not the having of a map, not even a "cognitive" map supposed to exist in the mind instead of on paper. A map is a useful artifact when the hiker is lost, but it is a mistake to confuse the artifact with the psychological state the artifact promotes.

Note that the perception of places and the perception of detached objects are quite different. Places cannot be displaced, whereas objects can be, and animate objects displace themselves. Places merge into adjacent places, whereas objects have boundaries. Orientation to hidden places with their attached objects can be learned once and for all, whereas orientation to movable objects has to be relearned continually. I know where the kitchen sink is, I think I know where the ski boots are stored, but I don't always know where my child is. One can only go to the last known locus of a detached object. Hidden objects can be moved without that event being perceived, and the unhappy state of the man whose car keys are seldom where he left them is notorious.

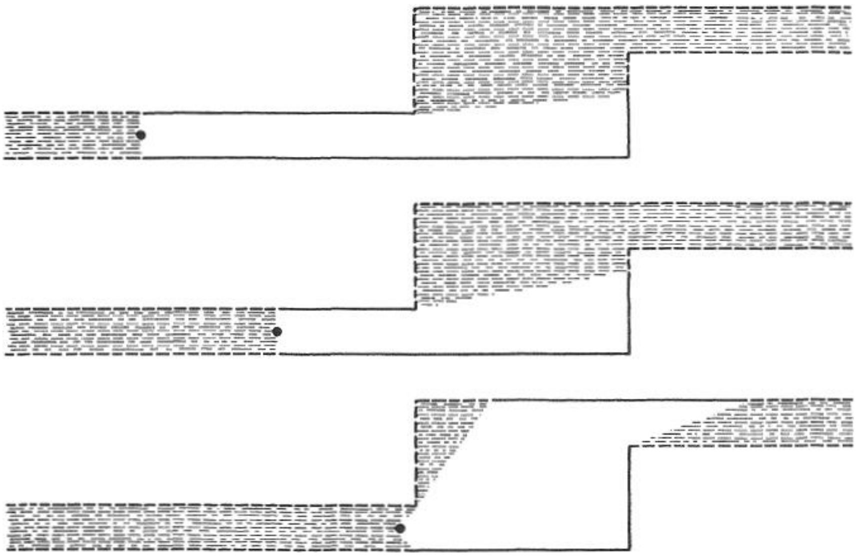

FIGURE 11.2 The opening up of a vista at an occluding edge, as seen from above. This is a plan view of a passageway that opens on a courtyard from which another passageway leads. As an observer moves along the corridor, the surfaces behind his head progressively go out of sight, and surfaces in front progessively come into sight at one occluding edge and then the other. The hidden portions of the ground are indicated by hatching. The hidden portions of the walls are indicated by dashed lines. The position of the observer is indicated by a black dot.

In the pages above I have formulated a theory of orientation to the places of the habitat. The perceiving of the world entails the coperceiving of where one is in the world and of being in the world at that place. This is a neglected fact that is neither subjective nor objective. To the extent that one has moved from place to place, from vista to vista, one can stand still in one place and see where one is, which means where one is relative to where one might be. One does not need a map with a circle on it labeled, "You are here." I suggest that this constitutes the state of being oriented.

The Problem of Public Knowledge

The hypothesis of reversible optical transformations and occlusions resolves the puzzle of how, although the perspective appearances of the world are different for different observers, they nevertheless perceive the same world. Perspective appearances are not the necessary basis of perception.

It is true that there is a different optic array for each point of observation and that different observers must occupy different points at any one time. But observers move, and the same path may be traveled by any observer. If a set of observers move around, the same invariants under optical transformations and occlusions will be available to all. To the extent that the invariants are detected, all observers will perceive the same world. Each will also be aware that his or her place in the world is different here and now from that of any other.

Points, of course, are geometrical concepts, whereas places are ecological layouts, but the above theory can also be put geometrically: although at a given instant some points of observation are occupied and the remainder unoccupied, the one set can go into the other.

The theory asserts that an observer can perceive the persisting layout from other places than the one occupied at rest. This means that the layout can be perceived from the position of another observer. The common assertion, then, that "I can put myself in your position" has meaning in ecological optics and is not a mere figure of speech. To adopt the point of view of another person is not an advanced achievement of conceptual thought. It means, *I can perceive surfaces hidden at my point of view but unhidden at yours.* This means, *I can perceive a surface that is behind another.* And if so, *we can both perceive the same world.*

The Puzzle of Egocentric Awareness

Psychologists often talk about egocentric perception. An egocentric perceiver is supposed to be one who can see the world only from his own point of view, and this habit is sometimes thought to characterize an egocentric *person.* Egoism is thought to come naturally to humans because they are innately aware of their private experiences and do not easily learn to adopt the point of view of others. This line of thinking now seems mistaken. Perception and proprioception are not alternatives or opposing tendencies of experience but complementary experiences.

The sensation-based theories of perception assume that the perspective appearances of the world are all that a newborn infant is given. They are the data for perception. Hence, the young child is necessarily egocentric, and cognitive development is a matter of progressing from subjective sensations to objective perceptions. The child's ego encompasses the world, and at the same time she is supposed to be confined to the awareness of her fleeting sensations. But there is a reason to be suspicious of all these speculations. The evidence about the earliest visual experiences of infants does not suggest that they are confined to surfaces seen-now-from-here, and the evidence definitely contradicts the doctrine that what they see is a flat patchwork of color sensations. I therefore suspect that the supposed egocentricity of the young child is a myth.

Hiding, Peeking, and Privacy

In Chapter 8 on affordances, I described how some of the places of an environment are *hiding* places. That is, they afford the hiding of oneself or of one's property from the sight of other observers. The phenomenon of seeing without being seen illustrates the application of optical occlusion to social psychology. The passage on hiding places in Chapter 8 should be reread.

The perceiving of occluded places and objects does occur and can be shared with other perceivers. To this extent, we all perceive the same world. But there is also ignorance of occluded things, and if you hide from me your private property, your hideaway in the hills, your secret lover, or the birthmark on your buttocks, then you and I do not perceive quite the same world. Public knowledge is possible, but so is its reciprocal, private knowledge.

Not only do babies like to play peek-a-boo and children to play hide-and-seek, but animals who are preyed upon hide from the predator, and the predator may hide from the prey in ambush. One observer often wants to spy upon others, to see without being seen. He peers through a peephole or peeks around the occluding edge of a corner. In opposition to this is the striving *not* to be seen by others, the need for privacy. Burrows, caves, huts, and houses afford not only shelter from wind, cold, and rain but also the state of being out of sight, or out of the "public eye."

The human habit of covering the body with clothing whenever one is in sight of others is a matter of hiding some skin surfaces but not others, depending on the conventions of the culture. To display the usually covered surfaces is improper or immodest. The providing of some information for the layout of these hidden surfaces, however, is the aim of skillful clothing designers. And the careful manipulation of the occluding edges of clothing with progressive revealing of skin is a form of the theatrical art called stripping.

Summary

The demonstration that reversible occlusion is a fact of visual perception has far-reaching implications. It implies that an occluding edge is seen as such, that the persistence of a hidden surface is seen, and that the connection of the hidden with the unhidden is perceived. This awareness of what-is-behind, and of the togetherness of the far side and the near side of any object, puts many of the problems of psychology in a new light.

The doctrine that all awareness is memory except that of the present moment of time must be abandoned. So must the theory of depth perception. The importance of the fixed point of view in vision is reduced. But a new theory of orientation, of way-finding, and of place-learning in the environment becomes possible. And the puzzles of public knowledge, of egocentricity, and of privacy begin to be intelligible.