

Perceptual-cognitive universals as reflections of the world

Roger N. Shepard

Department of Psychology, Stanford University, Stanford, CA 94305-2130

roger@psych.stanford.edu

Abstract: The universality, invariance, and elegance of principles governing the universe may be reflected in principles of the minds that have evolved in that universe – provided that the mental principles are formulated with respect to the abstract spaces appropriate for the representation of biologically significant objects and their properties. (1) *Positions* and *motions* of objects conserve their shapes in the geometrically fullest and simplest way when represented as points and connecting geodesic paths in the six-dimensional manifold jointly determined by the Euclidean group of three-dimensional space and the symmetry group of each object. (2) *Colors* of objects attain constancy when represented as points in a three-dimensional vector space in which each variation in natural illumination is canceled by application of its inverse from the three-dimensional linear group of terrestrial transformations of the invariant solar source. (3) *Kinds* of objects support optimal generalization and categorization when represented, in an evolutionarily-shaped space of possible objects, as connected regions with associated weights determined by Bayesian revision of maximum-entropy priors.

Keywords: apparent motion; Bayesian inference; cognition; color constancy; generalization; mental rotation; perception; psychological laws; psychological space; universal laws

Introduction

The ways in which genes shape an individual's perceptual and cognitive capabilities influence the propagation of those genes in the species' ecological niche just as much as the ways in which those genes shape the individual's physical size, shape, and coloration. A predatory bird has come to have not only sharp talons but also sharp eyes, and a small rodent has come to have not only quick feet but also quick recollection of the location of its burrow. Moreover, natural selection favors adaptation to any biologically relevant property of the world, whether that property holds only within a particular species' local niche or throughout all habitable environments. Thus, both the hawk and the ground squirrel have internalized the period of the terrestrial circadian cycle, whose 24-hour value is the same everywhere on earth and whose invariance is a consequence of a law – the conservation of angular momentum – holding throughout the universe.

From among the general properties that characterize the environments in which organisms with advanced visual and locomotor capabilities are likely to survive and reproduce, here I focus on the following three. (1) Material objects are generally conserved and, when they move (whether relative to the stable environment or to the self-moving observer), move in ways whose possibilities and geometrical simplicities are determined by the three-dimensional, Euclidean character of physical space. (2) The light scattered to an eye from an object's surface bilinearly conflates the invariant spectral reflectance properties of the surface itself and the momentary spectral composition of the illumination, which is subject to three principal degrees of freedom of linear transformation. (3) Objects that are of the same basic kind and, hence, that have the same biologically significant potential

(e.g., of being edible, poisonous, predatory, or suited to mating, parenting, and hence propagation of one's genes), generally form a connected local region in the space of possible objects, despite appreciable differences among individual objects of that kind in size, shape, position, motion, or color.

In perceptually advanced mobile organisms, then, genes that have internalized these pervasive and enduring facts about the world should ultimately prevail over genes that leave it to each individual to acquire such facts by trial and possibly fatal error. If so, psychological science may have unnecessarily restricted its scope by implicitly assuming that psychological principles, unlike the universal laws of physics, apply at most to the particular animals that happen to have evolved on one particular planet. When formalized at a sufficient level of abstraction, mental principles that have evolved as adaptations to principles that have long held throughout the universe might be found to partake of some of the generality of those prior principles (Shepard 1987a) – perhaps even attaining the kind of universality, invariance,

ROGER SHEPARD is the Ray Lyman Wilbur Professor of Social Science, Emeritus, of Stanford University and a recipient of the 1995 National Medal of Science. In the early 1960s, he devised the auditory illusion of endlessly rising pitch and developed the first of the methods of “nonmetric” multidimensional scaling, which he has used to reveal hidden structures and functional laws of generalization and perceptual and cognitive representation. Shepard then introduced the paradigm of “mental rotation” and with his students demonstrated the analog character of such imagined transformations. Shepard has sought universal mental principles that reflect universal features of the world.

and formal elegance (if not the quantitative precision) previously accorded only to the laws of physics and mathematics.

My own searches for universal psychological principles for diverse perceptual-cognitive domains have been unified by the idea that invariance can be expected to emerge only when such principles are framed with respect to the appropriate *representational space* for each domain. This idea was inspired, in part, by Einstein's demonstration that in extending physical principles beyond the biologically relevant scales of distance, velocity, mass, and acceleration, invariance could still be achieved – but only by casting those principles in terms of the appropriate four-dimensional space-time manifold. Invariance of the laws of physics was no longer restricted to inertial frames moving at velocities that are small relative to the speed of light (as in Newtonian mechanics, formulated with respect to three-dimensional Euclidean space), or even inertial frames moving at any possible (i.e., subluminal) velocity (as in special relativity, re-formulated with respect to [3+1]-dimensional Minkowski space). Only when reformulated yet again, with respect to the appropriately curved, [3+1]-dimensional Riemannian space, did the laws of physics finally become (in general relativity) invariant with respect to arbitrarily accelerated frames. Moreover, the motions of objects actually observed in the world were then explained, and explained most accurately (as confirmed, first, in accounting for the perihelion advance of Mercury's orbit and, subsequently, in other ways), not in terms of forces acting instantaneously across arbitrarily large distances in three-dimensional Euclidean space but solely in terms of the local geometry of the curved four-dimensional space-time manifold in the vicinity of the object itself. The paths of motion (like great circles on the surface of the Earth) were now simply the *geodesics*, the direct analogs of straight lines in the curved four-dimensional manifold.

But, for such biologically relevant properties of objects as their positions, motions, shapes, colors, and kinds, what sorts of representational spaces show promise of yielding invariant psychological principles? And if such representational spaces and associated psychological principles arose not accidentally but as adaptations to general properties of the world in which we have evolved, can an identification and analysis of such sources in the world point the way toward elegant and invariant formalizations of the corresponding psychological principles?

1. Representations of an object's position, motion, and shape

Position, motion, and shape are best considered together because, from the abstract, geometrical point of view that promises the most elegant and invariant formulation, the representations of these three attributes are inextricably interconnected. I focus initially and most extensively on the representations of positions and rigid motions between positions. Shape I can consider only briefly here, merely observing that the shape of an object may be understood in terms of the object's approximations to all possible symmetries, which in turn may be understood in terms of the object's self-similarities under all possible rigid motions.

The positions, motions, and shapes that are possible for an object depend on the kind of space within which that object is confined. On a biologically relevant scale (of size, ve-

locity, mass, and acceleration), one of the most general facts about the world in which we have evolved is that it is spatially three-dimensional and Euclidean. But how do we demonstrate that humans or other animals have internalized the geometry peculiar to this particular type of space?

The internalization of the circadian cycle was established when animals were raised in artificial isolation from the terrestrially prevailing 24-hour oscillation in illumination and temperature and were found, even so, to maintain a close approximation to their previous 24-hour activity cycle. (As the old quip has it: "You can take the boy out of the country, but you can't take the country out of the boy.") Similarly, the three-dimensionality of our world is so deeply entrenched in our mental makeup that while we may muse, "If only I had a larger office, I would have more room for my books," it does not occur to us to think, "If only I had a four-dimensional office, I would have more degrees of freedom for arranging them!" The very universality of the three-dimensionality of our world precludes our taking "the boy" or, indeed, the girl, the hawk, or the ground squirrel out of this three-dimensional "country," to see whether, in the absence of external support, any of these creatures would continue to perceive and to think three-dimensionally. We can, however, investigate whether an individual, though remaining in three-dimensional space physically, is able to take an object out of that space mentally, when only such a move could achieve compliance with another deeply internalized principle, such as the principle of object conservation.

Apparent motion, which is typically induced in an observer by alternately presenting two identically shaped objects in different static positions, provides one means of exploring this possibility. In the absence of any physically presented motion, the particular motion that is experienced must be a direct reflection of the organizing principles of the viewer's brain. The Gestalt psychologists, who were responsible for most of the early studies of apparent motion (see, e.g., Koffka 1931; 1935; Korte 1915; Wertheimer 1912), regarded such organizing principles as manifestations, in the neurophysiological medium of the brain, of minimization principles that operate in physical media generally – much as the spherical shape of a soap bubble arises from principles of conservation of matter (the enclosed volume of air) and minimization of surface area (the enclosing film of soap, with its surface tension). The uniquely powerful organizing principles of the brain are not, however, likely to be wholly explained by properties that grey matter shares with all matter. The neuronal circuits of the brain (unlike the molecules of such media as air or soap films) have been shaped by natural selection specifically to provide a veridical representation of significant objects and events in the external world.

1.1. Apparent motion achieves object conservation

Why, for example, does one experience a single object moving back and forth at all, rather than experiencing what is actually being physically presented in the laboratory – namely, two visual stimuli going on and off separately? Quite apart from questions about the particular type of movement experienced, the fact that any connecting movement is experienced is presumably the manifestation of an internalized principle of object conservation. It is simply more probable in our world that an enduring object abruptly moved from one position to a nearby position, than that one object suddenly ceased to exist and, at exactly the

same instant, a separate but similar object just as suddenly materialized in another position. Still, if the benefits of representing objects as enduring entities support the instantiation of a connecting motion, two questions remain: Out of the infinity of such possible motions, which particular motion will be instantiated? What formal characterization of that psychologically preferred motion will most elegantly reflect any simplicity, universality, and invariance of its ultimate source in the world?

1.2. Apparent motion is experienced in three-dimensional space

When identical two-dimensional shapes, such as the Cooper (1975) polygons adopted for illustration in Figure 1a, are alternately presented in orientationally different positions in their common two-dimensional plane, a single such shape is experienced as rigidly rotating about a fixed point in that plane (e.g., Farrell & Shepard 1981; Robins & Shepard 1977; Shepard 1981b; 1984). Similarly, when identical three-dimensional shapes, such as the Shepard-Metzler (Shepard & Metzler 1971) objects shown in Figure 1c, are alternately presented in their common three-dimensional space, a single such object is experienced as rigidly undergoing a rotational (most generally, a screw-like) motion in that space (Shepard 1984; Shepard & Judd 1976; see also Carlton & Shepard 1990a).

But what happens if the two alternately presented shapes are not identical but *enantiomorphic* – that is, mirror images of each other, like a right and left hand? Asymmetric shapes cannot be transformed into each other by any rigid motion confined to the plane or space in which they reside. They can be brought into congruence there only by a shape-reversing reflection of one of the two objects through some line or plane in their two- and three-dimensional spaces, respectively. Nevertheless, between mirror-image polygons in the plane (Fig. 1b), a rigid motion is still experienced. But it is necessarily experienced as a rotation out of the plane, through the three-dimensional space containing that plane (Shepard 1984). Presumably, we perceptually liberate the object from the two-dimensional plane for two reasons: having evolved in a three-dimensional world, we are just as




Figure 1. Pairs of alternately presented visual shapes (polygons like those used by Cooper 1975, or block models like those used by Shepard & Metzler 1971) that give rise to four different types of apparent motion: (a) a rigid 90° rotation in the picture plane, (b) a rigid 180° rotation out of the plane and through three-dimensional space, (c) a rigid screw displacement in three-dimensional space, and (d) nonrigid motion only.

capable of representing a rigid motion in three-dimensional space as in a two-dimensional plane. But only the motion in three-dimensional space can represent the shape conservation that is probable in the world – particularly for objects like those in Figure 1, bounded by straight edges or flat surfaces. (This is, incidentally, one reason for our use of stimuli composed of straight lines. The probability that an arbitrarily transformed object will give rise to straight lines in a two-dimensional projection is vanishingly small if nonrigid deformations are allowed. For curved free-form shapes, apparent motion is often experienced as a nonrigid deformation. Moreover, comparison of such shapes by mental rotation is far less accurate – see, for example, Rock et al. 1989.)

Between enantiomorphic solid objects portrayed as in three-dimensional space (Fig. 1d), however, viewers never report experiencing a rigid motion. Such a rigid motion is still mathematically possible – but only by breaking out of the three-dimensional space in which we and our object have been confined, so that we can rigidly rotate the object (now about a plane!) in a surrounding, more commodious four-dimensional space. Failing to achieve even a mental liberation from the only space we have known, we are destined to experience all motions as confined to that three-dimensional space and, hence, all transformations between enantiomorphic shapes as nonrigid. For shapes of the kind illustrated in Figure 1d, at least one of the “arms” of the object typically appears to rotate independently, as if connected to the rest of the object by some sort of swivel joint (a type of motion that, although less common than globally rigid motion, does occur in a world biologically enriched with joint-limbed animals and wind-fractured tree branches).

Similarly, computer generated projections of actual (as opposed to merely apparent) rotations of rigid structures give rise to the “kinetic depth” perception of rigidity for arbitrary rotations in three-dimensional space but not for arbitrary rotations in four-dimensional space (see, e.g., Green 1961; Noll 1965). These phenomena of real and apparent motion (as well as related phenomena of merely imagined motion, e.g., mental rotation) are consonant with the Kantian idea that we are constituted to represent objects and events only in Euclidean space of three (or fewer) dimensions. The modern evolutionary/mechanistic explication of this idea must be that the three-dimensional world simply has not exerted sufficient selective pressures toward the evolution of the more complex neuronal machinery that would be required to represent higher-dimensional spaces and the additional rigid transformations that such spaces afford.

1.3. Apparent motion traverses a kinematically simplest path

Even when a connecting motion is possible within three-dimensional space (as in Fig. 1c), the particular motion experienced is only one out of infinitely many possible rigid motions between the two presented positions. One might be tempted to guess that if apparent motion is guided by internalized approximations to principles holding at the biologically relevant scale in the external world, the most likely candidates for those external principles would be those of prerelativistic, Newtonian mechanics. This guess has proved untenable, however, in the face of several facts:

(1) Any rigid motion is compatible with Newton’s laws of motion, in the presence of arbitrary unseen forces. Hence, unless we exclude such forces, Newtonian mechanics itself

provides no basis for the selection of one path of motion over another.

(2) If we do exclude such forces, however, Newton's laws constrain an object's center of mass to traverse a straight line. But this is contrary to the now well-established finding that apparent motion tends to be over a curved path when the two positions in which the object is alternately presented differ in orientation (see, for example, Bundesen et al. 1983; Farrell 1983; Foster 1975b; Kolars & Pomerantz 1971; McBeath & Shepard 1989; Proffitt et al. 1988).

(3) The apparent motions that are most apt to be experienced, as well as the real motions that are discriminated most accurately and judged to be most simple, are those motions whose rotational component is about an axis determined by the geometry of the object's visible shape rather than by the physics of the object's invisible distribution of mass. In particular, the psychologically preferred axes of rotation are those of global or local symmetry of the shape as in Figure 2a – not the principal axes of inertia of the object as in Figure 2c (Carlton & Shepard 1990b). (The latter axes are not even directly determined by the object's visual shape, and can only be inferred by making an additional assumption, such as that the object is of uniform density.) Even an object, such as a cube, for which all possible rotational axes are inertially equivalent appears to rotate about a fixed axis when actually rotated about an axis of symmetry, as in Figure 2b, but appears to wobble when actually rotated about an axis that (though inertially equivalent) is not an axis of geometrical symmetry, as in Figure 2d (Shiffrar & Shepard 1991).

(4) Human infants reveal sensitivity to essentially geometrical constraints such as continuity, rigidity, and impen-




Figure 2. Axes of geometrical symmetry (a) favored by apparent motion and (b) around which real motion appears stable and is accurately compared, and nonsymmetry axes of physical inertia (c) avoided by apparent motion and (d) around which real motion appears to wobble and is less accurately compared. Figures 2a and 2c are from "Psychologically Simple Motions as Geodesic Paths: II. Symmetric Objects," by E. H. Carlton and R. N. Shepard, 1990, *Journal of Mathematical Psychology* 34:208. Copyright 1990 by Academic Press. Adapted by permission. Figures 2b and 2d are from "Comparison of Cube Rotations About Axes Inclined Relative to the Environment or to the Cube," by M. Shiffrar and R. N. Shepard, 1990, *Journal of Experimental Psychology: Human Perception and Performance* 7:48. Copyright 1990 by the American Psychological Association. Adapted by Permission.

etrability before manifesting sensitivity to constraints of physical dynamics based on gravity, mass distribution, and inertia (Spelke 1991).

(5) Even adults, from Aristotle to present-day college students, often manifest an "intuitive physics" that fails to comply with the constraints of Newtonian mechanics (McClosky 1983; Proffitt & Gilden 1989; Proffitt et al. 1990; see also Shepard 1987a, pp. 266-67), although in some such cases it may approximate constraints of kinematic geometry (see Shepard 1984; 1987a).

(6) Abstract geometrical constraints apply to a wider range of phenomena in the world than do concrete physical constraints and, for this reason, would presumably have had more opportunity for internalization through natural selection (as well as through learning). Things as lacking in mechanical rigidity as a constellation, a curl of smoke hanging in still air, or a shadow, all undergo transformations that (at least over sufficiently short periods of time) approximate geometrical rigidity relative to a moving or turning observer (Shepard 1984; Shepard & Cooper 1982). As Gibson observed, such self-induced geometrical transformations of the "ambient optic array" are probably the most ubiquitous of the transformations with which the visual systems of highly mobile animals must cope (e.g., Gibson 1979). We can understand then why apparent motion might be primarily governed not by the principles of Newtonian mechanics but, rather, by the more abstract and widely manifested constraints of kinematic geometry for three-dimensional space (Shepard 1984).

1.4. Kinematic simplicity is determined by geometry

Kinematic geometry is the branch of mathematics characterizing the motions that are geometrically possible and, among those, the motions that are in a purely geometrical sense most simple or natural – given a geometrical specification both of the object or set of objects and of any constraints on its possible motions. The objects may be geometrically specified to be shape-invariant under all transformations (i.e., rigid). The constraints on their motions may be geometrically specified to preclude mutual interpenetration; escape from their particular embedding space (having specified dimensionality, curvature, and global topology); or violation of the constraints on their relative motions imposed by specified mechanical interconnections (such as a one-degree-of-freedom hinge or slider, a two-degrees-of-freedom pivot, a three-degrees-of-freedom ball and socket joint, etc.). Kinematic geometry says nothing about physical mass, force, acceleration – and, hence, nothing about how much and what kind of effort would be required actually to carry out any particular specified motion, physically, for any given mass distribution within each component object (to say nothing of a specification of the friction at each joint or sliding surface, of the density and viscosity of the medium in which the objects might be immersed, or of how much and what kind of force can be applied before a physical component will bend, fracture, or break). The abstract constraints of geometry are thus conceptually separable from the more concrete constraints of physics: questions of whether a certain large table will fit through a particular door and, if so, what simple sequence of translations and rotations of the table will suffice are purely geometrical and quite distinct from questions of how many persons should be recruited for the job, or of

which geometrically possible sequences of rigid transformations will require the least physical effort.

For present purposes, we need consider only the simplest case of the motion of a single rigid object. Even for this simplest case, full mathematical characterization was not achieved until the last century (following the development of the relevant mathematical apparatuses of group theory, Lie algebras, quaternions, and differential geometry). Particularly relevant, here, is Chasles's (1830) theorem of kinematic geometry, according to which any two positions of an asymmetric shape in three-dimensional Euclidean space determine a unique corresponding axis through that space such that the object can be rigidly transported from either position to the other by a combination of a linear translation along that axis and a simple rotation about that same axis – that is, by the helical motion called a *screw displacement*. In particular cases, the translational or the rotational component may be null, leaving only the degenerate screw displacement of (respectively) a pure rotation, a pure translation, or (if both components are null) no motion at all.

If the two positions of an asymmetric object are confined to the Euclidean plane, as in Figure 1a, Chasles's theorem reduces to Euler's theorem. The two positions then determine a unique point in the plane such that the object can be rigidly carried from either position to the other by a simple rigid rotation in the plane about that point. (For generality and elegance, the degenerate case of pure translation is interpreted, in the abstract mathematical formalism, as a rotation of the object about a "point at infinity.")

Strictly, what is uniquely determined by the geometry of the two positions of an (asymmetric) object, is the geodesic path along which a rigid transformation can carry the object back and forth between those positions. Alternative motions along complementary segments of that same geodesic may be possible. Thus, a rotation can carry an object between two positions through either of two nonoverlapping paths around the same circle. Generally, apparent motion tends to be experienced over the shorter of two such alternative paths. But, when the presented positions of the object differ by close to 180° , the two alternatives are of nearly equal length and either motion may be experienced (see Farrell & Shepard 1981; Robins & Shepard 1977). (The case of objects possessing various symmetries, for which two positions of the object may be connected by different screw displacements around two or more distinct axes, will be considered later.) Even when the particular segment of the geodesic over which the motion is to be represented has been determined, kinematic geometry itself does not prescribe the time course of that motion – whether it must be fast or slow, accelerating or decelerating, and so on. In the physical world, the time course of an actual motion is determined by physical dynamics, based on the mass distribution and forces applied. In the mental world, however, the time course of the motion perceptually experienced in apparent motion or only imagined in mental rotation appears to be primarily determined by other, more general, invariant, and adaptively critical constraints, as I shall argue.

Of course, the screw displacements (including simple rotations) prescribed by kinematic geometry are not the only possible motions between two positions of an object in space or in the plane. There are always infinitely many possible motions, including infinitely many rigid motions in which the axes of rotation and translation can vary in orien-

tation from moment to moment and can depart from mutual alignment during the motion, as well as infinitely many more motions that do not preserve the rigid structure of the object. Natural selection has ensured that (under favorable viewing conditions) we generally perceive the transformation that an external object is actually undergoing in the external world, however simple or complex, rigid or nonrigid. Here, however, I am concerned with the default motions that are internally represented under the unfavorable conditions that provide no information about the motion that actually took place between two successive positions of an object. What I am suggesting is that when a simple screw displacement or rigid rotation is possible, that motion will tend to be represented because, of all transformations that conserve the object at the fullest level of shape, it is the geometrically simplest and hence, perhaps, the most quickly and easily computed. Certainly, within a general system suitable for specifying all possible rigid motions, such a motion requires the minimum number of parameters for its complete specification.

1.5. Geometry is more deeply internalized than physics

In accordance with Chasles's theorem, when an asymmetric shape is alternately presented in two orientationally different positions (as in Fig. 1c), under conducive conditions, human viewers generally do report the experience of a helical motion (Shepard 1984). The "conducive conditions" are primarily those in which the temporal interval between the offset of each stimulus and the onset of the other is short enough to yield a pattern of retinal stimulation consistent with some (necessarily rapid) actual motion, and the interval between the onset of each stimulus and the onset of the other is long enough, relative to the extent of the geometrically simplest rigid transformation, to permit completion of the (necessarily rate-limited) neural computations required for that transformation. If the two alternately presented positions of the object are confined to a plane (as in Fig. 1a), the experienced motion generally reduces to a simple rigid rotation around a fixed point in the plane, in accordance with the special case known as Euler's theorem. This single rigid rotation is geometrically simpler than the motion prescribed by Newtonian mechanics, which generally includes two components: a continuous motion of the center of mass (which is rectilinear in the absence of external forces), and an independent rotation about that moving center. Indeed, for a Newtonian motion in three-dimensional space, the axis of rotation need not retain an invariant orientation. Even in the absence of external forces, the axis of momentary rotation will itself wobble about the moving object's center of mass, unless the axis of rotation happens to coincide with a principal axis of inertia of the object. Only in the special case in which the two alternately presented positions of an object have identical orientations does the helical motion prescribed by kinematic geometry coincide with the rectilinear motion prescribed by Newtonian mechanics. Thus, the "intuitive physics" revealed by tests involving spatially extended bodies and rotational motions may deviate from classical physics (e.g., McClosky 1983; Proffitt & Gilden 1989; Proffitt et al. 1990) because whatever internalized knowledge of physical dynamics is tapped by such tests may be contaminated, to a variable degree across individuals and conditions of testing, by a more deeply internalized wisdom about kine-

matic geometry (Shepard 1984; 1987a; see also Freyd & Jones 1994).

In the method that McBeath and I introduced for quantifying the extent of the departure of apparent motion from a rectilinear path, a shape was alternately presented in different orientations on the left and right of a visual wall and observers adjusted the vertical height of a window in the wall so that the object appeared most compellingly to pass back and forth through that window, which was just large enough to accommodate the object. Figure 3a illustrates the two-dimensional display used in the initial study (see McBeath & Shepard 1989).

The obtained height-of-window settings uniformly implied a curvature away from the straight path, in the direction prescribed by kinematic geometry. As shown in Figure 3b, for linear separations of up to at least 3° of visual angle and for orientational differences of up to at least 90° between the alternately presented stimuli – for which the experience of motion over a particular path was still strong and well-defined – the settings were remarkably close to those prescribed by kinematic geometry. Even for larger separations and angular differences (viz. 180°), for which the experience of motion became weaker and less well-defined, the mean settings remained closer to the circular paths prescribed by Euler's theorem than to the rectilinear paths prescribed by Newtonian dynamics in the absence of external forces. Preliminary indications of similar deviations from rectilinearity have also emerged in subsequent unpublished explorations of the three-dimensional case, where the deviations are generally expected to be helical rather than merely circular. (For example, McBeath, using a computer generated full-color stereoscopic display, had viewers position a circular window anywhere in a two-dimensional wall that appeared to recede in depth, dividing a virtual room into left and right compartments within which the two positions of a three-dimensional object were alternately displayed.)

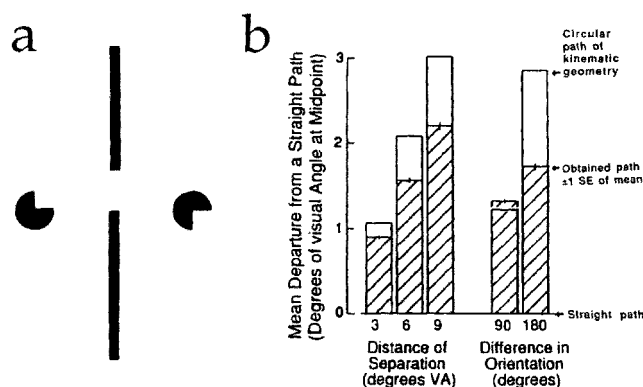


Figure 3. Depictions of (a) shapes alternately presented in different orientations on the left and right of a wall with a window whose height could be adjusted so that a single object appeared to pass back and forth through the window, and (b) mean displacements of the window above the height of a straight path of apparent motion that subjects produced for different linear and angular separations between the shapes. From "Apparent Motion Between Shapes Differing in Location and Orientation: A Window Technique for Estimating Path Curvature," by M. K. McBeath and R. N. Shepard, 1989, *Perception & Psychophysics* 46:334-35. Copyright 1989 by the Psychonomic Society.

1.6. Object symmetries entail alternative paths of apparent motion

For an object possessing some symmetry or symmetries, different screw displacements may be possible between two positions of the object about two or more distinct axes in space. A horizontal rectangular bar in the plane provides a simple illustration. Such a shape is identical to itself under 180° rotation (in the plane) about its center, and under 180° rotations (in space) about either a vertical or a horizontal axis through its center. As a consequence of these symmetries, when such a bar is alternately presented on the left and right, it may be experienced as rigidly moving over any one of seven different paths along five distinct geodesics between the two presented positions, and each of these motions is a screw displacement (if we include, as always, the degenerate screw displacements of pure rotation or pure translation).

Along one geodesic, there are two nonoverlapping 180° rotations in the picture plane around a point midway between the two positions in which the bar appears, one path through the upper portion of the plane, the other through the lower. Along a second geodesic, there are two nonoverlapping 180° rotations in depth about a vertical axis lying in the picture plane midway between the two presented positions, one through the three-dimensional space in front of the plane, the other through the space behind. In each of these first two cases, the two alternative motions correspond to the traversal of two complementary halves of a circular geodesic. Along a third geodesic, two distinct paths of rectilinear translation in the picture plane are geometrically possible between the two positions, one over the short segment of the horizontal line directly between the two side-by-side positions presented, the other over the infinitely longer path corresponding to the complementary part of that horizontal line (interpreted as the complete circle around a "center at infinity"). Finally, along the remaining geodesics, several distinct screw displacements are possible along this same line, in which the bar simultaneously translates and rotates 180° about the short segment of that axis in either direction, or in which the screw displacement entails (again) an infinitely longer translational component over the remaining part of the horizontal line. For these last geodesics, the longer paths of possible transformation, being infinitely longer, are not experienced, leaving just seven likely paths of geodesic transformation.

To obtain experimental evidence that these are the default paths of transformation between two such horizontally separated positions of a rectangular bar, Susan Zare and I primed motions over each of these four paths by appropriately adding a small symmetry-breaking extension to each rectangular bar, giving it the suggestion of one of the four possible L shapes, as shown in Figure 4 (see Carlton & Shepard 1990b, pp. 219-21). With the extension always attached to the upper left corner of the left bar, the symmetry of the right bar could be broken by attaching the corresponding extension to its upper left, upper right, lower right, or lower left (as shown in Fig. 4a, b, c, and d, respectively). The apparent motion tended accordingly to be experienced as a rectilinear translation along the horizontal axis common to the two rectangles (Fig. 4a), as a 180° rotation in depth about the vertical axis lying in the plane halfway between the two rectangles (Fig. 4b), or (less com-

pellingly, for reasons soon to be noted) as a 180° rotation in the plane about the horizontal line-of-sight axis orthogonal to the plane through a point halfway between the two rectangles (Fig. 4c), or as a 180° screw displacement along the horizontal axis common to the two rectangles (Fig. 4d).

The apparent rotation in the plane (corresponding to Fig. 4c) could also be induced by a form of *path-guided apparent motion* (cf., Shepard & Zare 1983). A low-contrast uniform gray static path was briefly exposed during the 5 msec interval between the offset of each bar and the onset of the other. The path in this case had the shape schematically indicated in Figure 4e by the area that (only for purposes of clear black and white reproduction here) is stippled and much darker than the very light, brief, uniform grey of the path actually presented.

At a random time, while the appropriate induced motion was being experienced, the symmetry-breaking extension (or, alternatively, the faint guiding path) was deleted from the cycling display of the two rectangular bars. Under optimal conditions, viewers typically continued for a few cycles to experience the kinematically simple motion that had been primed by the preceding extensions (or guiding path) before reverting to the experience of either the two most favored default motions, namely, the pure translation indicated in Figure 4a or the pure depth rotation indicated in Figure 4b (see Carlton & Shepard 1990b, p. 220).

The reason that the translation and depth rotation were favored over the rotation in the picture plane (even though both rotations were through the geometrically equivalent 180° angles) is presumably that transformations of the former two types would be more consistent with the retinally available information. For an extended bar, the absence of retinal excitation along any possible connecting motion is




Figure 4. Pairs of alternately presented rectangular bars with L-like extensions that prime four types of apparent motion: (a) rectilinear translation, (b) rotation in depth about a vertical axis, (c) rotation in the picture plane, and (d) a horizontal screw displacement about an axis through the two bars. Pairs of bars with briefly presented interstimulus guiding paths that induce two types of apparent motion: (e) rotation in the picture plane and (f) an up-and-down translation over an inverted V path. (From unpublished experiments by Shepard and Zare.)

less consistent with a rotation in the picture plane (Fig. 4c), for which an actual motion would have tended to stimulate fresh retinal receptors along the path, and more consistent with a translation or a rotation in depth (Figs. 4a and 4b), for which the two presented positions of the bars extensively overlap the path of motion. Even if a fleeting motion had actually occurred over the path corresponding to a rotation in the plane, the resulting weak excitations along the path would have been largely masked by the more forceful retinal “burning-in” of the bar in its more enduring end positions.

1.7. Conditions revealing the default paths of mental kinematics

As I have already noted, natural selection has favored neuronal machinery for swiftly representing whatever motion is actually taking place in the world – not just for representing simple screw displacements. But, to perceive geometrically more complex motions that depart from the default paths of transformation, two conditions must be met: the proximal information must unambiguously specify a more complex distal motion, and the information must impinge on the sensory surface at a rate that does not outstrip the rates of propagation and processing of the neuronal system behind that surface (a system that evolved in a pretechnological world in which most biologically relevant motions were presumably of relatively limited velocity).

Even apparent motion can be induced over a path that does not correspond to a kinematically simple screw displacement. Under appropriate conditions, brief interstimulus presentation of the path schematically illustrated in Figure 4f, for example, can induce a nonrotational experience of the bar translating upward, reversing, and translating back downward in a bouncing inverted-V trajectory between the left and right bar positions. But when the rate of alternation is increased just to the point where the interval between stimulus onsets (the stimulus-onset asynchrony or SOA) becomes too brief for the internal enactment of this kinematically complex motion, the experience tends to revert to the rigid rotation in the plane corresponding to the path depicted in Figure 4e. Presumably, this simple rotation is favored at the shorter SOA because it is the only default motion for which the presented path (Fig. 4f) provides approximate – although not perfect – support. With further reduction of the SOA (or with deletion of the guiding path), the motion usually reverts, once again, either to pure translation or to pure rotation in depth.

From the standpoint advocated by Gibson (1979), apparent motion may seem lacking in ecological validity in a world in which material objects do not go discontinuously in and out of existence. Yet, even in a natural environment, significant objects may be only intermittently visible – as when they are behind wind-blown foliage, for example. One’s life can then depend on whether two fleeting visual sensations are interpreted as a single predator moving left to right, or as two distinct objects, one stationed on the left and one stationed on the right. In the laboratory, moreover, the default motions that are experienced in the absence of external support are just the ones that reveal, in their most pristine form, the internalized kinematics of the mind and, hence, provide for the possibility of an invariant psychological law.

1.8. The emergence of invariant laws in representational space

Under appropriate conditions, the minimum time required for representation of a rigid motion between two positions of a stimulus has characteristically increased in an essentially linear manner with the magnitude of the spatial disparity of those positions. Thus, in the case of visual apparent motion, the SOA yielding the experience of a rigid transformation over a connecting path increases approximately linearly with the linear separation between the alternately presented stimuli (e.g., Corbin 1942 [see Shepard 1984, Fig. 5]; Miller & Shepard 1993) or, when the stimuli differ in orientation, with the angular difference between them (Shepard & Judd 1976). Similarly, in the case of mental rotation, the time required to determine whether two objects are identical in shape (as opposed to enantiomorphic) increases approximately linearly with the angular difference in their orientations (see, e.g., Cooper 1975; 1976; Shepard & Metzler 1971). (For overviews of many of the results that have been obtained both for apparent and imagined motion, see, e.g., Cooper & Shepard 1984; Shepard & Cooper 1982; and for an overview of a related phenomenon of “representational momentum,” see Freyd 1983.)

Several facts indicate that the slopes of these linear increases of time with distance are not determined by characteristic speeds with which corresponding objects move in the world. There do not seem to be any well-defined characteristic speeds: a bird may perch on a limb or swoop past, a stone may rest on the ground or be hurled. The apparent motion of an object can be experienced before the object itself has been identified as a type likely to move quickly or not at all (e.g., a mouse versus a stone). An object's velocity relative to the observer must, in any case, depend on the observer's own motion. Finally, the obtained slopes of the chronometric functions have generally depended much more on the type of task than on the type of objects presented, with fastest transformational rates found for apparent motion (Shepard & Judd 1976), slower rates for mental rotation (Shepard & Cooper 1982; Shepard & Metzler 1971), and, within mental rotation tasks, slowest rates when two externally presented objects are to be compared (Shepard & Metzler 1971) or when the objects are unfamiliar (Bethell-Fox & Shepard 1988), rather than when an externally presented object is to be compared with an internally represented, already well learned canonical object (Cooper 1975; 1976; Shepard & Metzler 1988).

Again, invariant laws require formulation in terms of more abstract regularities in the world. Neither the path over which an apparent motion is experienced nor the critical time required for the traversal of that path suggests a concrete simulation of the physically or biologically most probable motion of that particular object in that particular circumstance (Shepard 1984). Rather, natural selection seems to have favored the establishment of the identity (or nonidentity) of the two objects in the fastest possible way that preserves whatever is invariant in the structure of the object. Evidently, the fastest possible way for objects in three-dimensional space is via the simplest transformation permitted by the corresponding kinematic geometry of that space. Differences among the rates estimated in the different tasks may not so much reflect differences in typical behaviors of the objects presented as differences in the de-

mands on and external supports for internal computations in those tasks.

The formulation of an invariant chronometric law of linear increase of time with distance requires, of course, that we choose the psychologically appropriate definition of distance. Both for imagined transformation (Shepard & Metzler 1971) and for apparent motion (e.g., Attneave & Block 1973; Corbin 1942; Ogasawara 1936; Shepard & Judd 1976), the appropriate distance evidently is the extent of the relevant transformation in the three-dimensional world more than any distance on the two-dimensional retina. Moreover, invariance is not achieved by defining distance solely in terms of the two objects between which a rigid motion is to be imagined or experienced. Invariance can only be achieved relative to the particular path of motion mentally traversed or experienced on a given occasion, for example, out of all alternative paths that are also permitted by the symmetries of the particular object presented (Farrell & Shepard 1981; Metzler & Shepard 1974 [Fig. 16]; Shepard & Zare 1983).

An invariant chronometric law finally becomes possible when critical times are related to distances along the appropriate geodesic paths in the appropriate representational space. The rate of traversal of such a path is not invariant across different tasks, because natural selection has favored neuronal machinery that yields the fastest possible computation given the external support available, but the external support varies from situation to situation. Even for the same task, the rate is not invariant across different geodesic paths, because no global metric (but only what is called the connection – see Carlton & Shepard 1990a) can be established for the full space of possible positions. (In terms of the formal structure of kinematic geometry, this can be understood by considering that any finite rotation, however small, must dominate any finite translation, however large, because any finite translation is abstractly equivalent to an infinitesimal rotation about a “center at infinity” – Carlton & Shepard 1990a.) For any one given path in the space of possible positions, the linearity of transformation time nevertheless becomes an invariant by virtue of the additive nature of times of analog traversal through successive points along that geodesic. I turn now to a formal characterization of the abstract representational space of possible positions and the geodesics that I take to represent the default paths of apparent or imagined motions. Such a characterization is best developed, first, for the case of an idealized asymmetric object and, then, for the cases of an object's possessing or approximating various symmetries.

1.9. The manifold of positions of asymmetric objects, and its geodesics

Objects in three-dimensional space have three degrees of freedom of translation and, except for surfaces of revolution (such as a perfect cylinder, which has an axis of complete rotational symmetry), three additional degrees of freedom of rotation. The complete specification of the position of an asymmetric object at any given moment requires, therefore, the specification of six independent quantities, three for its location and three for its orientation. (Specification of the orientation of a rotationally symmetric ideal cylinder, in contrast, requires only two quantities rather than three, because all angular orientations about its central axis are in-

distinguishable.) Any rigid motion of an asymmetric object over time thus corresponds to the traversal of a one-dimensional path in an abstract six-dimensional space of the object's distinguishable positions. Moreover, because rotation of any object through 360° returns it to its original position, the three dimensions of orientation are all circular. The abstract six-dimensional space as a whole is accordingly curved and non-Euclidean.

Despite its globally curved, non-Euclidean structure, this six-dimensional space is approximately Euclidean in each local neighborhood – much as the surface of the earth, although globally spherical, approximates a flat Euclidean plane within each sufficiently small region (corresponding, for example, to a single state or country). Spaces that thus approximate Euclidean space in each local neighborhood but may have a globally curved structure are called manifolds. The six-dimensional manifold of object positions has a particular mathematical structure (called, again, its connection) such that the paths in the manifold prescribed by kinematic geometry are the geodesics – the analogs, for a curved space, of straight lines in Euclidean space. (For successive stages in the development of these ideas in connection with the perceptual representation of positions and motions of objects, see Foster 1975b; Shepard 1981b; 1984; Shepard & Farrel 1985; and, most fully, Carlton & Shepard 1990a; 1990b.)

Geodesics are the one-dimensional curves in a manifold that are most simple and uniform in that, like straight lines in Euclidean space, the entire curve can be generated by iteratively applying the same local translational operation that carries any point on the curve into another nearby point on the curve, thus extending the curve in the most natural way. For the geodesics on the surface of a sphere (the *great circles*), for example, a step in the direction that takes one from a point to a nearby point on the geodesic will, with sufficient iteration, take one clear around the circle; equivalently, a straight tape smoothly applied to the surface in the local direction of the curve at any point will eventually return to that starting point, having covered the entire great circle.

As a reflection of the intimate connection between positions and motions that I mentioned at the outset, the set of distinguishable positions of an asymmetric object and the set of rigid displacements of such an object are representable by the same manifold. Once we have selected any one position of an asymmetric object as its canonical reference position, application of any screw displacement (whose rotational component does not exceed 360°) will carry the object into a unique position, and every possible position can be obtained in this way. The correspondence between distinguishable positions and screw displacements is not strictly one-to-one, however. As already remarked, for two objects differing only in orientation, there are two distinct rotations, which will carry one into the other around complementary segments of the geodesic circle. I shall soon return to the consequences of this for the structure of the manifold.

1.10. Formal characterization in terms of group theory

The structure of the set of positions of an object, the set of rigid displacements of the object, and the corresponding manifold with its geodesics can be elegantly formulated in terms of group theory. A group is a set of elements, which

in the present case would correspond to rigid displacements of an asymmetric object in space, that meet the following four conditions:

1.10.1. Closure. To any ordered pair of elements from the set there is a uniquely corresponding single element, called their *product*, that is also a member of the set. (Thus, for the two screw-displacement transformations, T_1 and T_2 , there is a single such transformation, T_3 , that carries the object to the same position as the transformation T_1 followed by the transformation T_2 : $T_1 \cdot T_2 = T_3$.)

1.10.2. Associativity. An ordered subset of three elements corresponds to the same product element whether a partial product is first formed from the first two elements or from the last two elements, before forming a final product with the remaining element. (Thus, for the ordered set of transformations, T_1 , T_2 , and T_3 : $[T_1 \cdot T_2] \cdot T_3 = T_1 \cdot [T_2 \cdot T_3]$.)

1.10.3. Existence of identity element. The set of elements contains a unique element whose product with any given element is just that given element. (Thus the degenerate transformation, here denoted I , that leaves the position of an object unchanged has no effect beyond the effect of any given transformation, T_1 , that it precedes or follows: $T_1 \cdot I = I \cdot T_1 = T_1$.)

1.10.4. Existence of inverse. For every element in the set, there is a unique element in the set, called its inverse, such that the product of the element and its inverse is the identity element. (Thus, for every transformation, T_1 , there is a compensating inverse transformation, T_1' , that restores the object to its initial position: $T_1 \cdot T_1' = I$.)

A familiar example of a group is the set of integers under addition. The group-theoretic “product” in this case is simply the (algebraic) sum of any two integers. Clearly, we have associativity: $(a + b) + c = a + (b + c)$; an identity element (zero); and an inverse for any element n (namely, the integer $-n$). As already implied, the set of elements of a group has dual interpretations – as the set of operations (e.g., the set of continuous displacements in space, or the set of discrete displacements along the number line by addition of positive or negative integers), or as the set of objects obtainable from a canonical element by those operations (e.g., the set of positions of an object in space obtainable by rigid displacements from a reference position, or the set of integers obtainable by integer shifts from – i.e., algebraic additions to – a reference integer, such as zero).

The relevant group for the representation of distinguishable positions or rigid displacements of an asymmetric object in three-dimensional Euclidean space is the Euclidean group, E^+ . (The “+” is used here to indicate the restriction to rigid transformations confined within the three-dimensional space, thus excluding reflections between enantiomorphic shapes, such as a left and right hand, that could otherwise be obtained by rigid rotation through a higher-dimensional embedding space.) Because a general screw displacement includes a translational and rotational component, the Euclidean group is composed of the group of linear translations and the group of orthogonal rotations. In group-theoretic terms (see Carlton & Shepard 1990a), E^+ is expressible as the *semidirect product* of the three-dimensional translation group, \mathbb{R}^3 , and the three-dimensional rotation group $SO(3)$:

$$E^+ = \mathbb{R}^3 \otimes SO(3) \quad (1)$$

The manifold of distinguishable positions (or, equivalently, rigid displacements) of an asymmetric object in three-dimensional space is isomorphic to the Euclidean group, E^+ .

The concept of the product of two groups may be clarified by considering the simpler product of the group of rigid translations along a line (or, in the discrete case, the group of integers under addition) and the group of rigid rotations about a circle (or, in the discrete case, the group of positive clock-face integers or months of the year 1 through 12, modulo 12). Each of the elements of the *direct product* of these two groups is composed of one element from each of the two component groups (where either element can be the identity element). The direct product of such a rectilinear and circular group is, naturally enough, a cylindrical group. Elements of such a group, by virtue of their rectilinear and circular components, can take us from any point on the surface of the cylinder to any other. In such a direct product group, the elements are *commutative*, that is, the product of two elements is independent of their order so that from a given point on the surface of a cylinder, we get to a given other point whether we first translate the appropriate distance parallel to the axis of the cylinder and then rotate through the appropriate angle about that axis, or whether we first rotate through that angle and then translate over that distance.

In the case of a semidirect product group, however, not all elements will commute in this way. The Euclidean group is necessarily a semidirect product group because rotations in three-dimensional space are generally noncommutative: for an asymmetric shape such as σ (the letter b tipped 90° clockwise in the picture plane), a 90° clockwise rotation followed by a 180° rotation around a horizontal axis yields the result “d” while the same rotations performed in the reverse order – first 180° around a horizontal axis followed by a 90° clockwise rotation in the picture plane – yields the different result “p.” (A more complete account of semidirect products is provided by Carlton & Shepard 1990a.)

Each subgroup of a group, such as the subgroup of pure translations \mathbb{R}^3 and the subgroup of pure rotations $SO(3)$ of the Euclidean group E^+ , individually satisfies the already-stated conditions for a group. The Euclidean group also contains other, more restricted subgroups such as the group of translations along a horizontal axis of three-dimensional space, or the group of rotations about a vertical axis of that space.

Of greatest relevance, here, is the set of *one-parameter subgroups* of the Euclidean group. These correspond to the geodesics in the manifold of distinguishable positions, and are straight lines in the three-dimensional translation subgroup, \mathbb{R}^3 , circles in the three-dimensional rotation subgroup $SO(3)$ and, more generally, helical curves in the full six-dimensional Euclidean group E^+ . (The designation of these subgroups as “one-parameter” corresponds to the fact that a single parameter suffices to specify a location along a one-dimensional geodesic.) In an analogous but more easily imagined, lower-dimensional, and direct-product case, a tape started at an arbitrary angle will wind helically around the surface of a cylinder, which also has a straight (axial) component, the analog of \mathbb{R}^3 , and circular (angular) component, the analog of $SO(3)$.

For pure rotations of an object in space, we need con-




Figure 5. Flattened depiction of one two-dimensional section through the three-dimensional manifold, $SO(3)$, of orientations of a marked cube. From “Representation of the Orientations of Shapes,” by R. N. Shepard & J. E. Farrell, 1985, *Acta Psychologica* 59:109. Copyright 1985 by Elsevier Science Publishers. Reproduced by permission.

sider only the great-circle geodesics in the three-dimensional submanifold corresponding to $SO(3)$. Figure 5 illustrates, by means of the orientations of a labeled cube, a two-dimensional section through this submanifold. The portrayal of this submanifold as a flat disk is only for convenience of illustration in a flat picture. The intrinsic metric of this two-dimensional submanifold is actually that of a spherical surface, thus providing for the great-circle shapes of the geodesics (see Carlton & Shepard 1990a). Moreover, diametrically opposite points around the perimeter of the disk correspond to the same orientation of the object (as shown in the figure by agreement in orientations of the letter B on the back of the cube) and such pairs of points, although widely separated in the figure, should be regarded as the same points.

We are now in a position to clarify further the relation between the spatial representation of distinguishable positions of an asymmetric object and the representation of its rigid displacements. The hemispherical surface illustrated, in flattened form, in Figure 5 includes points corresponding to rotations of only up to 180° from the orientation of the cube represented by the central point (with the F-marked face upright and in front) taken as its canonical orientation. This is sufficient for the representation of all distinguishable orientations falling on geodesics in this surface because, for every rotation through more than 180° (the longer way around a geodesic circle in this surface), there is a rotation through less than 180° (the shorter way around that circle) that is included in the surface and that results in exactly the same orientation of the object. So, although the two possible transformations (the longer and shorter ways around the circle) are distinct, the results of these two transformations are identical. For the complete representation of the three-dimensional subgroup of distinct rota-

tions, $SO(3)$, then, each two-dimensional hemispherical section, such as that illustrated in Figure 5, must have its missing half added, to form a complete sphere. In the complete manifold of rotations, then, diametrically opposite points correspond to distinct rotations (the shorter or longer ways around the same geodesic circle) but, in the corresponding manifold of distinguishable positions, such diametrically opposite points, because they correspond to indistinguishable orientations, are identified (treated as the same point). (Counterintuitively for us, who have evolved to deal with macroscopic objects, such an identification is not needed for an important class of microscopic objects, viz., fermions, which include such basic constituents of matter as electrons and protons. As was first called to my attention by Eddie Oshins, according to an empirically verified prediction of quantum mechanics, these particles do not become physically identical to themselves until rotated through two complete 360° turns!)

1.11. Formal characterization of positions and motions for a symmetric object

An object possessing one or more symmetries entails a modification of the manifold of distinguishable positions. By definition, whereas an asymmetric object becomes distinguishable from itself under a rotation through any non-zero angle short of 360° , a symmetric object becomes identical to itself under some other rigid transformation, such as a 180° rotation in the case of a rectangle. Consequently, some widely separated points of the manifold of distinguishable orientations for an asymmetric object (such as the points corresponding to 180° – different orientations of the asymmetric polygon in Fig. 6a) must be mapped onto

the same point of the manifold of distinguishable orientations of a symmetric object (such as the single point corresponding to any two 180° – different orientations of the centrally-symmetric polygon in Fig. 6c). As illustrated at the bottom of the figure, the great circle corresponding to one complete picture-plane rotation of the asymmetric polygon (Fig. 6a) is thus twisted (through the intermediate curve shown in Fig. 6b) into a double-wound circle (Fig. 6c) in which each pair of orientations of the polygon separated by 180° maps into the same point (Shepard 1981b; Shepard & Farrell 1985). One complete 360° rotation of a centrally symmetric object (like the polygon in Fig. 6c) is thus represented by two complete excursions around a geodesic circle in the space of distinguishable positions of that object (the circle depicted at the bottom of Fig. 6c).

For an object possessing a symmetry, the submanifold of orientations is necessarily replaced by a *quotient manifold*. Designating these manifolds by the names of their corresponding groups, we can more specifically say that the manifold

$$SO(3) \text{ is replaced by } SO(3)/S(O) \quad (2)$$

where $S(O)$ is the manifold corresponding to the symmetry group of the object (see Carlton & Shepard 1990b). The symmetry group of the object is, simply, the subgroup of rigid transformations that leaves the object indistinguishable from its initial state. Thus the symmetry group of the square is a subgroup of the Euclidean group that includes rotations through 90° and 180° in the plane, as well as 180° rotations in space about vertical, horizontal, and diagonal axes of the square.

Quantitative evidence from a number of experiments (including experiments on real and merely imagined motion, as well as experiments on visual apparent motion) now indicates that psychologically preferred paths of rigid transformation do correspond to geodesics in the appropriate manifold – including the appropriate quotient manifold $SO(3)/S(O)$ for objects with various symmetries (e.g., Farrell & Shepard 1981; Shepard 1981b; Shiffrar & Shepard 1991; see also Carlton & Shepard 1990b, pp. 219-21). As will be noted, such manifolds and geodesics can even be recovered by applying methods of multidimensional scaling to psychological data.

1.12. Formal characterization for approximations to various symmetries

Most of the objects that we encounter in the world are neither completely asymmetric nor exactly symmetric. Instead, they more or less approximate various global or merely local symmetries. Just as a strict symmetry of an object corresponds to the transformation (rotation or reflection) that carries that object exactly into itself, a symmetry that is only approximate corresponds to the transformation that achieves a local maximum of correlation in shape between the object and itself – with the degree of approximation measured by the magnitude of the correlation at that local maximum. Only a perfect sphere is identical to itself under every rotation and reflection about its center and, hence, is wholly symmetric. (Thus there is a more abstract, purely geometrical basis of the spherical shape of the soap bubble invoked by the Gestalt psychologists.) A person's face, body, and brain only approximate but do not achieve strict bilateral symmetry. Complete asymmetry, on the other hand, can

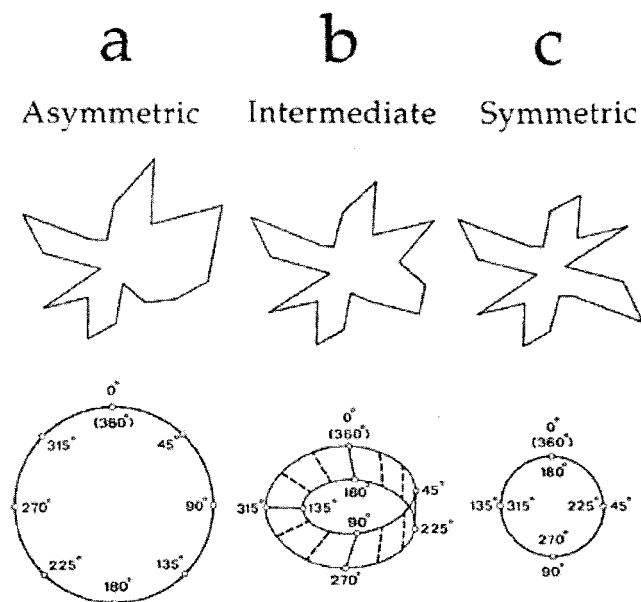


Figure 6. Illustrative polygons (above) and their corresponding geodesic paths of rotation in the picture plane (below) for three degrees of approximation to central symmetry: (a) 0%, (b) 75%, and (c) 100%. From "Psychophysical Complementarity," by R. N. Shepard, 1981b. In: M. Kubovy & J. Pomerantz (Eds.), *Perceptual Organization*, p. 317. Hillsdale, NJ: Lawrence Erlbaum. Copyright 1981 by Lawrence Erlbaum Associates. Adapted by permission.

never be attained. Any shape (including the “random” polygon in Fig. 6a) necessarily resembles itself to greater or lesser degrees under various angles of relative rotation. (Hence, the perfect circle depicted below that polygon does not precisely correspond to the internal representation of the set of distinguishable orientations of that particular polygon. Strictly, that circle is a kind of average representation of possible orientations for a total ensemble of polygons generated by the same random rules.)

Indeed, the shape of any particular object can be defined in terms of its degrees of approximation to all possible symmetries of that object, via the correlation between the object and itself under each possible rotation and reflection (Shepard 1981b; 1988). Although degrees of approximation to symmetries thus appear to be fundamental in human perception and cognition, the classical development of the group-theoretic basis of symmetry in mathematics has treated each type of symmetry as a discrete feature that an object possesses either wholly or not at all.

A formal, quantitative treatment of approximations to symmetries can, however, be given in terms of representational space (Carlton & Shepard 1990b; Farrell & Shepard 1981; Shepard 1981b; 1988; Shepard & Farrell 1985). Approximations to symmetries are regarded as inducing deformations in the original manifold of distinguishable orientations corresponding to $SO(3)$, for an ideally asymmetric object (or ensemble of random objects), toward the manifold, corresponding to $SO(3)/SO$, for each type of symmetry that a given object approximates.

Farrell and I sought empirical support for such a spatial representation of the orientations of polygons possessing various degrees of approximation to central symmetry, that is, to self-identity under 180° rotation in the two-dimensional plane. The minimum SOAs for the experience of rigid rotational motion between two alternately presented orientations (Farrell & Shepard 1981) and the times required for the discrimination of sameness or difference of two simultaneously presented orientations (Shepard & Farrell 1985) were both consistent with the representations of these shapes in their corresponding manifold of distinguishable positions (see Carlton & Shepard 1990b; Shepard 1981b). Specifically, multidimensional scaling of the discrimination times (using the INDSCAL method of Carroll & Chang 1970) yielded points in four-dimensional space falling close to the particular geodesics prescribed (Shepard & Farrell 1985), namely, closed curves forming the edge of a one-sided Möbius band (as illustrated in two-dimensional projection at the bottom of Fig. 6b); and the minimum SOAs for rigid apparent motion were as predicted for motions between the two alternately presented orientations over just these geodesic paths (Farrell & Shepard 1981).

1.13. Formal connections between the representations of positions, motions, and shapes

That the same manifold can represent both the distinguishable positions of an object in space and the possible rigid displacements of the object between its distinguishable positions holds also for objects that approximate various symmetries. This is the basis of the inextricable connection noted between the representations of the positions and kinematically simplest motions of an object. Shapes, however, can have many more than six degrees of freedom.

Clearly then, shapes cannot be fully represented as individual points in the manifold of positions/displacements, a manifold that has no more than six dimensions (and fewer, for objects, such as a cylinder or a sphere, with complete rotational symmetry). An isomorphism does, nevertheless, hold between the shape of any object and the conformation of the corresponding manifold of positions/displacements for that object. The conformation is dictated by the object's degrees of approximation to all possible symmetries (Shepard 1981b; 1988).

Ultimately, shapes themselves should be formally representable as points in a higher-dimensional manifold of all possible shapes. The full development of such a representation must provide for a detailed, parametric characterization of the degrees of approximation of a shape to possible symmetries in three-dimensional Euclidean space. Just as the position of any given object can be represented, historically, as the result of the simplest rigid transformation that might have carried the object into its given position from a prespecified canonical position, the shape of any given object might be interpreted, historically, as the result of the simplest nonrigid deformation that might have brought the object into its present shape from some prespecified, simplest canonical shape.

Leyton (1992) has achieved significant progress toward a group-theoretic account of how objects may be perceived and represented in terms of the derivational history that each implies. In the spirit of the approach I have outlined here, the appropriate representational space might provide, in general, for the interpretation of any object as having emerged from some more symmetrical, canonical progenitor through the traversal of a symmetry-breaking geodesic in that space. Unlike the manifold of positions and rigid motions, the space of possible shapes and nonrigid motions would be not only higher-dimensional but also anisotropic and inhomogeneous. In a possible, if remote analogy with general relativity, for which a test particle follows a geodesic toward a gravitational singularity in the space-time manifold, the cognitive interpretation of a given shape might be regarded as following a geodesic backward toward a point of maximum symmetry (and, perhaps, minimum entropy) in the manifold of possible shapes.

2. Representation of an object's color

The problem of color has inspired major efforts by some of the greatest scientists of all times, including Newton, Young, Helmholtz, Maxwell, and Schrödinger, to name just a few of the most illustrious physicists. So much attention to color might seem difficult to justify from an evolutionary standpoint. The perception and representation of the positions and motions of objects in space is clearly essential for our survival and reproduction. But the perception and representation of colors, though doubtless contributing to our discrimination of some biologically relevant objects (such as red berries against green leaves) and our recognition of, or learned attachment to, others (such as a face with blue, green, or brown eyes, or surrounded by yellow, red, brown, or black hair), evidently is not essential for many animals, including humans.

Originally, the investigation of color was probably motivated, instead, by the challenge of reconciling the seemingly unanalyzable subjective experience of colors with

such seemingly colorless concepts of physical science as space, time, particles, or waves – including the electrochemical events in what has aptly been styled “the dark chamber of the skull” (as by B. P. Browne, quoted in William James, 1890 [p. 220 of 1950 edition]). The challenge remains (Shepard 1993), and is even augmented by the need to provide an evolutionary explanation for the ways in which the internal representation of colors differs from the physical characteristics of external surfaces and of the electromagnetic radiations that they reflect in the world.

In this regard, two facts about the human perception of an object’s color are perhaps most fundamental: first, the color appearance of an object’s surface is essentially invariant despite enormous variations in the spectral composition of the light that falls on that surface and, hence, of the light that the surface scatters back to our eyes. Second, although these physical variations are also potentially of high dimensionality, we can match the color appearance of any such surface by adjusting just three chromatic components produced by a suitable color mixing apparatus. The color appearances of surfaces thus correspond to relatively fixed points in a three-dimensional *color space*.

Schematically, this color space can be thought of as approximating the idealized spherical solid portrayed in Figure 7. We can describe this space either in terms of three cylindrical coordinates of lightness, hue, and saturation (as shown in Fig. 7a), or in terms of three rectangular coordinates of lightness-versus-darkness, redness-versus-greenness, and blueness-versus-yellowness (as indicated in Fig. 7b). But *what in the world* is the source of the three-dimensionality of this color representation? And *what in the world* is the source, in this representation, of the circularity, discovered by Newton, in the continuum of hues? For, this circularity presents us with the psychophysical puzzle that the hues corresponding to the most widely separated of the visible physical wavelengths, namely red

and violet, appear more similar to each other than they do to a hue of intermediate wavelength, such as green.)

As before, the answers that may first spring to mind may not necessarily be the correct ones. In the case of motion, the most deeply internalized constraints evidently are determined less by Newtonian mechanics and the mass distribution of each object than by the more abstract kinematic geometry of three-dimensional Euclidean space and the symmetry groups of objects. Similarly in the case of color, I suggest that the three-dimensionality and circular structure of the representation derives less from anything in the intrinsic spectral characteristics of surfaces or of their reflected light than from the more abstract constraints of the universally linear way in which illumination from an invariant stellar source is transformed by a planetary environment, and by the prevailing three-dimensional structure of the planetary transformations. I begin with a consideration of the universal linearity of spectral transformation and the selective pressure toward its internal representation in diurnal animals with highly developed visual systems.

2.1. Formal characterization of the linearity underlying color constancy

The invariant physical characteristic of a surface underlying its perceived color is its spectral reflectance function $S^x(\lambda)$. This function specifies for each wavelength of incident light the fraction of that light that will be scattered back to any receptive eyes. Accordingly, the amount of light reaching an eye from a point x on an environmental surface, $P^x(\lambda)$, is expressible as the product of the amount of light of that wavelength in the ambient illumination, $E(\lambda)$, and the spectral reflectance of the surface for that wavelength at point x , $S^x(\lambda)$:

$$P^x(\lambda) = E(\lambda) S^x(\lambda) \quad (3)$$




Figure 7. Schematic illustrations of human color space showing (a) its three cylindrical dimensions of lightness, saturation, and hue, and (b) its three opponent-process rectangular dimensions of light-dark, red-green, and blue-yellow. From “The perceptual organization of colors: An adaptation to the regularities of the terrestrial world?” by R. N. Shepard, 1992. In: J. Barkow, L. Cosmides, and J. Tooby (Eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, pp. 467–97. New York: Oxford University Press. Copyright 1992 by Oxford University Press. Adapted by permission.

Empirical data and theoretical considerations (concerning universal quantum mechanical interactions between photons and surface molecules), reviewed by Maloney (1986), indicate that the spectral reflectance functions S_x of wavelength for natural surfaces can be approximated as linear combinations of a small number, n , of *reflectance basis functions*:

$$S^x(\lambda) = \sum_{j=1}^n \sigma_j^x S_j(\lambda), \tag{4a}$$

where σ_j^x is the weight for the j th of the n basis reflectance functions for surface point x (see also Brill 1978; Buchsbaum 1980; Sälström 1973).

Other empirical data and theoretical arguments (to

which I shall return) indicate that the spectral distributions E for natural conditions of illumination can similarly be approximated as linear combinations of a small number, m , of *lighting basis functions*:

$$E(\lambda) = \sum_{i=1}^m \epsilon_i E_i(\lambda), \tag{4b}$$

where ϵ_i is the weight for the i th of the m basis lighting functions for the ambient illumination (see Maloney & Wandell 1986).

Substitution of Equations 4a and 4b into Equation 3 then yields a dimensionally reduced linear model governing the way illumination and surface properties are combined in the proximal stimulus P_x . Figure 8 is my schematic illustra-




Figure 8. Illustration, for two conditions of terrestrial filtering (T1 and T2), of how the spectral composition of the light from an unvarying source, U, is linearly transformed first by terrestrial filtering, T, and then by scattering from a surface, S, before reaching the observing receptor, R.

tion of how the spectral composition of the light scattered to the eye from a surface (here, a green leaf) differs between two conditions of terrestrial filtering of the illumination, in which a cloud moves to block the longer-wavelength (redder) rays direct from a setting sun (T1), or to block the shorter-wavelength (bluer) rays scattered from the molecules of air (T2). In either case, the distribution of the unvarying solar light (U) is linearly transformed (UT) by the spectral distribution of the momentary terrestrial condition of filtering (T1 or T2), and that light is then linearly transformed again (UTS) by the spectral reflectance distribution of the surface (S) in the process of being scattered to the observing receptor (R). In order to achieve color constancy, the visual system must arrive at a correct characterization of the invariant spectral reflectance distribution $S(\lambda)$ of the external surface (of the leaf) despite the contamination of the spectral distribution of the proximal stimulus (UTS) by the terrestrial filtering $T(\lambda)$ of the illumination.

Using such a linear model, Maloney and Wandell (1986) showed how (under quite general conditions) a visual system with a sufficient number of chromatically distinct types of photoreceptors (such as the red-, green-, and blue-sensitive cones in the human retina) can achieve a disentanglement of the characteristics of the surface (S) and the characteristics of the illumination (UT) and thus attain color constancy. Because the linearity of the way in which the spectral properties of illumination and surface combine is universal, this linearity should tend to be internalized in the visual systems of organisms wherever natural selection has favored color vision. But we are still left with the question of what it is in the world that determines the dimensionality of color representation.

2.2. The representation of surface colors should have the dimensionality of natural illumination

According to Maloney (1986), the number of degrees of freedom of spectral reflectance of natural surfaces (n in Eq. 4a) falls somewhere between five and eight. A visual system that completely recovers the chromatic characteristics of such surfaces by means of the computation described by Maloney and Wandell (1986) would require a number of chromatically distinct types of photoreceptors that is one greater than this number of degrees of freedom, that is, between six and nine. The conclusion is clear: a visual system, like ours, that has only three types of color receptors (the red, green, and blue cones) and, hence, that is restricted to three dimensions of color representation cannot fully capture the intrinsic reflectance properties of natural surfaces.

Suppose, however, that the dimensionality of color representation has been favored not because it captures the full spectral reflectance properties of natural surfaces but because it is the minimum dimensionality needed to compensate for natural variations in illumination and, thus, to achieve *constancy* of whatever chromatic aspects of the surfaces are represented. Then, even though we may not perceive everything that could be perceived about each surface, we at least perceive each surface as the same under all naturally occurring conditions of illumination.

Available evidence indicates that the number of degrees of freedom of terrestrial lighting (m in Eq. 4b) is essentially three. Principal components analyses have revealed that the great variety of spectral energy distributions of natural illumination measured for different atmospheric conditions

and times of day can be well approximated as linear combinations of just three basic functions (see Judd et al. 1964; and other studies cited in Maloney & Wandell 1986, Note 17). Moreover, the three dimensions of spectral variation in natural illumination have identifiable sources in the world (Shepard 1992):

(1) There is an overall light-versus-dark variation ranging from the full and direct illumination by midday sun and unobstructed sky to whatever portion of that same illumination (uniformly reduced across all wavelengths) reaches an object only by scattering from achromatic clouds, cliffs, or moon.

(2) There is a red-versus-green variation depending on the balance between the long (red) wavelengths, which are selectively passed by atmospherically suspended particles (a particularly significant factor when the sun is close to the horizon) or which are selectively blocked by water vapor, and the remaining band of visible wavelengths, which (ranging from yellows through blues) center on green.

(3) There is a blue-versus-yellow variation depending on the balance between the short (blue and violet) wavelengths, which are selectively scattered (e.g., to a shaded object) by the molecules of the air itself, and the remaining band of visible wavelengths of light directly from the sun (as might reach an object through a small “window” in a leafy canopy), which (ranging from greens through reds) center on yellow.

Possibly, then, the light-dark, red-green, and blue-yellow *opponent processes*, proposed by Hering (1887/1964) and Hurvich and Jameson (1957), on quite different (psychophysical and, subsequently, neurophysiological) grounds, may not have to be accepted as an arbitrary design feature of the human visual system. Such a three-dimensional representation of color may have emerged as an adaptation to a pervasive and enduring feature of the world in which we have evolved. At the same time, this feature may be the nonarbitrary source of the transformation of the rectilinear continuum of physical wavelength into Newton's circle of hues: Very schematically, the two colors in either the red-green pair or the blue-yellow pair, corresponding to the two extremes of variation on an independent dimension of terrestrial filtering, are analogous to diagonally opposite corners of a square (see Shepard & Carroll 1966 [Fig. 6, p. 575]) or diametrically opposite points on a circle (as in Fig. 11 in the next section of the present article). As such, the two opposite colors in either of these pairs must be further apart than the colors in any other pair, including red and blue, which, although corresponding to points close to the extreme ends of the physical continuum of visible wavelengths, are perceptually represented in a way that is more analogous to points separated by one edge of a square or by only a quarter of a circle (Shepard 1992; 1993).

If the linear transformations of the illumination that occur in nature have just three degrees of freedom, then three dimensions are required to compensate for those transformations and, thus, to maintain constancy in the apparent colors. Indeed, three such dimensions are needed to maintain constancy even in just the apparent lightnesses of surfaces, without regard to chromatic color (Shepard 1990; 1992). That is, for every terrestrially induced linear transformation on the illumination, a compensating (inverse) transformation must be internally performed to achieve invariance of the final internal representation of the colors – including even their ordering with respect to achromatic

lightness. (Even in the purely achromatic, i.e., grey-scale, representation of, say, a red and a blue surface, the red surface would appear a lighter grey than the blue surface in the light of the setting sun but the blue surface would appear lighter grey than red in the light scattered only from the sky – unless the input is first analyzed into separate chromatic channels and appropriately transformed, before being reduced to the final grey-scale representation.)

Figure 9a indicates how one possible linear transformation (for simplicity of illustration, a two-dimensional transformation representable by a diagonal matrix) would affect the amounts of light of long and short wavelengths reflected back from each of a number of colored surfaces (indicated by the dots). These amounts might correspond to what would be picked up by red and blue cones in the human retina. Under a shift in natural illumination in which, for example, clouds that were blocking light directly from a low sun shift to block, instead, light scattered from the sky, the light scattered back to the eye from each surface (indicated by a filled circle) contains a reduced portion of its original short-wavelength (blue) component and increased portion of its original long-wavelength (red) component (as indicated by the arrow carrying the filled circle to the position of a corresponding open circle). An inverse linear transformation (mapping the rectangle with dashed outline back to

the square with solid outline) will reinstate the original configuration of dots and, hence, achieve constancy of appearance of the surface colors. (The chromatic information about the surfaces is not in any sense used up in correcting for the illumination. Because only a small subset of the many visible surfaces – the individual points in Figure 9a – is sufficient for making the correction, the correcting transformation still provides for the representation of the colors in the whole set in the low-dimensional representation.)

The general case of a linear transformation that is both three-dimensional and nondiagonal is more difficult to illustrate for the whole set of points representing individual surfaces, but a rough idea of such a transformation may be gained from the more schematic depiction of the more general linear transformation between a cube and a parallelepiped in Figure 9b.

The dimensionality required is the same regardless of the particular transformation used to approximate the optimally color-constant transformation. This transformation could have the simplest (diagonal) form of the transformation described by Land and McCann (1971). It could have the more color-constant general linear form of the transformation proposed by Maloney and Wandell (1986; see also the revised approach described by Marimont & Wandell 1992). Or it could have some still more sophisticated form that takes account of surface orientation, shading, and shadows (see, e.g., Sinha & Adelson 1993); specular reflections (from glossy surfaces – see, e.g., Tominaga & Wandell 1989); or even, when the geometry permits the inference that the light falling on the object is identical to the light reaching the eye directly from the visible source, spectral properties of the illumination itself.

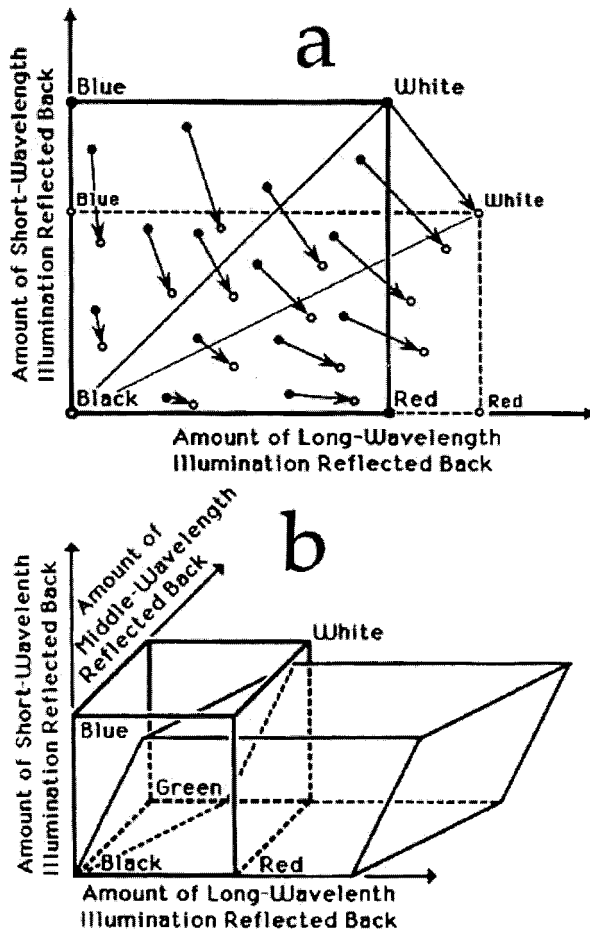


Figure 9. Schematic illustration of the effects of a terrestrial transformation on the amounts of light of different wavelengths scattered back to an eye: (a) for just two dimensions and a diagonal transformation, and (b) for a nondiagonal transformation in three dimensions. (See text for explanation.)

2.3. Formal characterization of the representation of invariant colors

As suggested by the preceding discussion of the spectral transformations of light by atmospheric filtering and surface reflection, and of the inverse transformations required to achieve invariance, these spectral transformations, like the rigid transformations of objects in space, constitute a mathematical group. Krantz (1975a; 1975b) has already presented an extensively developed group-theoretic formulation for the appearances of colored lights. From an evolutionary perspective, however, it was the invariant characteristics of light-reflecting objects – not the variable light or sources of light – that were of primary biological significance for the survival and reproduction of our ancestors in the pretechnological world. The linearity of the transformations of filtering and reflection ensures that the appropriate group for representing variations in the spectral composition of the light reaching our eyes from surrounding surfaces is a linear group, instead of the Euclidean group appropriate for rigid motions of objects in space. Of potential value, therefore, would be the further development of such a group-theoretic formulation of the representation of surface colors at a level of detail comparable to that provided in the group-theoretic representations of lights by Krantz (1975a; 1975b), of positions and motions by Carlton and Shepard (1990a; 1990b), of nonrigid deformations by Leyton (1992), or of musical intervals by Balzano (1980).

The formalization of the structures underlying psychological representation at a suitably abstract level can reveal deep analogies between disparate domains. In the domain

of color, just as in the domains of position, motion, deformation, and musical pitch, transformations have an abstract group-theoretic representation. Different domains require different groups, such as the Euclidean group for changes in position of an invariant shape, and the linear group for changes in the spectral composition of light reflected from an invariant surface. Nevertheless, they share some fundamental properties. In the representation of position or motion and in the representation of color, alike, the formal characterization reveals, for example, how prevailing structural constraints yield dimensional reduction of the representational space. Thus the symmetry group $S(O)$ of a surface of revolution, such as a cylinder or a sphere, entails, through substitution of the appropriate quotient manifold (Eq. 2), a reduction from a six-dimensional to a five- or a three-dimensional space of distinguishable positions, respectively (Carlton & Shepard 1990b). Similarly, a restriction on the degrees of freedom of terrestrial filtering permits a reduction in the dimensionality of the representation for surface colors, from a space of six or more dimensions needed to capture the full reflectance characteristic of the surfaces of natural objects, to the three-dimensional space sufficient for the minimal invariant representation of their intrinsic colors.

2.4. Generality of the principles of color representation

Adaptation to the degrees of freedom of natural illumination does not of course ensure color constancy under conditions departing from those that have prevailed during terrestrial evolution. Modern technology has produced spectrally unnatural light sources under which human vision may not be color constant – as demonstrated in the vision laboratory, or under mercury vapor street lamps at night (where our companions may take on a ghastly aspect or we may fail to recognize our own car). Nor is an essentially three-dimensional variation of natural daylight the only factor that can determine the dimensionality of a species' color space. For nocturnal or deep-sea animals, the sensitivity afforded by achromatic rod vision may outweigh the benefits of cone-based color constancy. Even for many diurnal animals (including new world monkeys and human dichromats), the gain in color constancy attainable by the addition of a third class of wavelength-selective cones may be only marginal. Finally, runaway sexual selection may lead not only to the evolution of uniquely colored markings or plumage but also to the emergence of additional classes of retinal receptors and dimensions of color space tuned to the representation of such colors (Shepard 1992).

Still, the converging evolution of three-dimensional color representation in diverse visually dependent animals – evidently including most humans as well as the birds and the bees – may not be accidental. The speculation that I have favored is that this three-dimensionality may be an adaptation to a property that has long prevailed on our planet. We may need three dimensions of color not because the surfaces of objects vary in just three dimensions but because we must compensate for the three degrees of freedom of natural lighting in order to see a given surface as having the same intrinsic color regardless of that illumination.

The reduction specifically to three dimensions of color, though justified here in terms of the variations of natural illumination prevailing on earth, may hold more generally. On any planet capable of supporting visually advanced

forms of life, illumination is likely to originate primarily from an essentially invariant stellar source. Moreover, the atmospheric and surface conditions necessary for such life are likely to provide only a limited spectral window of transmitted wavelengths of that light. Hence, the principal variations of the light reaching significant objects on or near the surface of such a planet are likely to be a variation in the overall quantity of that light and independent variations at the short-wavelength edge and the long-wavelength edge of the spectral window. Although additional, more subtle and spectrally selective variations may be to some extent present, these three variations – in the overall level of the transmitted light, and the extent of its reach into the short and the long wavelengths – seem likely to predominate in planetary environments generally and to exert the greatest influence through natural selection.

Whether or not my conjecture as to the nonarbitrary source and possible generality of the tendency toward three-dimensional color representation is ultimately supported, the universally linear way in which the spectral composition of light is transformed by scattering and filtering in the external world seems likely to have favored, wherever color vision has evolved, the internal implementation of compensating linear transformations on the proximal stimulus. Only in this way can significant external objects under any naturally varying illumination yield a color-constant internal representation, whatever its dimensionality may be.

3. Representation of an object's kind

The preceding examples concerned abilities to identify stimuli as distal objects that – despite wide variations in position and lighting – are nevertheless identical in intrinsic shape or color. My third and final example concerns an ability that does not require spatial or color vision and, hence, that is still more fundamental and ubiquitous. This is the ability to recognize that even when the distal objects themselves are not identical, they may nevertheless be objects of the same basic kind and, hence, likely to have the same significant consequences for the perceiver. For example, whether a newly encountered plant or animal is edible or poisonous depends on the hidden genetic makeup of the natural kind of that object.

Under the term *basic kind* I mean to subsume not only such *natural kinds* as animal, vegetable, and mineral species, but also such *basic level categories* (Rosch et al. 1976) as knife, bowl, or chair (for humans) or trail, burrow, or nest (for animals of some other species). Objects of the same basic kind are thus objects that provide the same functions or affordances (in the sense of Gibson 1979). A basic kind typically includes objects that, although more or less similar, may be readily discriminable from each other: an apple may be red or green; a trail may be level or steep; a chair may have a low or high back. Generalization from one object to another is not a failure of discrimination, therefore, but a cognitive act of deciding that two objects, even if readily distinguishable, may be similar enough to be of the same kind and, hence, to offer the same significant consequence or affordance.

This simple idea yields a quantitative explanation of a very general empirical regularity that is latent in generalization data of the sort that specifies, for all pairs of n stimuli, the probability that a response learned to one of the

stimuli in the pair will be made upon presentation of the other stimulus of that pair. The latent regularity emerges when such data are submitted to multidimensional scaling, a method that finds the unique mapping of objects or stimuli into a space of minimum dimensionality such that the data have an invariant monotonic relation to corresponding distances in that space (see Kruskal 1964a; Shepard 1962a; 1962b; 1980). The resulting *generalization gradients*, which describe how the probability of a response learned to one stimulus falls off with the distance from it of each other stimulus in the obtained spatial representation, have uniformly approximated a function of simple exponential decay form. See Figures 10a and 10b, respectively, for generalization gradients that I obtained in this way (Shepard 1962b; 1965) for spectrally pure colors (hues) based on generalization data from pigeons (Guttman & Kalish 1956) and on a related type of similarity data from humans (Ekman 1954). The smooth curves are simple exponential decay functions fitted to the points by adjustment of a single parameter, the slope constant (which can equivalently be regarded as a scaling factor for the distances in each spatial representation). (For a number of gradients of generalization obtained in this way for other visual and auditory stimuli, see Shepard 1987b; and, for confirmation that the shape of the obtained functions is determined by the data and not by the multidimensional scaling method, see Shepard 1962b; 1965.)

Figure 11 displays the points corresponding to the spectral colors in the configuration (obtained from the human data) that yielded the approximately exponential function shown in Figure 10b. Four observations concerning this configuration are relevant here. First, as the close fit to the subsequently drawn circle indicates, the obtained configuration of points closely approximates Newton's color circle – the equatorial great circle of hues schematized in the earlier Figure 7a. Second, the implied psychophysical mapping from the rectilinear continuum of physical wavelengths to the circular continuum of perceived hues emerges as a consequence of my requirement that the law of gener-

alization be not only invariant but monotonic – a requirement that was met, as can be seen in Figure 10b. (The pigeon data on which Figure 10a was based did not span a wide enough range of wavelengths to reveal this circularity.) Third, as I have already noted, the circularity of perceived hues is consistent with the opponent-processes representation of colors (Fig. 7b; see Shepard 1993; Shepard & Cooper 1992), which I conjecture to have arisen as an adaptation to the three degrees of freedom of natural illumination. Fourth, circular components, though historically ignored by most psychophysicists, arise in many representational spaces, including, in addition to those for color and for position and motion (considered here, and in Shepard 1978b), the chroma circle and the circle of fifths for musical pitch (Balzano 1980; Krumhansl & Kessler 1982; Shepard 1964b; 1965; 1982a; 1983), and the circadian, circalunar, and circannual components of time (e.g., see Enright 1972; Winfree 1980).

3.1. Formal characterization of generalization based on possible kinds

I originally derived the proposed universal law of generalization for the simplest case of an individual who, in the absence of advance knowledge about particular objects, encounters one such object and discovers it to have an important consequence. From such a learning event, the individual can conclude that all objects of this kind are consequential and that they therefore fall in some region of representational space that overlaps the point corresponding to the object already found to be consequential. Apart from its overlap with this one point, however, this *consequential region* remains of unknown location, size, and shape in representational space.

Although it is not essential to the basic theory, in the interest of keeping the initial formulation as sharp as possible, I propose for the present to proceed on the working hypothesis that the region in representational space corresponding to a basic kind is a *connected region*. Between the




Figure 10. Generalization gradients for spectral hues obtained by applying multidimensional scaling to human and animal data: (a) based on the solution obtained by Shepard (1965) for the pigeon generalization data collected by Guttman and Kalish (1956), and (b) based on the solution obtained by Shepard (1962b) for the human similarity data collected by Ekman (1954). The distance, D , for each point is the Euclidean distance between the two colors in the multidimensional scaling solution based on generalization data, G ; and the smooth curve in each plot is a one-parameter exponential decay function fitted by Shepard (1987b). From "Toward a Universal Law of Generalization for Psychological Science," by R. N. Shepard, 1987, *Science* 237:1318. Copyright 1987 by the American Association for the Advancement of Science. Adapted by permission.




Figure 11. Multidimensional scaling configuration for Ekman's 14 spectral colors, obtained by Shepard (1962b) and corresponding to the plot shown in Figure 10b. The circle was subsequently drawn through the points to bring out the resemblance to Newton's color circle. The three-digit numbers indicate the wavelengths (in nanometers) of the corresponding stimuli. From "The Analysis of Proximities: Multidimensional Scaling With an Unknown Distance Function. II," by R. N. Shepard, 1962, *Psychometrika* 27:236. Copyright 1962 by the Psychometric Society. Adapted by permission.

points corresponding to any two objects of that kind, then, there is always a continuous path in the representational space that falls entirely within the consequential region for that kind. Thus, an apple could be continuously changed into any other apple, a chair could be continuously changed into any other chair, and a capital A could be continuously changed into any other capital A in such a way that at each step of the metamorphosis, the object would retain its recognizability as an apple, chair, or letter A, respectively.

The characterization of basic level categories in terms of a dichotomous criterion of connectedness, rather than in terms of some graded measure of correlation (of the sorts proposed by Rosch et al. 1976, and others) has two potential advantages: it can provide for the possibility of a sharp boundary between objects that though similar, belong to different natural kinds (only one of which, for example, manufactures a toxin); and it can provide for the possibility that objects of the same kind may nevertheless differ arbitrarily and widely in some of their features (an animal can vary enormously in size, shape, or coloration, for example, and still be a dog). Connectedness need not hold for non-basic (e.g., superordinate or ad hoc) categories. There may be no continuous series between two such pieces of furniture as a sofa and a floor lamp for which every object along the way is also recognizable as a piece of furniture; and there may be no continuous series between two letters of the alphabet such as "B" and "C" for which every intermediate shape is also recognizable as a letter of the alphabet. Even for what I am calling basic kinds, my current working assumption of connectedness is only provisional.

I begin by considering an individual who has just found a particular, newly encountered object to have a significant consequence. This individual can only estimate the likelihood that a second, subsequently encountered object also has that consequence as the conditional probability that a random and (provisionally) connected region that happens to overlap the point corresponding to the first object, also overlaps the point corresponding to the second. The gradient of generalization then arises because a second object that is closer to the first in the representational space is more likely to fall within such a random region that happens to overlap the first.

To obtain an unbiased quantitative estimate of the probability that the new stimulus is consequential, the individual must use Bayesian inference. In effect, such an individual integrates over all candidate regions in representational space – with whatever prior probabilities, $p(s)$, are associated in that individual with the different possible sizes, s , for such regions. (In the absence of advance information, these prior probabilities are naturally assumed to be independent of the locations of the corresponding regions in the representational space.) For a test stimulus corresponding to a position x in the representational space, the generalization $g(x)$ from a training stimulus (taken, without loss of generality, to be centered at the origin of an arbitrary coordinate system) to a new stimulus at location x is then given by

$$g(x) = \int_0^{\infty} p(s) \frac{m(s,x)}{m(s)} ds, \quad (5)$$

where $m(s)$ denotes a (volumetric) measure of the region of size indexed by s , and $m(s,x)$ denotes a corresponding measure of the overlap between two regions of that size, one centered at x and one centered at the training stimulus (i.e., at the origin).

The results of such integration turn out to depend remarkably little on the prior probabilities assigned (Shepard 1987b). For any choice of the probability density function $p(s)$ having finite expectation, integration yields a decreasing concave upward gradient of generalization. For any reasonable choice, integration yields, more specifically, an approximately exponential gradient. For the single most reasonable choice in the absence of any advance information about size – namely, the choice of the probability density function entailed by Bayesian inference from minimum knowledge or *maximum entropy* priors (see Jaynes 1978; Myung 1994) – integration yields exactly an exponential decay function (Shepard 1987b). Specifically, the maximum entropy assumption leads to a generalization function of the simple form

$$g(x) = \left(e^{-k|x|} \right) \quad (6)$$

where the single parameter k depends only on the expectation of $p(s)$.

Once again, invariance emerges only when formulated with respect to the appropriate, abstract representational space. To refer back to the domains of position, motion, and color, there is greater generalization between rectangles differing in orientation by 90° than between rectangles differing by somewhat less than 90° , and there is greater generalization between surfaces reflecting the shortest and longest visible wavelengths (violet and red) than between

either of these and a surface reflecting an intermediate wavelength (e.g., green). Clearly, generalization cannot be monotonic with distance in the usual physical space (of angle or wavelength). But generalization can become both invariant and monotonic with distance in the psychologically appropriate representational space, in which angles and wavelengths alike map into closed curves (Shepard 1962b; 1965; 1981a; Shepard & Farrell 1985).

Invariance in the law of generalization has thus been obtained by separating the *psychological* form of generalization in the appropriate psychological space from the *psychophysical* mapping from any specified physical parameter space to that psychological space. The psychophysical mapping, having been shaped by natural selection, would favor a mapping into a representational space in which regions that correspond to basic kinds, though differing widely in size and shape, have not, on average over evolutionary history, been systematically elongated or compressed in any particular directions or locations in the space. From what they learn about any newly encountered object, animals with a representational space for which biologically relevant kinds were consistently elongated or compressed in this way would tend to generalize too much or too little in certain directions of that space, relative to other species that had evolved an innate representational space that was appropriately regularized for the biologically relevant basic kinds in our world.

Ultimately, I expect the approach to generalization based on inference from maximum entropy priors, like the approaches I have already outlined for the representations of position, motion, shape, and color, to find a grounding in the theory of groups. This is because entropy, taken as a measure of the absence of knowledge (following Shannon 1948), can have a well-defined meaning only in relation to a space that (as I put it above) is properly “regularized” or that (in the words of Wiener 1948) has a “fundamental equipartition.” To take the simplest example, if we have no knowledge about the location of a point in a one-dimensional space, we can only suppose that every location on the line is equally probable. (This is the “principle of indifference” so successfully employed in physics by Maxwell and Boltzmann – see Jaynes 1978.) Accordingly, the distribution that maximizes entropy in this case is, in fact, the uniform distribution. But if we were to transform this space by a nonlinear transformation (such as $x^* = x^2$ or $x^* = \log x$), what had been a uniform and maximum entropy distribution in the original space would no longer be so in the transformed space, and vice versa. Without going further into this deep and subtle matter here, I simply note that, in the opinion of one of the leading proponents of the maximum entropy approach in physics, “This problem is not completely solved today, although I believe we have made a good start on it in the principle of transformation groups” (Jaynes 1978).

3.2. Extensions of the generalization theory

3.2.1. Determinants of the metric of representational space.

A distinction that has been found basic to the understanding of similarity assessments and to discrimination and classification performances is the now widely recognized distinction between psychologically integral and separable relations among stimulus dimensions (e.g., see Garner 1974; Lockhead 1966; Shepard 1964a; 1991). This distinction has also been found to have a natural basis in the idea of consequential regions (Shepard 1987b; 1991; Shepard & Ten-

enbaum 1991). To the extent that the extensions of such regions along two or more dimensions have been positively correlated over evolutionary history, the integration over all possible regions, with their associated maximum-entropy weights, yields surfaces of equal generalization that approximate ellipsoids, implying the L_2 norm and associated Euclidean metric for that multi-dimensional representational space. To the extent that the extensions of such regions along the different dimensions have been uncorrelated, the integration over possible regions yields surfaces of equal generalization that approximate cross polytopes (a diamond-shaped rhomb in two dimensions, a triangular-faced octahedron in three), implying the L_1 norm and associated “city-block” metric for that subspace. In both of these multi-dimensional cases, integration still yields the exponential type of decay of generalization with distance in representational space originally derived for the one-dimensional case (for which the Euclidean and city-block metrics are equivalent). (The most appropriate group-theoretic representation is expected to be different, however, for conjunctions of integral and of separable dimensions.)

3.2.2. Generalization over discrete features. Although the derivation of the exponential gradient of generalization has been outlined here for the case of a continuous representational space, the theory is not restricted to the continuous case. When the objects possess only discrete (or even binary-valued) features, the analogs of the consequential regions in the continuous case become consequential subsets, and the analog of the volumetric size, $m(s)$, of a region becomes the (finite) number of objects in such a subset. Nevertheless, summation (the discrete analog of the integration used in the continuous case) still yields an exponential type of fall-off of generalization with distance, where distance is now defined in terms of the sum of the weights of the features that differ between the two objects or, if the features are all equal in weight, simply in terms of the number of differing features (Russell 1988; see also Gluck 1991; Shepard 1989).

3.2.3. Classification learning. Over a sequence of learning trials in which different objects are found to have or not to have a particular consequence, Bayesian revision of the prior probabilities associated with the various candidate regions yields a convergence to the true consequential region (Shepard & Kannappan 1991; Shepard & Tenenbaum 1991). Moreover, it does so in a way that agrees with results for human categorization (e.g., Nosofsky 1987; 1992; Shepard & Chang 1963; Shepard et al. 1961): The learning proceeds more rapidly when the consequential set of objects forms a region in the representational space that is connected rather than disconnected (Shepard & Kannappan 1991). The learning also proceeds more rapidly when the consequential set is compact in terms of the Euclidean metric if the dimensions are integral, but more rapidly when the consequential set is based on shared features (or conjunctions of features) if the dimensions are separable (Shepard & Tenenbaum 1991). (For related simulations, see Nosofsky et al. 1992; 1994; and for a similar Bayesian approach in which, however, the underlying hypotheses are taken to be Gaussian distributions rather than the sharply bounded regions posited here, see Anderson 1991.)

3.2.4. A law of discriminative reaction time. As I noted in the discussion of critical times in imagined and apparent

motion, natural selection has favored the ability to make decisions not only accurately but swiftly. But, whereas the time required to determine that two things are identical despite their apparent difference linearly increases with their transformational separation in the space of possible positions (as in mental rotation), the time required to determine that two things are different despite their apparent similarity nonlinearly decreases with their separation in the space of possible objects.

Specifically, latency of a discriminative response, like probability of generalization, falls off according to a decreasing, concave-upward function of distance between stimuli in representational space. But, whereas generalization probability, which cannot exceed one, approximates an exponential decay function of distance, discrimination latency, which is unbounded, is expected (under idealized conditions) to grow without limit as the difference between the stimuli approaches zero. In practice, such a function cannot be precisely determined for very small differences; experimental subjects would eventually either simply make a random guess or leave the experiment to terminate a potentially interminable trial. Nevertheless, functions that have been obtained do often approximate a reciprocal or hyperbolic form (e.g., Curtis et al. 1973; Shepard 1981a; 1989; see also Shepard et al. 1975). Such a form can be theoretically derived within the framework of the generalization theory. Suppose, for example, (1) that the internal representations corresponding to candidate regions overlapping either stimulus become activated, each with probability per unit time proportional to that region's associated prior probability of being consequential, and (2) that the first such representation to be activated – which overlaps one but not the other of the two stimuli – precipitates the discriminative response. Integration over all possibilities then yields, for the expected latency of discrimination, a reciprocal type of dependence on distance in representational space (see Shepard 1987b).

3.2.5. The generality of generalization. Presumably, things having the potential for particular, associated consequences belong to distinct kinds (including physical elements, chemical compounds, and biological species) and do so not just in the human or even the terrestrial environment but throughout the universe. If so, the exponential law of generalization, the reciprocal law of discriminative reaction time, and the Euclidean and city-block metrics of representational space may have arisen not just for the humans or animals we have studied on earth. Such laws and such

metrics may have arisen wherever sufficiently advanced forms of life may have evolved. (This remains true even if biological species are themselves in part the product of mind – as suggested by the genetic algorithm simulations of Todd & Miller 1991.)

4. Conclusion

Perhaps psychological science need not limit itself to the description of empirical regularities observed in the behaviors of the particular, more or less accidental collection of humans or other animals currently accessible to our direct study. Possibly we can aspire to a science of mind that, by virtue of the evolutionary internalization of universal regularities in the world, partakes of some of the mathematical elegance and generality of theories of that world. The principles that have been most deeply internalized may reflect quite abstract features of the world, based as much (or possibly more) in geometry, probability, and group theory, as in specific, physical facts about concrete, material objects. By focusing on just three perceptual-cognitive examples – concerning the representation of the colors of objects, the kinds of objects, and the positions, motions, and shapes of objects – I have tried to indicate how psychological principles of invariant color, optimum generalization, and simplest motion may achieve universality, invariance, and mathematical elegance when formulated in terms of points, connected subsets of points, and geodesic paths in the appropriate abstract representational spaces.

ACKNOWLEDGMENTS

This article was originally published as the lead article in the first issue of the *Psychonomic Bulletin & Review* 1994, 1:2–28. The research and the drafting of this article were supported by the National Science Foundation (Grant No. DBS-9021648). Final revisions were supported, as well, by the Santa Fe Institute, and benefitted significantly from interdisciplinary discussions with many scientists there. Other colleagues who have especially contributed to or influenced the particular work reviewed here include Eloise Carlton, Lynn Cooper, Leda Cosmides, Joyce Farrell, Jennifer Freyd, Sherryl Judd, Laurence Maloney, Michael McBeath, Jacqueline Metzler, Geoffrey Miller, In Jae Myung, Robert Nosofsky, Margaret Shiffrar, Joshua Tenenbaum, Brian Wandell, and Susan Zare. Laurence Maloney and Vijoy Abraham greatly facilitated the conversion of the earlier version of the text and the figures to the present, edited electronic form, and Phineas de Thornley Head, of *Behavioral and Brain Sciences*, kindly entered a number of my final corrections to the electronic version.