

Orientation in Physical Reasoning: Determining the Edge That Would Be Formed by Two Surfaces

John R. Pani
Georgia Institute of Technology

Colin T. William
Emory University

Gordon T. Shippey
Georgia Institute of Technology

Physical reasoning is strongly influenced by various parameters of orientation. The authors report 3 experiments in which this phenomenon was explored for a particularly elementary transformation: the formation of a line from the intersection of 2 planes. Participants perceived pairs of planar surfaces (disks) in a variety of orientations in 3-D space and indicated the orientations of the edges that would result if the surfaces interpenetrated. The ranges of error and response time were large. Performance depended on whether the orientation of the edge that would be formed was the same as components of the orientations of the perceived surfaces, the degree to which the orientation of the edge would be canonical in the environment, and whether the angle between the surfaces would be perpendicular. The results are discussed in the context of a general approach to orientation in perception and physical reasoning.

When the relations among things in the world might be seen in one way but instead are seen in another, visual perception may be said to be a type of description (Pani, in press). Such situations occur consistently in perceptions that involve orientation. As a simple example, if the shape in Figure 1 is oriented as shown at the left, it is seen to have a pointed top and bottom and uniform orientations of edges and surfaces about one of its object axes. When the same shape is oriented as shown at the right, it appears to have a flat top and bottom and nonuniform orientations of edges and surfaces. Thus, a single object is seen to have different qualitative properties, those of a regular “dipyramid” or those of an “antiprism,” depending on the object axis that is used to determine the orientations of edges and surfaces (Pani, in press-a; Pani, Zhou, & Friend, 1997; see also Hinton, 1979). In this instance, the selection of an object axis depends on the orientation of the object to the vertical (Mach, 1906/1959; Rock, 1983).

Recently, it has become clear that spatial organization in terms of the orientations of things is critical to high-level perception and physical reasoning. Variation in orientation has a profound effect on the ability to perceive or imagine simple rotations (Pani, 1993; Pani & Dupree, 1994; Pani, William, & Shippey, 1995; Parsons, 1995; Shiffrar & Shepard, 1991), projective transformations (e.g., the casting

of a shadow; Pani, Jeffres, Shippey, & Schwartz, 1996), and elementary 3-D shapes (as in Figure 1; Hinton, 1979; Pani et al., 1997), as well as to reason about, or remember, common environmental occurrences (e.g., Chase, 1986; Hecht & Proffitt, 1995; Tversky, 1981).

It is well established that certain orientations are cognitively simpler than others. Vertical and horizontal orientations are perceived, remembered, and produced more accurately and efficiently than oblique orientations (e.g., Lehtinen-Railo & Jurmaa, 1994; Palmer & Hemenway, 1978; Pani & Dupree, 1994; Tversky, 1981). Parallel and perpendicular (“normal”) orientations are similarly canonical for human perceivers (Pani, 1993; Pani et al., 1996; Rock, 1983; Wertheimer, 1950). But investigations of canonical forms of orientation in the area of physical reasoning continue to provide new information about the role of orientation in spatial cognition and present a number of challenges for contemporary theories of perception, spatial organization, mental imagery, and reasoning.

In the present article, we report three experiments that were designed to explore the effects of orientation on the imagination of a particularly elementary spatial transformation. Participants were shown pairs of surfaces, as illustrated in Figure 2, and were asked to demonstrate the orientation of the line of intersection if the surfaces were to expand or move together so that they interpenetrated. In other words, participants in the experiments indicated the orientation of the line of intersection of two planes.

This task has a number of important properties. First, it is an instance of physical reasoning that requires perception (of the surfaces) and imagination (of the line of intersection). Within this domain, it significantly extends the set of tasks for which variation in orientation has been examined. The transformation involved does not require a change of object orientation, as rotation does, or of shape, as projection does.

John R. Pani and Gordon T. Shippey, College of Computing, Georgia Institute of Technology; Colin T. William, Department of Psychology, Emory University.

We thank Ulric Neisser for his help with this project and Carolyn Mervis for her comments on earlier versions of this article.

Correspondence concerning this article should be addressed to John R. Pani, who is now at the Department of Psychology, University of Louisville, Louisville, Kentucky 40292. Electronic mail may be sent via Internet to jrpani01@homer.louisville.edu.

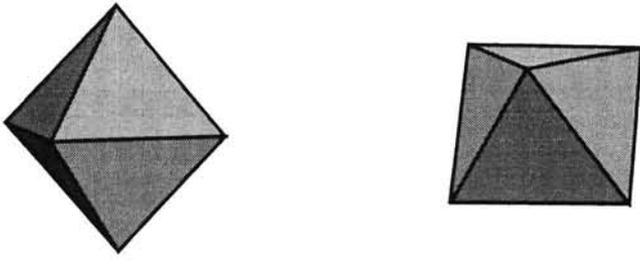


Figure 1. The regular octahedron tends to be seen either as a dipyrmaid (at the left) or as an antiprism (at the right) depending on which object axis is made salient.

Nonetheless, the intersection of planes shares with rotation and projection the property of being fundamental to the geometry of the everyday world (e.g., Gasson, 1983). For example, it is a constraint on nature that is observed in the formation of edges from surfaces. Intuitions that reveal a mind's eye that respects fundamental laws of geometric structure suggest that a person should perform the task simply by imagining one surface to "move over and pass through the other" (i.e., a pure translation). When the imagined surfaces intersect, the individual may "look" at the edge that was formed and note its orientation explicitly. (For discussion of intuitions about imagery, see Finke, 1989; Kosslyn, 1980; Pani, 1996; Pylyshyn, 1981; Shepard, 1978, 1984.)

We report, below, that people cannot always tell what edge would result from the joining of a pair of surfaces. For some pairs, adult participants typically can indicate accurately and efficiently the orientation of the edge that would be formed. For other pairs, the same individuals may take a very long time to decide and then give an answer that is incorrect by 40° or more. The pattern of performance tells much about the spatial organization of edges and surfaces and its relations to physical reasoning.

Before presenting the experiments, we discuss concepts important for understanding the perception of orientation (see Pani, in press-a). We then propose a specific model for determining the line of intersection of two planes. In the General Discussion, we present a theoretical perspective that encompasses this and other instances of physical reasoning.

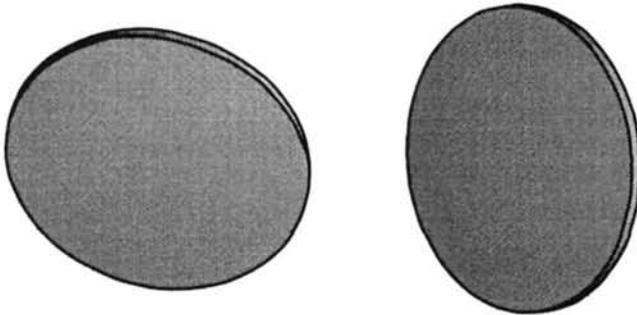


Figure 2. An example of a pair of disks for which an individual might determine an implied line of intersection.

In this view, successful physical reasoning typically depends on finding a fit between the description of a situation given in perception and the description needed for reasoning. A look at successful descriptions, and the fit between them, characteristically reveals three nested sets of properties. At the most specific level are the aligned orientations (parallel, perpendicular, vertical, and horizontal). A more general set of properties includes the symmetries (i.e., invariance across transformation). And the most general set of properties includes the singularities (maximum, minimum, same, and orthogonal). A complete description of a physical problem that uses these properties will be relatively efficient, distinctive, nonarbitrary, and pragmatically useful.

Orientations of Edges and Surfaces

Slant and Direction of Slant

Description of the orientation of an object must include specification of properties of the object, such as a major axis, for which orientation will be determined. One way to begin describing the orientation of a flat surface is to identify a direction in which it faces. A *surface normal* is generally used for this, and the description of orientation becomes a matter of determining the orientation of the normal. A common descriptive system used for this purpose refers to two orthogonal angles in *spherical coordinates*. Thus, the solar panel in Figure 3 is considered to face roughly upward and south.

Spontaneous perception of surface orientation appears to be coded in terms of such orthogonal angles, sometimes called the *slant of a surface* and the *direction of the slant* (e.g., Gibson, 1950; Sedgwick, 1986; Stevens, 1983a). This perception, as with any use of spherical coordinates, requires a reference system for determining the two angles. There

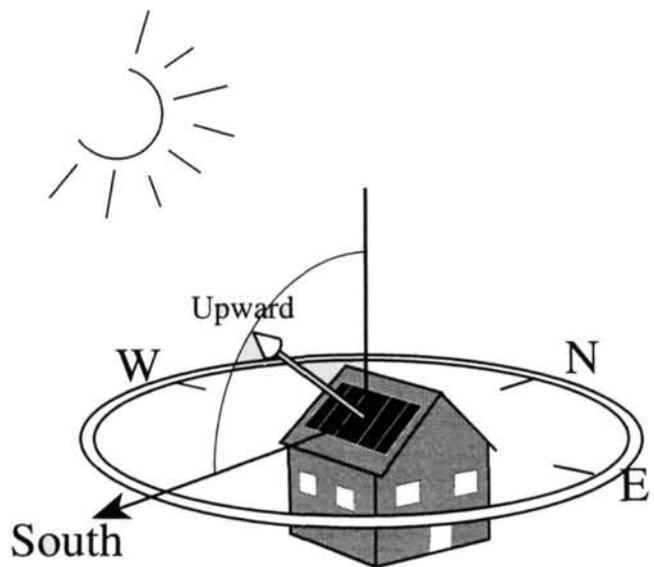


Figure 3. The direction that a surfaces faces is perceived in terms of its slant (upward in this case) and its direction of slant (south).

must be a polar axis and an equatorial plane that (typically) is orthogonal to the axis. If one is concerned with the direction that a surface faces, then the slant of the surface is the angle of a surface normal between the polar axis and the equatorial plane. Direction of slant is the direction within the equatorial plane that the surface normal points; it can be measured by projecting the normal to the equatorial plane and taking the angle of the projection relative to a reference point in the plane. When the polar axis is the line of sight, the equatorial plane is the frontal plane of the viewer, and the direction of slant is called "tilt" (Koenderink, van Doorn, & Kappers, 1992; Stevens, 1983a).

In many instances, the polar axis is not a single line but rather a single direction. In the local terrestrial environment, for example, there is an infinite number of verticals, but they all have the same direction. Similarly, the equatorial plane is not a single plane but rather a single planar orientation. In the local terrestrial environment, this orientation is the *horizontal*. The reference point for direction of slant varies widely. In large scale geography, it is often north, but on a baseball field it might be the centerline of the field.

A surface normal is an intrinsic property of a surface: It can be determined in the absence of context (for discussion of intrinsic geometric properties, see Osserman, 1995). In many instances, the context of an object, including its physical interactions, produces structure in the object that permits alternative descriptions of the object (Pani, in press-a). Of particular interest here, when a surface occurs in the context of a polar reference system, such as the gravitational vertical, it is not necessary to describe the orientation of the surface using the direction in which it faces (i.e., a surface normal, as in Figure 3). Orientation can be measured for lines on the surface itself, as illustrated in Figure 4. Consider that if a ball is rolled down a flat hill, it follows a particular line. This is the line of least action, and it

minimizes the angle of the surface relative to the polar axis while maximizing the angle relative to the equatorial plane (see Figure 4). If balls are started at different positions along the top of the hill, a family of parallel lines is produced, and any one of these *slant lines* may be used to measure the slant of the hill.

If steps are carved in the hill, they intersect the horizontal plane and face in the direction that the hill faces (i.e., they are orthogonal to it; see Figure 4). This family of parallel *direction lines* can give the direction of slant of the hill. Note that slant lines and direction lines will always be orthogonal to each other. Overall, slant and direction lines are *orientation lines*: lines on a surface that exemplify, and can be used to measure, its slant and direction of slant (Pani, in press-a).

Singular Orientations

There are sets of orientations of lines and planes, and, thus, of edges and surfaces, that are geometrically singular with respect to a polar reference system, as illustrated in Figure 5. These are orientations that share components with the reference system, and one consequence of this singularity is that these orientations are canonical for human perceivers. There is one orientation in which surfaces are parallel to the equatorial plane and perpendicular to the polar axis. For this one orientation, slant is zero and direction of slant does not exist. In the local terrestrial environment, these are the horizontal surfaces. There is a set of orientations in which surfaces coincide with the polar axis and are perpendicular to the equatorial plane. In the local terrestrial environment, these are the vertical surfaces (e.g., walls and doors).

The three Cartesian planes are singular surface orientations within a polar reference system (see Figure 5). One of the Cartesian planes is the horizontal. The other two are contained within the set of vertical surfaces. The frontal plane has a direction of slant of zero. The sagittal plane, being orthogonal to the frontal, also has a singular orientation (sideways).

There is one orientation in which lines are parallel to the polar axis and perpendicular to the equatorial plane. For this one orientation, slant is maximized (or minimized; see Figures 3 and 4), and direction of slant does not exist. In the local terrestrial environment, these are the vertical lines. There is a set of orientations in which lines coincide with the equatorial plane and are perpendicular to the polar axis. In the local terrestrial environment, these are the horizontal lines (e.g., the edges of tables and window sills).

The three Cartesian axes are singular orientations of lines within a polar reference system (see Figure 5). One of the Cartesian axes is the vertical. The other two are contained within the set of horizontal lines. One of them points in the direction of slant of zero. The other is orthogonal to this and points sideways. The Cartesian lines and planes, overall, are embedded within the polar reference system as a set of canonical orientations (once a value for zero direction of slant is selected).

Surfaces with Cartesian orientations can be considered aligned with a polar reference system. A *partially oblique*

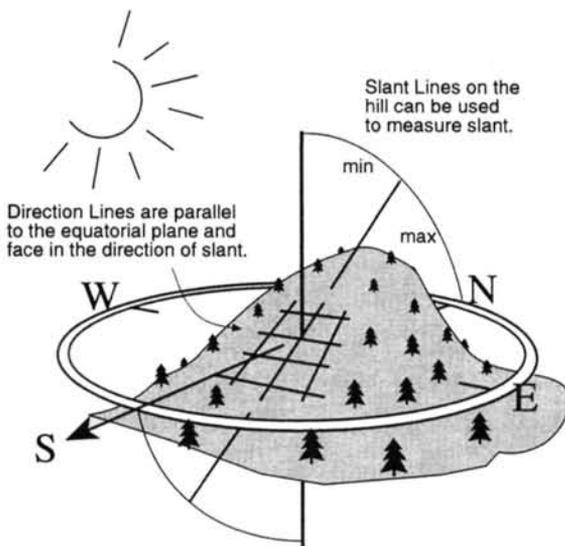


Figure 4. Orientation lines can be used to determine the slant and direction of slant of a surface.

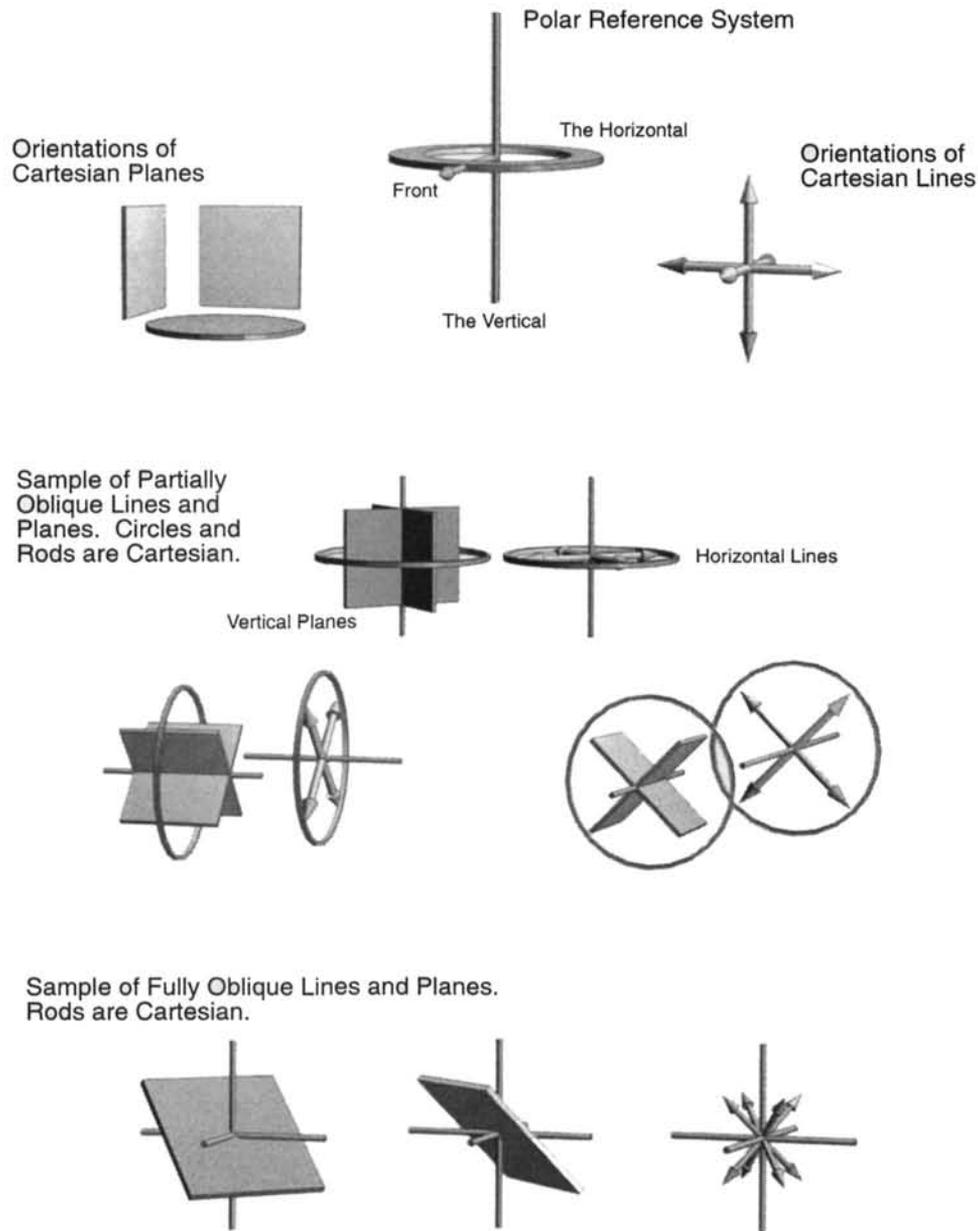


Figure 5. Orientations of lines and planes in a polar reference system. The edges of the rectangular surfaces are instances of orientation lines.

surface has singular slant or direction of slant, but not both (see Figure 5). One set of orientation lines on a partially oblique surface will be parallel to a Cartesian axis. The other set of orientation lines will intersect the orthogonal Cartesian plane but will be oblique to the canonical axes in that plane. A *fully oblique surface* has nonsingular slant and direction of slant. Neither its slant lines nor its direction lines are parallel to a Cartesian axis.

Similarly, lines with Cartesian orientations can be considered aligned with a polar reference system (see Figure 5). A *partially oblique line* has singular slant or direction of slant, but not both. The line will be contained within a Cartesian

plane but oblique to the axes in that plane. A *fully oblique line* has nonsingular slant and direction of slant. It is aligned with no Cartesian reference element.

Choice of the Polar Axis

The existence of singular orientations, and their possible relevance to human perceivers, makes especially important the question of the direction of the polar axis for perceivers. It is common in discussions of vision to consider the polar axis for surface orientation to be the line of sight (e.g., Gibson, 1950; Koenderink et al., 1992; Marr, 1982; Sedg-

wick, 1986; Stevens, 1981, 1983a, 1983b). One good reason for this is that 3-D orientation relative to the line of sight theoretically can be recovered from gradients of texture and linear perspective on surfaces (see Braunstein, 1976; Cutting & Millard, 1984; Gibson, 1950; Stevens, 1981). It is possible that the polar axis is the line of sight in early vision but not after perception of 3-D layout has been completed (see Stevens, 1983b; see also Epstein, Babler, & Bownds, 1992; an alternative hypothesis is offered by Sedgwick, 1983).

The present data suggest that the polar axis for normal perception and imagination of surface orientation is the environmental vertical. Slant, in this case, is what Gibson (1950) referred to as gravitational slant (also geographical slant; Gibson & Cornsweet, 1952; Kinsella-Shaw, Shaw, & Turvey, 1992; Proffitt, Bhalla, Gossweiler, & Midgett, 1995). A vertical orientation of the polar axis is consistent with numerous studies that have found a primacy of the vertical as a reference direction. A study by Pani and Dupree (1994) is particularly relevant to consideration of the present research. It was known that the perceived orientation of an axis of rotation had a large effect on the ability to determine the outcome of certain simple rotational motions. These authors found that it was the orientation of the axis to the environment, and especially whether the rod was aligned with the vertical, that was important in this task (see also Rock, 1973, 1983; Sedgwick & Levy, 1985; see also Corballis, Nagourney, Shetzer, & Stefanatos, 1978; Hinton & Parsons, 1988). Indeed, the vertical remained a strong reference axis even when performance would have improved considerably if an alternative reference axis had been adopted (e.g., an egocentric one).

The Edge That Joins Two Surfaces

If two surfaces intersect, each surface has an orientation with respect to the other. The simplest expression of this object-relative orientation is just the angle between the surfaces, called the *dihedral angle*. The apex of this angle is the line of intersection, as shown in Figure 6. The magnitude of the angle is measured in a plane that is perpendicular to the line of intersection. Thus, the line of intersection of the two surfaces, and the orthogonal plane for measuring the angle between the surfaces, is a special case of using a polar reference system to determine orientation. That is, the angle

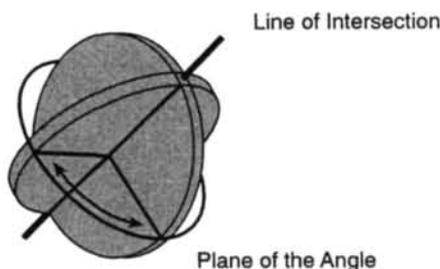


Figure 6. The dihedral angle.

of one surface to another is given by considering the surfaces to be vertical within a local reference system. The magnitude of the angle is the direction of slant of one surface relative to the other.

In engineering, a determination of the intersection of two planes is based on the geometry of the dihedral angle: Determine two surface normals, between which the dihedral angle can be found. The line of intersection is given as the perpendicular to the two normals (and to the dihedral angle; see Figure 6). Because an equivalent process might be carried out intuitively, the dihedral angle is a type of *base organization* for determining the intersection of two planes: It is a widely applicable organization, or description, of a physical phenomenon that readily serves as a basis for reasoning about it (Pani, in press-a, in press-b). In other words, the dihedral angle can be used as a mental model that underlies attempts to determine lines of intersection (for discussion of mental models, see Gentner & Stevens, 1983; Johnson-Laird, 1983).

Imagination of the Intersection of Planes

Degree of Fit

If two separated surfaces are perceived, and a task is set to determine their implied line of intersection, a goal-directed process is set in motion. The dihedral angle is a base organization of planar intersection that individuals could attempt to apply in this task. We suggest, however, that the dihedral angle often is not recognized immediately. This is because determination of a line of intersection begins with perception of the orientations of the surfaces relative to the environment (at least in the tasks we examined). The dihedral angle, in contrast, depends on object-relative description of orientation.

We propose that the need to begin reasoning in this task with perception of the surfaces relative to the environment encourages a heuristic procedure. Most generally, it is a heuristic search for the single line that is common to both surfaces. In particular, the individual attends to the orientations of the two surfaces, ultimately focusing on one of them (perhaps the one with the least canonical orientation). The slant and direction of slant of this surface are encoded in terms of the two families of orientation lines (see Figure 4). The individual then shifts attention to the second surface and determines whether any of the encoded orientation lines coincide with that surface. There will be such a coincidence just when an orientation line from the first surface is parallel to the line of intersection. This is a *fit between descriptions* because it comes from an attempt to work transitively from lines that are available in perception to lines that are needed in reasoning. In many cases there is not such a fit. We call the binary variable that codes a coincidence or noncoincidence between the orientation lines of one surface and the face of the other surface the *degree of fit* between the surfaces. Examples of pairs of surfaces that do or do not have this fit with the line of intersection are illustrated in Figure 7.

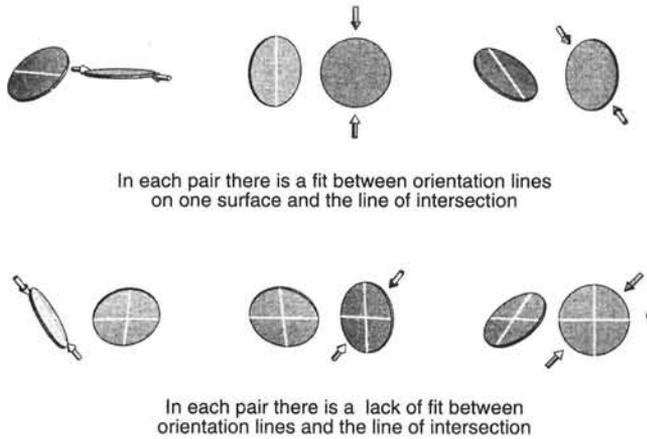


Figure 7. Degree of fit between a pair of surfaces and the implied line of intersection.

Orientation of the Line of Intersection

For any polar reference system, the singular lines and planes are ready candidates for the orientations of lines and planes that are needed in physical reasoning (e.g., Pani et al., 1996; Pani et al., 1995). All else equal, determining the line of intersection should be easiest when it is a Cartesian axis of the environment; determining the line should be more difficult when it is partially oblique, and even more so when it is fully oblique.

Dihedral Angle

A third variable concerns the orientations of the surfaces to each other. If people are sensitive to the organization of the dihedral angle, determining a line of intersection should be easier when that angle is perpendicular. In other words, it should be easier to obtain a purely object-relative organization of the surfaces and to find the line of intersection as the apex of the dihedral angle when the two surfaces are perpendicular to each other (e.g., Pani et al., 1996). Note that the proposed effects of these three variables are consistent with each other. The three variables concern separate geometric relations whose canonical values can be used to guide the search for lines of intersection.

Joint Surface Orientation

A fourth variable concerns the degree to which the orientation of each of the surfaces is canonical in the environment. Participants might identify an implied line of intersection well if both surfaces were aligned to the environment (i.e., if both were Cartesian). Performance might deteriorate as partially and then fully oblique surfaces were incorporated into the display pairs. The worst performance would occur when both surfaces were fully oblique to the environment. We call this gradation in the orientation of the surfaces *joint surface orientation*, and a scaling of it is summarized in Figure 8. Surface orientation is not as directly related to lines of intersection as are the three

variables discussed above. Hence, our view of the possible importance of canonical surface orientation in this task is that the cognitive processes suggested by the other three variables would be easier as the surfaces involved had more canonical orientations. For example, the process of ascertaining whether orientation lines from one surface coincided with the other surface (the degree of fit discussed above) might be more efficient as the second surface had a more canonical orientation.

Experiment 1

The first experiment was designed to examine people's ability to determine the line of intersection of two planes represented by physically present surfaces. Participants viewed pairs of wooden disks for as long as they wished. In each instance, participants indicated the orientation of the implied line of intersection in terms of one of the 13 canonical orientations in 35° to 45° intervals. Thus, there was no time pressure, and extreme accuracy was not required.

It is important to note that the four physical variables—degree of fit, orientation of the line of intersection, angle between the surfaces, and joint surface orientation—do not permit a factorial experimental design. For instance, it is impossible to select two surfaces that are both aligned to the environment but that intersect in a line that is oblique to the environment. Our response was to sample systematically from the universe of pairs of surface orientations and then to base the primary analyses of the experimental data on multiple regression.

Method

Participants. Twenty Emory University undergraduate students (11 women, 9 men) participated for course credit in an introductory course in psychology.

Materials. Participants viewed pairs of white wooden disks displayed side by side in front of a large black cardboard surface. Each disk was 7.5 cm in diameter and 0.375 cm thick. The disks were suspended from behind with wooden rods, out of view. The centers of the disks were 13 cm apart, 10 cm in front of the background, and 38 cm above a standard laboratory table. Participants viewed the disks and the background through a circular

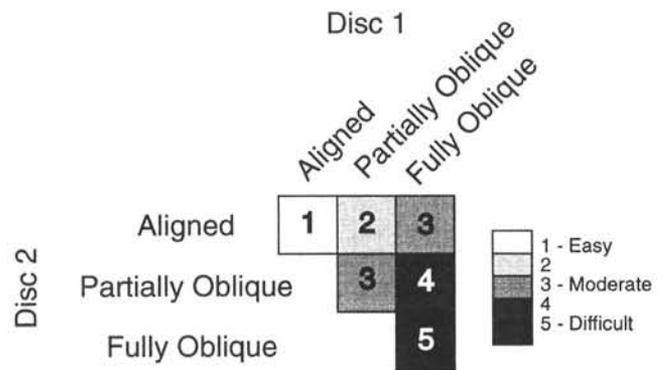


Figure 8. Scaling of joint surface orientation.

porthole, 30 cm in diameter, in a large gray screen. The screen stood 55 cm in front of the background. The porthole could be covered from behind to block the participant's view. The experiment was conducted in a brightly lit laboratory room. The apparatus was positioned in such a way that the disks were well illuminated from above (and slightly in front) by a bank of florescent lights. In addition, the porthole allowed much light to enter from the front.

The display pairs were sampled from the infinitely large universe of intersecting planes according to a three-step process, as indicated in Figure 9. First, the orientations of the lines of intersection were selected. In particular, we selected the 13 canonical directions that include the three fully aligned directions (i.e., the Cartesian axes), six partially oblique directions (i.e., two directions within each of the Cartesian planes), and the four fully oblique directions. (In particular, we chose the 13 axes of cubic rotational symmetry.) With the lines of intersection established, it was necessary to select planes that intersected at those lines. The simplest way to think of the set of planes that intersect at a line is to imagine the pages of an open book splayed out in all directions. The pages of the book all have different orientations, but they intersect in one line, the spine of the book. We selected four representative surfaces for each of the 13 orientations, equally spaced in their orientations about the line. Thus, each surface was at a 45° angle (i.e., oblique) to its two immediate neighbors and 90° from the remaining surface. For each of these sets of four surfaces, two of the surfaces were aligned to the environment as much as possible; the remaining two surfaces were maximally oblique to the environment. (The 52 surfaces generated in this way are described in the Appendix.) In the third step of generating the displays, all six possible pairs from each set of surfaces at each of the 13 axes of intersection were selected. This process produced 78 displays.

We decided that 78 displays were too many for individual participants. If one considers just the unique combinations of values of the four independent variables represented in this experiment, there were 14 unique types of display among the 78 displays. For each participant, many of the displays that were redundant with respect to the 14 unique types of display were eliminated. A constraint on this reduction, however, was that each participant received all available display types for all three lines of intersection aligned with the environment, all three principal planes that could contain a partially oblique line of intersection, and three out of four fully oblique lines of intersection. In the end, the number of trials for each participant was reduced to a more manageable 33. The full set of 78 displays was used by counterbalancing the sampling of displays across participants. A separate random order of presentation of the displays was used for each participant.

A small multicolored wooden house was used to familiarize participants with concepts important to the experiment. The house was T-shaped, with symmetric sloped roofs above each leg of the T. Each surface on the house was painted a different color for easy

reference. The house was designed to include edges that represented various aligned, partially oblique, and fully oblique lines of intersection.

Participants demonstrated the orientations of the lines of intersection in any direction in 3-D space, in intervals of 35° to 45° . The response apparatus consisted of a white polyhedron, 4.75 cm in diameter, with faces in all 13 canonical directions (26 faces in all). Each face of the polyhedron was drilled with a hole perpendicular to the face so that participants could insert a thin dowel, 25 cm in length, through the center of the block in exactly the 13 orientations of the lines of intersection. (Tape wrapped around the dowel kept it from sliding all the way through the polyhedron.) The polyhedron was held at the height of the disks by a stand made of gray pipe, 1.25 cm in diameter. A horizontal piece of pipe held the polyhedron at the rear (and, thus, was out of view). The back end of that pipe was connected to a vertical piece that secured the whole assembly to a wooden stand.

Procedure. Participants were tested individually in single sessions. Each session began with the participant reading a set of written instructions. The instructions used numerous illustrations and references to the multicolored model house to demonstrate the intersection of planes in 3-D space. The instructions went on to explain the participant's task: to determine the line of intersection between the two planes represented in the display and to match the orientation of that line to one of the possible orientations of the stick in the polyhedron. The instructions mentioned two possible methods for thinking about the line of intersection of two planes represented by surfaces. One could imagine the surfaces to move until they interpenetrated, or one could imagine the surfaces to expand until they intersected. We assumed, however, that participants would use whatever cognitive process they found most effective, and there were no restrictions placed on the participants. For example, the participants could use their hands to represent the surfaces if they wished. The participants were told to take as much time as they needed in each trial. The important consideration was that they be as accurate as they could be.

After a participant read the instructions, the experimenter reviewed the information in them and demonstrated the use of the response apparatus. The experimenter then asked the participant to use the response apparatus to indicate the orientation of the experimenter's pencil held at a variety of orientations (with equal representation given to aligned, partially oblique, and fully oblique orientations). If the participant responded incorrectly, the experimenter corrected the participant. After six consecutive correct reproductions of the orientations of the pencil, the experimenter asked the participant to look at a pair of disks in the display apparatus, find the line of intersection, and indicate it with the response apparatus. Again, if the participant answered incorrectly, the experimenter corrected the participant. There were two of these practice trials, with displays chosen from the set of displays that would not be shown to the participant later. With the practice trials completed, the experimenter informed the participant that the experimental trials would be timed but that it was most important to be accurate.

Each trial proceeded as follows. First, the experimenter covered the porthole and arranged a pair of disks in the display apparatus. When the participant was ready to proceed, the experimenter revealed the display and began timing the participant by pressing a computer key. The experimenter sat silently behind the participant and waited until the participant inserted the stick completely into the polyhedron. The experimenter then pressed the computer key, and the duration of the response was recorded. The experimenter recorded the participant's response and then covered the porthole

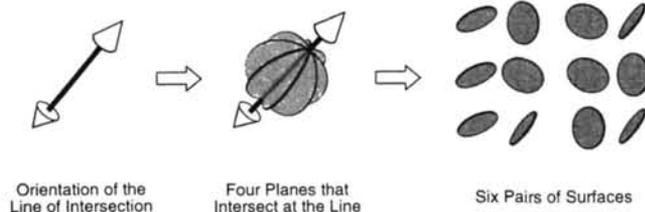


Figure 9. Logical steps in the generation of experimental displays.

and prepared for the next trial. The participant was given a 5-min break halfway through the set of trials.

Results

It was not unusual for participants to take 40 s or more to respond to particular types of display. Similarly, a majority of responses for certain types of display were errors (i.e., incorrect by at least 35°). For a task such as this, response time is best considered a measure of the overall level of difficulty experienced by the participants (e.g., Pani et al., 1995). For this reason, all response times were included in relevant statistical analyses.

Mean error, mean response time, and a combination of these, called *overall performance*, were dependent variables in analyses of multiple regression. Error was computed by considering the correct answer and the participant's answer to be axes running through the center of a sphere. If an answer was incorrect only by an angle in a vertical plane (i.e., slant) or in a horizontal plane (i.e., direction of slant), then the error was simply the angle of separation in that plane. If an answer was incorrect in both a vertical and a horizontal plane, these error values were considered legs of a spherical triangle, and the error was calculated as the hypotenuse of the triangle. (Although we think this the best way to measure angular error, we did verify that other measuring schemes do not significantly alter the nature of the results.)

Both error and response time (RT) were averaged across participants for each of the 33 display types received by every participant. The measure of overall performance was computed because of apparent trade-offs between speed and accuracy for individual types of display. The measure of overall performance was produced by converting mean error and mean RT to *z* scores for each of the 33 display types and then taking the means of the *z* scores.

The 33 error values ranged between 9.00° and 54.75°. Mean error across the entire set of displays was 29.03°, with a standard deviation of 13.00°. The 33 RT values ranged between 9.90 s and 41.80 s. Mean RT across the entire set of displays was 25.88 s, with a standard deviation of 8.76 s.

The multiple regression analysis for each dependent variable included four independent variables. These were degree of fit (with a fit assigned 1 and a nonfit assigned 2), orientation of the line of intersection (with aligned given 1, partially oblique given 2.5, and fully oblique given 3), joint orientation of the surfaces in a pair (scaled from 1 to 5; see Figure 8), and the angle of intersection (with 90° assigned 1 and 45° assigned 2). Stepwise multiple regression was used for statistical analysis; the criterion for entry at each step was that a variable be statistically significant ($p < .05$) at that step and that it have the smallest value of p among the set of unentered variables.

The nonlinear scaling for orientation of the line of intersection was based on the results of previous research (e.g., Pani et al., 1995). In general, however, we realized that there are many alternatives for the categorical breakdown and for the scaling of most of the independent variables.

Thus, we explored a variety of alternatives in order to verify that the obtained results were not due to arbitrary decisions. The results of the alternative analyses were not essentially different from those reported here.

Correlation matrices and values of R for all three dependent variables are presented in Table 1. For error, degree of fit entered into the regression first, $r = .715$, and accounted for 51% of the variance among the means. The angle between the surfaces remained statistically reliable after the first step and entered next into the regression equation. The R for this analysis was .771, accounting for 60% of the variance among the means, $F(2, 30) = 22.0$, $p < .001$. For RT, joint surface orientation entered the regression first, $r = .792$. The angle between the surfaces entered next. The R for this analysis was .824, accounting for 68% of the variance among the means, $F(2, 30) = 31.7$, $p < .001$. In the analysis of overall performance, joint surface orientation entered the regression first, $r = .745$. Degree of fit entered next, $R = .826$. Angle between the surfaces then entered the regression. The R was then .878, accounting for 77% of the variance among the means, $F(3, 29) = 32.7$, $p < .001$.

Among the lines of intersection that were aligned with a Cartesian axis, performance with the vertical direction was generally superior. For error, there was an overall effect of the direction of the aligned axes, $F(2, 38) = 3.28$, $p < .05$. Closer examination indicated that the vertical led to better performance than the other two directions taken together, $F(1, 38) = 5.63$, $p < .05$. For RT, the overall effect of the direction of the aligned axes approached significance, $F(2, 38) = 2.63$, $p = .09$. Participants were faster with the vertical than with the other two directions taken together, $F(1, 38) = 4.25$, $p < .05$.

Discussion

There was a wide range of performance on this task. For some pairs of surfaces, it was easy for the participants to determine the orientation of the line of intersection, with mean RT of around 10 s and mean error under 10°. (Note that RT included the time to manipulate the response

Table 1
Correlation Matrices and Values of R From Regression of the Three Measures of Performance on the Four Orientation Parameters

Variable	1	2	3	4	R
1. DegreeFit	—	.447	.196	.461	
2. Intsect		—	.144	.745	
3. Angle			—	.102	
4. JointSurf				—	
Error	.715	.453	.423	.474	.771
RT	.407	.707	.305	.792	.824
Overall	.660	.683	.428	.745	.878

Note. All values are statistically reliable, $p < .05$. Boldface correlation coefficients contributed to the value of R . DegreeFit = degree of fit; Intsect = orientation of the line of intersection; Angle = angle between the surfaces; JointSurf = joint surface orientation; RT = reaction time.

apparatus.) For other pairs of surfaces, mean RT was over 30 s and mean error was greater than 40°. In a task with such a wide range of difficulty, it would seem that error (or accuracy) is the single most important measure to consider, and degree of fit accounted best for error. Nonetheless, joint surface orientation accounted best for the time taken to respond, and the two variables together accounted for the combined measure of overall performance. Angle between the surfaces was a statistically reliable predictor for all three measures of performance (error, RT, and overall performance).

Experiment 2

There were three primary goals of the second experiment. The first was to obtain a measure of the accuracy with which participants perceived the orientations of the individual surfaces in each trial. These data could be used to ascertain whether the pattern of results for the determination of surface intersection might be due to anisotropies in the accuracy of perceived surface orientation (over and above those anisotropies that correlate with the scale of joint surface orientation). The second goal was to display the surfaces so that they were situated in a more normal 3-D space (i.e., not an empty space behind a round aperture). For this purpose, the surfaces were displayed as stereographic computer simulations of shaded textured objects poised above a textured table. The third purpose was to use a display and response system that permitted the analog indication of the orientation of the line of intersection to be closer to the surfaces. This would greatly reduce the need for short-term memory of the surfaces and the supposed line of intersection because perception, imagination, and response could occupy the same local space. In Experiment 2, therefore, the orientation of the line of intersection was indicated with a stereographic shaded double-headed arrow that was manipulated, by computer keys, just over the top of the pair of disks.

Method

Participants. Twenty-four Emory University undergraduate students (14 women, 10 men) participated for course credit in an introductory course in psychology. None of the participants had been in the earlier experiment.

Materials. The pairs of surfaces and the analog response elements were displayed through computer images created with a state-of-the-art 3-D graphics package (Infini-D, ver. 2.6). The computer images were presented stereographically on a computer monitor 896 by 672 pixels at a resolution of 72 dots per inch. The stereographic aspect of the displays consisted of the presentation of two views of a simulated 3-D scene, one view for each eye. Each view was contained in a computer image that was 448 pixels square and centered within a lateral half of the monitor (with a distance of 15.8 cm between the centers of the two views). Participants viewed the displays through a set of four front-surface mirrors set up as two sideways periscopes. When one looked through the inner pair of mirrors, ocular convergence appropriate for viewing a point on the screen actually provided views of the centers of the two laterally separated images (for more detail, see Pani et al., 1995). The images were accurate in size, linear perspective, and stereo

disparity for sets of objects of given sizes located in a single 3-D space 25 cm behind the screen and viewed from a distance of 51 cm.

The mirrors were held securely by a stand at the height of the display images, 42.5 cm from the monitor. (The periscopic geometry and the distance between the eyes and the mirrors gave an effective viewing distance that averaged 51 cm.) There was a small screen between the viewer and the mirrors with a 2.9 cm square aperture in front of each inner mirror. Participants sat in a chair and looked into the inner mirrors as if through a pair of binoculars fixed in space. The only view was of the portion of the monitor that held the display images. The stand was constructed so that participants could easily reach around it with both arms in order to press keys on the computer keyboard. The room was darkened to eliminate reflection off of the monitor.

Pairs of surfaces were rendered as textured blue disks, illustrated in Figure 10. The surfaces appeared to be 9 cm in diameter and 0.2 cm thick. They were placed side by side slightly above a simulated white table with black random texture lines. The centers of the surfaces appeared to be 10.86 cm apart and 6.1 cm above the table.

The display pairs were primarily the same as those used in Experiment 1. One difference was that pairs with fully oblique lines of intersection intersected at a line that had a $\pm 45^\circ$ slant to the vertical (rather than the $\pm 55^\circ$ of the cubic axis). This change standardized all of the axes that were oblique to the vertical at a single value. An additional difference from the first experiment was that certain pairs were eliminated so that there would be fewer trials. Pairs with intersections at the line of sight were not used, and there were fewer replications of the pair types that had fully oblique lines of intersection. (There was a total of 8 such pairs, rather than 12.) This reduced the total set of surface pairs to 70. This still was too many displays for individual participants. Using the sampling criteria developed in Experiment 1, each participant received a set of 26 trials from the complete set. The full set of 70 displays was used by counterbalancing the sampling of displays across participants. A separate random order of presentation of the displays was used for each participant.

The initial display in each trial was the pair of disks and the table. At the press of a computer key, a small gray disk with an arrow pointing forward appeared centered above the disks. The participants used four keys on the number pad to adjust the orientation of this surface normal so that it matched the orientation of each disk, illustrated at the left in Figure 10. The 8 and 2 keys adjusted the slant of the indicator in 5° steps. The 4 and 6 keys adjusted the direction of slant of the indicator. As this method required a minimum of 2,666 images to be stored on disk (just for the surface

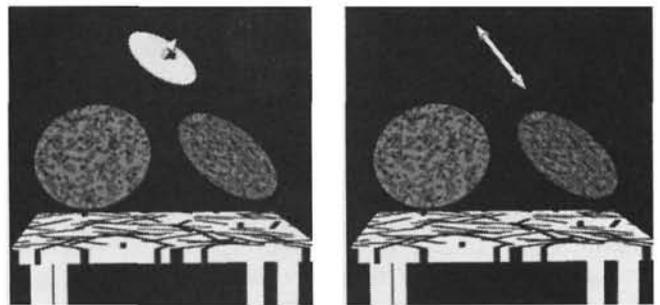


Figure 10. Examples of displays used in Experiment 2. Actual displays were stereograms. Actual disks were blue. The orientations of the surface normal and the double-headed arrow were manipulated by the participants.

normal), the surface normal was not programmed to turn around and face backward (which would have greatly increased the number of images required). Participants were instructed, and practiced, to always think of the surface as facing forward (or else 90° to the line of sight).

After indicating the orientation of each of the surfaces, participants indicated the orientation of the line of intersection. A key press displayed a double-headed arrow centered above the disks, illustrated at the right in Figure 10. The same four keys on the number pad changed the orientation of the arrow in 22.5° increments. So that operation of this indicator would be the same as operation of the surface normal, the arrow also could not be turned around backward. Beginning at the line of sight, the arrow could be moved only as far as the frontal plane. In Experiment 3, this constraint on the arrow was removed and was found to make no difference in the experimental results.

Procedure. The instructions were quite similar to those of Experiment 1. Demonstrations included use of the keyboard to administer the trials. There were six practice trials, during which the experimenter monitored and discussed the participant's use of the keys. With the practice trials completed, the experimenter informed the participant that the experimental trials would be timed but that it was most important to respond correctly.

Each trial proceeded as follows. First, the participant pushed a start key, and the pair of disks appeared on the monitor. When the key was pushed again, the surface normal appeared, facing the participant. The participant matched the orientation of the surface on the left by using the number pad keys to adjust the orientation of the normal, taking as much time as he or she needed. Pressing an end key recorded the final setting and reset the normal. The participant then adjusted the normal to match the orientation of the surface on the right. Pressing the end key recorded the final setting and removed the normal from the screen. A further press of the start key displayed the arrow. The participant used the number pad to adjust the arrow to an orientation that matched the line of intersection, as best the participant could determine. Again, there was no time limit. Pressing the end key blanked the screen. When the participant was ready, a further press of the start key began a new trial.

Results

Error in matching the surface normal to the orientations of the individual surfaces was scored with the system that was used for the line of intersection in Experiment 1. That is, error just in slant or in direction of slant, relative to the vertical, was simply recorded (i.e., as number of steps times the 5° step size). When there was error in both slant and direction of slant, the error was computed as the angle subtended by the hypotenuse of a right spherical triangle with legs corresponding to the angular error in slant and direction of slant. The overall mean of the unsigned error in matching the surface normal to the individual surfaces was 8.4°, with a standard deviation of 3.4°. Mean unsigned error in surface normal matching as a function of the individual surface orientations is presented in Figure 11. There were no statistically significant differences in error between surfaces that slanted forward and those that slanted backward.¹

The correlation between the a priori scale of joint surface orientation (as used in Experiment 1) and the mean error in matching the surface normal to the disks in a display was .890, $t(24) = 9.56$, $p < .001$. That is to say, the more that two surfaces in a pair were oblique, the larger the mean error

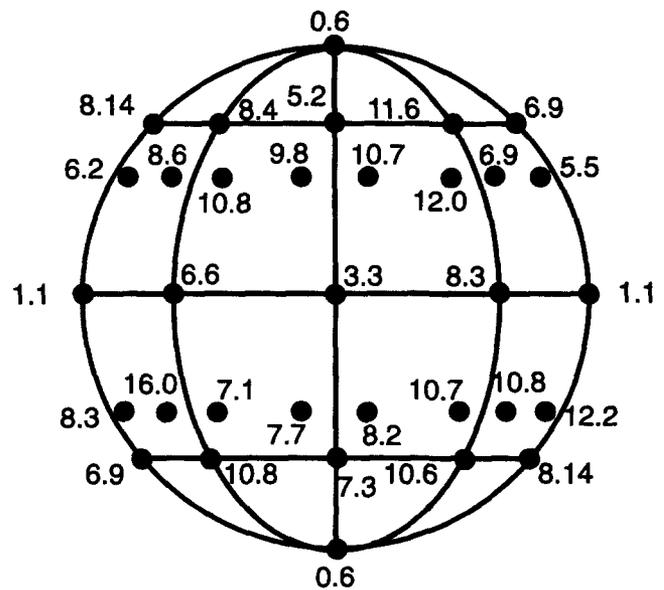


Figure 11. Mean error in degrees for matching the surface normal to the orientations of individual surfaces in Experiment 2.

in matching their orientations. For all three measures of performance in determining the line of intersection (i.e., error, RT, and overall performance), the a priori scale of joint surface orientation correlated with performance more highly than did the error in matching the individual surfaces in a pair. For RT and overall performance (but not for error), this difference between correlation coefficients was statistically significant, respectively, $t(23) = 2.44$, $p < .05$; $t(23) = 2.18$, $p < .05$; two-tailed tests. In the following summaries, therefore, only the a priori scale of joint surface orientation is discussed explicitly.

The range of performance again was quite large, with relatively fast and accurate response for some pairs of surfaces and slow and inaccurate response for other pairs. The 26 error values ranged between 2.81° and 59.38°. Mean error across the entire experiment was 25.67°, with a standard deviation of 13.66°. The 26 response time values (for determining the line of intersection) ranged between 8.97 s and 50.29 s. Mean response time across the entire set of displays was 30.85 s, with a standard deviation of 12.87 s.

Correlation matrices and values of R relating error, RT, and overall performance to the four independent variables are presented in Table 2. For error, orientation of the line of intersection entered the analysis first, $r = .773$. Angle between the surfaces entered next into the regression, $R =$

¹Inspection of Figure 11 does not suggest a replication of the result of Koenderink et al. (1992), in which error for matching surface orientation was greater for slant away from the frontal plane than for tilt around the line of sight. Their display, however, was a continuously curved surface, roughly spherical, that was viewed monocularly. This is quite different from flat surfaces, with clearly visible edges, viewed stereoscopically. In addition, the control of the surface normal in their experiment used the line of sight as the polar axis, whereas our experiment used the vertical.

Table 2
Correlation Matrices and Values of R From Regression
of the Three Measures of Performance on the
Four Orientation Parameters

Variable	1	2	3	4	R
1. DegreeFit	—	.587	.326	.761	
2. Intsect		—	.135	.697	
3. Angle			—	.088	
4. JointSurf				—	
Error	.591	.773	.428	.582	.874
RT	.564	.728	.498	.726	.888
Overall	.614	.798	.492	.695	.920

Note. All values are statistically significant, $p < .05$ (except the single value of $r < .1$). Boldface correlation coefficients contributed to the value of R . DegreeFit = degree of fit; Intsect = orientation of the line of intersection; Angle = angle between the surfaces; JointSurf = joint surface orientation; RT = reaction time.

.832. Degree of fit then entered, $R = .874$. These three variables together accounted for 76% of the variance among the error means, $F(3, 22) = 23.7, p < .001$. For RT, orientation of the line of intersection entered the regression first, $r = .728$. Angle between the surfaces then entered the regression, $R = .832$. Joint surface orientation then entered, $R = .888$. These three variables together accounted for 79% of the variance among the means of RT, $F(3, 22) = 27.3, p < .001$. In the analysis of overall performance, orientation of the line of intersection entered the analysis first, $r = .798$. Degree of fit entered next, $R = .856$. Angle between the surfaces then entered the analysis, $R = .920$. These three variables together accounted for 85% of the variance among the means on the measure of overall performance, $F(3, 22) = 40.6, p < .001$.

Discussion

This experiment used a display and response system in which the analog indication of the line of intersection occupied the same local space as the surfaces. This permitted immediate comparison of the surfaces, the imagined line of intersection, and the analog indicator, thus greatly reducing the need for spatial short-term memory. Nonetheless, there continued to be a wide range of performance in the determination of intersections of planes. For some pairs of surfaces, mean RT was less than 10 s and mean error was less than 10° . For other pairs of surfaces, mean RT was over 40 s and mean error was greater than 40° .

The major difference between these results and those of Experiment 1 was that orientation of the line of intersection joined with the angle between the surfaces in correlating significantly with all three measures of performance (i.e., error, RT, and overall performance). Degree of fit continued to be related to error and to overall performance. Joint surface orientation continued to be related to RT. In fact, joint surface orientation approached statistical significance for overall performance also, $p = .10$. With more statistical power, all four independent variables might have contributed to accounting for this overall measure. As it was, the R

for overall performance reached a high value (.920; $R^2 = .847$).

The accuracy with which participants perceived the orientations of the displayed surfaces correlated with the a priori scale of joint surface orientation. However, the a priori scale had slightly higher correlations with the ability to determine lines of intersection.

Experiment 3

An impressive result of the previous experiment is that three orientation parameters contributed independently to variations in performance. There was an opportunism such that participants benefited from canonical values on a variety of parameters of orientation. In Experiment 3, we tested whether this opportunism would extend to a new parameter. Consider that the effects of degree of fit and orientation of the line of intersection imply a tendency to perceive and imagine things relative to the reference system of the permanent environment. The influence of the angle between the surfaces suggests the importance of object-relative orientation. In many other studies, it has been observed that perception and reasoning can be influenced by salient *local reference systems* that arise either from object frames (e.g., an automobile or elevator) or from configurations of objects (Koffka, 1935; McAfee & Proffitt, 1992; Palmer, 1980; Pani & Dupree, 1994). Just lining up objects can create a local configuration that acts as a reference axis and influences spatial organization (Attneave, 1971; Kubovy, 1994; Mach, 1906/1959; Palmer, 1989). We thus became interested in the question of whether a local reference system would influence reasoning about the intersections of planes. In Experiment 3, the disks within display pairs were not always separated along a horizontal line, as they were in Experiments 1 and 2. Rather, the positions of the disks were considered to form an axis. As much as the geometry of 3-D space permitted, this position axis was varied independently in its orientation to the line of intersection and to the Cartesian axes of the environment. If the position axis could serve as a salient reference axis, it would be especially useful when it was aligned with the line of intersection. For illustration, a position axis aligned with the line of intersection is shown in Figure 12.

Method

Participants. Twenty-four Emory University undergraduate students (14 women, 10 men) participated for course credit in an introductory course in psychology. None of the participants had been in the earlier experiments.

Materials. The physical apparatus and the method for construction and presentation of displays were the same as those used in Experiment 2. Construction of the set of displays began with selection of seven types of display from the first two experiments. Two relatively easy types were chosen with lines of intersection aligned with the environment (i.e., vertical and frontal-horizonal). Three moderately difficult types were chosen with partially oblique lines of intersection (one each from the horizontal, frontal, and sagittal planes). One rather difficult type and one very difficult type were chosen from the fully oblique lines of intersection. For each of

these seven types of display, the disks were positioned in 3-D space so that the position axis was parallel, perpendicular, or oblique to the line of intersection. Independently, the position axis could be parallel, perpendicular, or oblique to the Cartesian planes of the environment. The disks were made slightly smaller than in Experiment 2 and were moved somewhat closer together (see Figure 12).

The geometry of 3-D space does not permit a complete factorial combination of the two position axis variables. Accordingly, multiple regression analysis again was used to determine their relative effects. Overall, there were 5 different position axis orientations for each of the display types with a line of intersection aligned with the environment, 7 different position axis orientations for each of the display types with a partially oblique line of intersection, and 6 different position axis orientations for each of the display types with a fully oblique line of intersection, making 42 categories of display in all. Each participant received displays from all 42 categories. For generality, two different versions of each of the partially oblique display types were used across participants (e.g., with a line of intersection slanting left or slanting right). Four different versions of the fully oblique pairs were used across participants (e.g., with a line of intersection slanting upper-left or lower-right). A separate random order of presentation of the displays was used for each participant.

Procedure. The procedure was identical to that in the previous experiment, with two notable exceptions. Participants did not judge the orientations of the individual surfaces (i.e., the surface normal did not appear on the screen). In addition, the double-headed arrow could be moved through a full 360° arc in either direction for slant or for direction of slant.

Results

The critical analyses concerned the orientation of the position axis: orientation of the position axis to the environment (with values of 1, 2, and 3 for aligned, partially oblique, and fully oblique, respectively), orientation of the

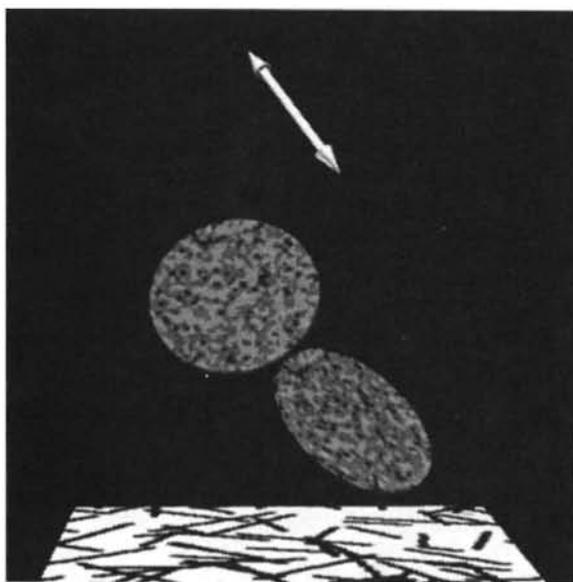


Figure 12. A display from Experiment 3 in which the position axis is aligned with the line of intersection.

Table 3
Correlation Matrices Relating the Three Measures of Performance to the Orientation Parameters for the Position Axis

Variable	1	2	3
1. PosIntsect	—	.048	.708
2. PosEnviro		—	.739
3. PosJoint			—
Error	.032	.160	.135
RT	.188	.057	.167
Overall	.119	.118	.163

Note. None of the correlations that included the performance measures were statistically significant. PosIntsect = orientation of the position axis to the line of intersection; PosEnviro = orientation of the position axis in the environment; PosJoint = orientation of the position axis jointly to the line of intersection and the environment; RT = reaction time.

position axis to the line of intersection (with values of 1, 2, and 3), and a composite scale of both orientations (ranging from 1, for dual alignment, to 5, for maximally oblique). Correlation matrices relating these variables to the performance measures are given in Table 3. None of these correlations was statistically significant.

Mean performance pooled across variation of the position axes for each type of display that was sampled from Experiment 2 was compared with the appropriate mean data from that experiment. We did this to see whether the full range of movement of the double-headed arrow made a difference to performance in Experiment 3. There was no appreciable difference. The correlation between the comparable mean error in Experiments 2 and 3 was .94, $t(5) = 6.16$, $p < .001$; the correlation between mean RT was .88, $t(5) = 4.14$, $p < .01$.

Discussion

Effects of orientation often can be influenced by the introduction of local reference systems, with varying degrees of strength of the local system (e.g., Koffka, 1935; McAfee & Proffitt, 1991; Palmer, 1980; Pani & Dupree, 1994). In this experiment, a local axis formed by the positions of the two surfaces did not substantially influence performance, either positively or negatively. The negligible effects of this local reference system suggest that the environment-relative and object-relative determination of orientation is of overwhelming importance in this task. These results also have the more practical implication that the frontal-horizontal positioning of the surfaces in Experiments 1 and 2 was not important to the outcomes of those experiments.

General Discussion

If reasoning about fundamental relations in 3-D space were carried out through mental imagery that was an internalization of these relations, individuals in the present experiments would have accurately imagined any of the

pairs of surfaces to intersect at an edge. Clearly this did not occur. Although some edges were determined quickly and accurately, other edges definitely were not. The observed difficulties cannot be attributed to shortcomings in the displays or to extrinsic, arbitrary, or unfamiliar response scales (see Loomis, Da Silva, Fujita, & Fukusima, 1992; Proffitt et al., 1995). As in related experiments that have concerned physical reasoning, the displays either were physical objects or realistic stereo views. The methods of responding involved physical demonstration of the orientations (e.g., Pani, 1993; Pani et al., 1995; Pani et al., 1996; Pani et al., 1997).

The ability to accurately determine the lines of intersection of physically represented planes was affected by the degree of fit between perceived surface orientations and the implied lines of intersection, the orientations of the lines of intersection, and the angles between the surfaces. In our view, degree of fit affected the transition from spontaneous perception of the surfaces to determination of the lines of intersection. The importance of the orientation of the line of intersection was due to the consistent salience of singular orientations within prominent reference systems (the principal directions of the environment, in this case). The importance of the angle between the surfaces was due to the effectiveness of perpendicular angles in determining the object-relative orientations of surfaces. When the dihedral angle was well understood, the line of intersection could be found as one of its components. Interestingly, the time required to complete the task was affected most directly by the degree to which the surfaces had canonical orientations relative to the environment.

We think it is important that variation on three different parameters of orientation independently affected the accuracy of performance on this task. On the other hand, this flexibility did not extend to a fourth variable. The attempt in Experiment 3 to add a useful local reference system by manipulating the relative positions of the surfaces was not effective. The position axis had little effect on performance, even when it was aligned with the line of intersection. Of course, it is always possible to increase the salience of local reference systems (see Palmer, 1989; Pani & Dupree, 1994). However, it seems that the environmental reference system and the organization of the dihedral angle are especially strong in this task.

Reasoning About the Intersection of Two Planes

Physical Reasoning

The results of the experiments reported in this article support and extend a view of physical reasoning that we have been developing (e.g., Pani, 1994, in press-a, in press-b; Pani et al., 1996). Consider in somewhat more detail the basic elements of the task. First, there is linguistic identification of perceived objects, their properties, and hypothetical events, in such terms as *surface*, *plane*, *orientation*, *intersection*, and *line*. These terms refer to things, properties, and relations in 3-D space and are generally associated with environmental, sometimes egocentric, but rarely retinal,

coordinate systems (e.g., Pani, in press-a; Pani & Dupree, 1994; Rock, 1983). Reasoning includes application of this linguistic statement of the problem to the description of the physical situation given in spontaneous perception.

One part of the application of the problem statement is identification of a *target organization* for the physical phenomenon, a cognitive goal toward which processing will be aimed. Such target organizations can vary in specificity. For example, the description of an organization might be as general as "a single line in 3-D space that would be common to both planes." When a physical phenomenon is sufficiently well defined, however, the individual has a *base organization* of the phenomenon that is applied to the spontaneously perceived situation. A base organization is a description of the essential relations that define a physical structure or event (e.g., a rotation, a projective transformation, or the joining of two surfaces). With the statement of the problem as a guide, the individual attends to particular aspects of the perceptual description of the situation and tries to extend them into the base organization that is needed. If the base organization is established, there can be what Wertheimer (1959) would call "productive thinking": A variety of physical relations and events can be determined readily.

Seeing a base organization in a particular physical arrangement can be quite challenging. One problem related to orientation is that the spatial reference system that is needed for the base organization may not be the same as the one used in spontaneous perception. For example, a line of intersection within a dihedral angle is part of the object-relative specification of the orientation of two surfaces, but orientation of the surfaces is first specified relative to the environment (at least in the task studied in the present experiments). Hence, the dihedral angle often is not seen, and cognitive procedures are turned toward building an organization that will succeed in the task. At this point, physical reasoning may become quite an active, heuristic, and time-consuming process.

Procedures for Finding Planar Intersections

One very general procedure that is available for determining lines of intersection is a search for the one line that is common to both surfaces. But even such a constrained search in 3-D space can be quite difficult (e.g., Pani et al., 1996). A heuristic version of this search is reflected in the variable of degree of fit. Individuals perceive the surfaces, encode a candidate set of lines in the orientation lines of one surface, and check whether any of those lines coincide with the other surface. We are not suggesting that this would be an explicit strategy. Rather, we are attempting to cast in logical terms the individual's spatial intuitions that in some cases the two surfaces appear to fit together.

The potential for noticing such a fit involves taking lines that originate in goal-directed perception and considering them as candidates for a line that is needed in reasoning. Noticing the fit involves a kind of transitive relation, in which a salient element in one spatial description is seen to work successfully in another description.

The occurrence of such a transitive process also is

responsible for the importance of variation in the orientation of the line of intersection. Canonical orientations within a salient reference system are ready candidates for elements needed in physical reasoning. When a person searches for the one line common to two surfaces, and that line is a canonical axis of the environment, the line is seen, remembered, and indicated easily. If the line already has been determined, so that transitivity is not strictly necessary, there still is reinforcement of the line, making it more salient and stable.

The importance of the angle between the surfaces in the determination of planar intersections stems from the fact that people are substantially more skilled in seeing and imagining perpendicular angles than oblique angles (see below). The presence of perpendicular angles apparently helped participants to organize the object-relative dihedral angles in these experiments. A line of intersection was then available through two properties of the dihedral angle: The line is at the apex of the angle, and it is perpendicular to the plane that contains the angle. Note again that the organization of the dihedral angle has an important difference from degree of fit and from the orientation of the line of intersection. Visualizing the dihedral angle involves seeing the surfaces and the line of intersection as a purely object-relative configuration that is embedded, as a whole, within the environment (much as an object has a major axis with an orientation to the environment; see Pani, in press-a).

The degree of fit, the orientation of the line of intersection, and the dihedral angle are partially independent physical variables. Each of them has a component that may be aligned with the line of intersection. Further, each can reinforce the others when they are mutually aligned. When none of these variables imply the line of intersection, the individual is left with relatively exhaustive spatial searches and comparisons, and these can become quite complex and difficult (e.g., Pani et al., 1996). On the other hand, we suspect that people can gain skill in this task (indeed, it is probably a good test of spatial ability). Our hypothesis is that people become more adept at seeing the purely object-relative dihedral angle (see also Pani & Dupree, 1994, Experiment 2).

Canonical Physical Properties

The easier problems in physical reasoning are characterized by a simplicity of organization. This simplicity includes the fundamental structures used to describe spatial relations (e.g., polar reference systems), the base organizations of physical phenomena (e.g., the dihedral angle), and the application of this understanding to particular physical problems (e.g., use of various kinds of alignment in determining lines of intersection). Simplicity of organization can be explained in terms of the importance for cognition of certain types of physical properties. These fall into three nested sets. The most inclusive set is the *singular properties*. Included within the singularities are the *symmetries*. *Aligned orientations* are important instances of singularity and symmetry in the domain of spatial relations. Dependence on singularity, symmetry, and alignment leads to descriptions of spatial

relations that are distinctive, nonarbitrary, efficient, and pragmatically useful.

Singularity

Singular properties are qualitatively unique values that arise in comparisons along dimensions of variation. Fundamental properties that are uniquely singular are *maximum*, *minimum*, *same* (or *invariant*), and *orthogonal* (or *independent*). A more specific set of singular terms that apply to spatial structures includes, among many others, *straight*, *circular*, *aligned*, *parallel*, *perpendicular*, *vertical*, *horizontal*, *symmetric*, *midpoint*, *center*, *top*, and *front* (see Goldmeier, 1982).

Singular properties are nonmetric. One basis of the cognitive value of singular properties is that they are nonmetric (see also Hummel & Biederman, 1992; Kosslyn, 1994). Consider, for example, the term perpendicular. If two sticks are hinged so that they can form some version of a T, a perpendicular angle can be formed by swiveling the sticks about the hinge. It will not be necessary to use a protractor. Rather, the lines are perpendicular when the angles on both sides of the central leg are equal (i.e., when the spaces on both sides are the same shape). This is an ordinal judgment that converges on a single value. In like manner, a protractor is not needed to judge when two sticks are parallel. Parallel is the single value at which the angular difference between the sticks disappears (i.e., if measured, it is at zero).

The nonmetric aspect of singular properties tends to put them at the same level as categories of properties. For example, parallel and perpendicular are coordinate with acute and obtuse. The latter terms are categories of angles, whereas parallel and perpendicular are each a single angle. They are, as it were, categories with single members. Nonmetric properties at a categorical level will be distinctive, and this will make them salient and memorable.

Singular properties are nonarbitrary. A second advantage for cognition of singular properties is that they are nonarbitrary. For example, a dihedral angle is measured in the equatorial plane of the line of intersection of the two surfaces (see Figure 6). This equatorial plane is perpendicular to the line of intersection. Consider, however, that a different plane could be used to measure the angle (e.g., one at 45° to the line). In this case, however, the measured angle would be different. The equatorial plane is the one nonarbitrary plane for measuring the angle between two surfaces because it is the plane for which the angle is maximized. In a similar way, the slant of a surface is given by the line within the surface that has the minimum angle to the vertical (i.e., that is steepest). Other lines could be chosen, and then the slant of the surface would be different. The slant line is a minimum (or maximum, relative to the horizontal) and is the one nonarbitrary selection for determination of the slant.

Symmetry

An important subset of the singular properties are the symmetries. In the sciences in general, symmetry refers to the regularity of physical structures, to those properties of

things that are invariant across conceivable transformations (e.g., Cederberg, 1989; Hargittai & Hargittai, 1994; Stewart & Golubitsky, 1992; Weyl, 1952). Different types of symmetry correspond to different transformations, and some symmetries are associated with multiple transformations. For example, the bilateral symmetry of mammals is an instance of reflection symmetry (i.e., invariant across reflection), and the prototypic flower has both rotation and reflection symmetry (i.e., invariant across rotation and reflection). The salient uniformity of the orientations of edges and surfaces in generalized cylinders, such as the shape at the left in Figure 1, can be considered a type of symmetry (Pani, 1994; Pani, in press-a; Pani et al., 1997; Zerroug & Nevatia, 1994).

Aligned orientations have symmetry. For example, the transitivity from an orientation line to a parallel line of intersection is an instance of translational symmetry between the two lines. A perpendicular angle between two planes can be recognized, or constructed, by the reflection symmetry that it contains (see Figure 6).

Symmetry is a repetition of structure throughout an object or system, and where there is repetition (or redundancy), there is the possibility of efficient description (e.g., Attneave, 1954; Garner, 1974; Gell-Mann, 1994; Leeuwenberg, 1971; Pani, in press-a). Symmetry in the organizations of physical phenomena, in perceived situations, and in the fit between them reduces the set of terms needed to describe the world and its possibilities. Thus, the organization of the dihedral angle not only is distinctive and nonarbitrary for being singular, it also is part of a relatively simple structure because it is measured in a plane that is symmetric with respect to the line of intersection and the two surfaces (see Figure 6).

Aligned Orientations

Aligned orientations include parallel, perpendicular, vertical, and horizontal. A somewhat broader class of aligned properties includes coincident and colinear, in which both orientation and position in space are encoded. All of these properties are examples within the domain of spatial relations of singular and symmetric terms. Hence, these terms refer to properties that are distinctive, nonarbitrary, and simple (i.e., regular). Parallel orientations have the advantages of relations that include *same*: collapsing of descriptive parameters into single values and cognitive processing that uses simple continuation or repetition. Perpendicular orientations have a variety of advantages related to their being orthogonal to their reference orientations as well as being a symmetric midpoint, or balance, within the range of orientations.

In general, aligned orientations permit efficient description of the nature of physical phenomena as well as the efficient application of this understanding in particular situations. Hence, they are ubiquitous in fundamental descriptive structures (e.g., the polar reference system), in base organizations (e.g., the dihedral angle and the axis and planes of rotational motion), and in the fit between spontaneous perception and cognitive organization (e.g., the transitivity and reinforcement observed in the present experiments).

Pragmatism

The physical structures and systems that people comprehend readily are a small subset of those that are physically possible (e.g., Hinton, 1979; Pani et al., 1996; Pani et al., 1995; Pani et al., 1997; Proffitt, Kaiser, & Whelan, 1990). Why would this be so? One answer is that we understand those things that are relatively distinctive, nonarbitrary, and simple, as attested by the frequency of singularity, symmetry, and alignment in our physical understanding. Certainly this is true. On the other hand, no cognitive system, especially when it is capable of learning, should sacrifice effective use of information for efficiency of operation. Perhaps humans have not made such a sacrifice.

We suggest that the singularity and symmetry that make certain physical relations distinctive and simple also are a basis for the physical relations that people find useful (Pani, in press-a). For example, we understand a small subset of the physically possible rotations (Pani et al., 1995), but it is just these rotations that are used in hinges, pulleys, wheels, knobs, cranks, and dials. The other rotations tend to characterize tumbling boxes and falling leaves. They are sometimes interesting to look at, but the character of the rotation is of little utility. More generally, it is spatial relations that are partitioned into orthogonal dimensions, that are balanced, aligned, symmetric, and that allow a fit between things that provide a basis for useful construction across a variety of physical contexts. Both farmers and bees have mastered the art of creating regular lattices, but in neither case was aesthetic beauty the goal. It is a happy coincidence for human cognition that what tends to be distinctive and simple tends also to be useful.

Related Statements Concerning Physical Reasoning

The view of physical reasoning expressed in this article has a number of important elements in common with other modern statements concerning perception and cognition. The concept of a base organization of a physical phenomenon has important similarities to mental models (e.g., Gentner & Stevens, 1983; Johnson-Laird, 1983). Other authors have pointed to the importance of a cooperation between perception and problem solving (e.g., Hayes, 1973; Larkin, 1989; see also Pani, 1996). Our emphasis on the value of alignment, symmetry, and singularity resonates with numerous other suggestions that canonical forms of structure permit essential simplification for higher cognition (e.g., Biederman, 1987; Goldmeier, 1982; Palmer, Rosch, & Chase, 1981; Rosch, 1975; Tversky, 1981; Wertheimer, 1950). Perhaps the view most like our own is that of Proffitt, Gilden, Kaiser, and their colleagues (e.g., Kaiser, Proffitt, Whelan, & Hecht, 1992; Proffitt & Gilden, 1989; Proffitt, Kaiser, & Whelan, 1990). Their concept of unidimensional heuristics in reasoning has important similarities to our suggestion that people depend on singularity, symmetry, and a fit between descriptions.

More generally, there is a growing effort to distinguish between instances of reasoning that people find intuitively simple and ostensibly similar instances that require ad-

vanced expertise (Cheng & Holyoak, 1985; Clement, 1982; Johnson-Laird, 1983; Kahneman, Slovic, & Tversky, 1982; McCloskey, 1983; Pani, in press-a; Proffitt & Gilden, 1989). An important job for cognitive psychology is to identify the boundaries between what is intuitive and what is complex across a variety of domains of human experience. So far, it appears that everyday understanding of physical phenomena is confined to identifiable classes of phenomena with high degrees of singularity and invariance. Moving outside that domain involves a commitment to advanced expertise.

The theoretical underpinnings of cognitive organization in higher vision, memory, mental imagery, and physical reasoning have become, in many ways, a development of theories of structural description. This is not a turning away from grounding perception theory in an understanding of the perceptual world (Pani, in press-a). Rather, it is a specification of the properties of the world that are in fact detected and encoded at various levels of cognitive structure and activity. Such a specification permits tracing the transitions from perception through reasoning and back again to intelligent behavior.

References

- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*, 183–193.
- Attneave, F. (1971). Multistability in perception. *Scientific American*, *225*, 62–71.
- Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Braunstein, M. L. (1976). *Depth perception through motion*. New York: Academic Press.
- Cederberg, J. N. (1989). *A course in modern geometries*. New York: Springer-Verlag.
- Chase, W. G. (1986). Visual information processing. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance: Vol. 2. Cognitive processes and performance* (pp. 1–71). New York: Wiley.
- Cheng, P. N., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391–416.
- Clement, J. (1982). Students' preconceptions in introductory mechanics. *American Journal of Physics*, *50*, 66–71.
- Corballis, M. C., Nagourney, B. A., Shetzer, L. I., & Stefanatos, G. (1978). Mental rotation under head tilt: Factors influencing the location of the subjective reference frame. *Perception & Psychophysics*, *24*, 263–273.
- Cutting, J. E., & Millard, R. T. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General*, *113*, 198–216.
- Epstein, W., Babler, T., & Bownds, S. (1992). Attentional demands of processing shape in three-dimensional space: Evidence from visual search and precuing paradigms. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 503–511.
- Finke, R. A. (1989). *Principles of mental imagery*. Cambridge, MA: MIT Press.
- Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
- Gasson, P. C. (1983). *The geometry of spatial forms*. Chichester, England: Ellis Horwood.
- Gell-Mann, M. (1994). *The quark and the jaguar: Adventures in the simple and the complex*. New York: Freeman.
- Gentner, D., & Stevens, A. L. (1983). *Mental models*. Hillsdale, NJ: Erlbaum.
- Gibson, J. J. (1950). The perception of visual surfaces. *American Journal of Psychology*, *63*, 367–384.
- Gibson, J. J., & Cornsweet, J. (1952). The perceived slant of visual surfaces—optical and geographical. *Journal of Experimental Psychology*, *44*, 11–15.
- Goldmeier, E. (1982). *The memory trace: Its formation and its fate*. Hillsdale, NJ: Erlbaum.
- Hargittai, I., & Hargittai, M. (1994). *Symmetry: A unifying concept*. Bolinas, CA: Shelter Publications.
- Hayes, J. R. (1973). On the function of visual imagery in elementary mathematics. In W. G. Chase (Ed.), *Visual information processing* (pp. 177–214). New York: Academic Press.
- Hecht, H., & Proffitt, D. R. (1995). The price of expertise: Effects of experience on the water-level task. *Psychological Science*, *6*, 90–95.
- Hinton, G. (1979). Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science*, *3*, 231–250.
- Hinton, G. E., & Parsons, L. M. (1988). Scene-based and viewer-centered representations for comparing shapes. *Cognition*, *30*, 1–35.
- Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, *99*, 480–517.
- Johnson-Laird, P. N. (1983). *Mental models*. Cambridge, MA: Harvard University Press.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Kaiser, M. K., Proffitt, D. R., Whelan, S. M., & Hecht, H. (1992). Influence of animation on dynamical judgments. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 669–690.
- Kinsella-Shaw, J. M., Shaw, B., & Turvey, M. T. (1992). Perceiving “walk-on-able” slopes. *Ecological Psychology*, *4*, 223–239.
- Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, *52*, 487–496.
- Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt, Brace.
- Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S. M. (1994). *Image and brain*. Cambridge, MA: MIT Press.
- Kubovy, M. (1994). The perceptual organization of dot lattices. *Psychonomic Bulletin & Review*, *1*, 182–190.
- Larkin, J. H. (1989). Display-based problem solving. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon*. Hillsdale, NJ: Erlbaum.
- Leeuwenberg, E. L. J. (1971). A perceptual coding language for visual and auditory patterns. *American Journal of Psychology*, *84*, 307–349.
- Lehtinen-Railo, S., & Jurmaa, J. (1994). Effect of visual experience on locational judgements after perspective change in small-scale space. *Scandinavian Journal of Psychology*, *35*, 175–183.
- Loomis, J. M., Da Silva, J. A., Fujita, N., & Fukusima, S. S. (1992). Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 906–921.
- Mach, E. (1959). *The analysis of sensations* (5th ed.) (C. M. Williams & S. Waterlow, Trans.). New York: Dover. (Original work published 1906)
- Marr, D. (1982). *Vision*. New York: Freeman.

- McAfee, E. A., & Proffitt, D. R. (1991). Understanding the surface orientation of liquids. *Cognitive Psychology*, *23*, 483–514.
- McCloskey, M. (1983). Naive theories of motion. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 299–324). Hillsdale, NJ: Erlbaum.
- Osserman, R. (1995). *Poetry of the universe*. New York: Anchor Press.
- Palmer, S. E. (1980). What makes triangles point: Local and global effects in configuration of ambiguous triangles. *Cognitive Psychology*, *12*, 285–305.
- Palmer, S. E. (1989). Reference frames in the perception of shape and orientation. In B. E. Shepp & S. Ballesteros (Eds.), *Object perception: Structure and process* (pp. 121–163). Hillsdale, NJ: Erlbaum.
- Palmer, S. E., & Hemenway, K. (1978). Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 691–702.
- Palmer, S. E., Rosch, E., & Chase, P. (1981). Canonical perspective and the perception of objects. In J. Long & A. D. Baddeley (Eds.), *Attention and performance*. *9*, 135–151. Hillsdale, NJ: Erlbaum.
- Pani, J. R. (1993). Limits on the comprehension of rotational motion: Mental imagery of rotations with oblique components. *Perception*, *22*, 785–808.
- Pani, J. R. (1994). The generalized cone in human spatial organization. *Spatial Vision*, *8*, 491–502.
- Pani, J. R. (1996). Mental imagery as the adaptationist views it. *Consciousness and Cognition*, *5*, 288–326.
- Pani, J. R. (in press-a). Descriptions of orientation and structure in perception and physical reasoning. In E. Winograd, R. Fivush, & W. Hirst (Eds.), *Ecological approaches to cognition: Essays in honor of Ulric Neisser*. Hillsdale, NJ: Erlbaum.
- Pani, J. R. (in press-b). Descriptions of orientation in physical reasoning. *Current Directions in Psychological Science*.
- Pani, J. R., & Dupree, D. (1994). Spatial reference systems in the comprehension of rotational motion. *Perception*, *23*, 929–946.
- Pani, J. R., Jeffres, J. A., Shippey, G., & Schwartz, K. (1996). Imagining projective transformations: Aligned orientations in spatial organization. *Cognitive Psychology*, *31*, 125–167.
- Pani, J. R., William, C. T., & Shippey, G. T. (1995). Determinants of the perception of rotational motion: Orientation of the motion to the object and to the environment. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1441–1456.
- Pani, J. R., Zhou, H., & Friend, S. M. (1997). Perceiving and imagining Plato's solids: The generalized cylinder in spatial organization of 3D structures. *Visual Cognition*, *4*, 225–264.
- Parsons, L. M. (1995). Inability to reason about an object's orientation using an axis and angle of rotation. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1259–1277.
- Proffitt, D. R., Bhalla, M., Gossweiler, R., & Midgett, J. (1995). Perceiving geographical slant. *Psychonomic Bulletin & Review*, *2*, 409–428.
- Proffitt, D. R., & Gilden, D. L. (1989). Understanding natural dynamics. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 384–393.
- Proffitt, D. R., Kaiser, M. K., & Whelan, S. M. (1990). Understanding wheel dynamics. *Cognitive Psychology*, *22*, 342–373.
- Pylyshyn, Z. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, *88*, 16–45.
- Rock, I. (1973). *Orientation and form*. New York: Academic Press.
- Rock, I. (1983). *The logic of perception*. Cambridge, MA: MIT Press.
- Rosch, E. (1975). Cognitive reference points. *Cognitive Psychology*, *7*, 532–547.
- Sedgwick, H. A. (1983). Environment-centered representation of spatial layout: Available visual information from texture and perspective. In A. Rosenthal & J. Beck (Eds.), *Human and Machine Vision* (pp. 425–458). New York: Academic Press.
- Sedgwick, H. A. (1986). Space perception. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance*. Vol. 1. *Sensory processes and perception* (pp. 21-1–21-57). New York: Wiley.
- Sedgwick, H. A., & Levy, S. (1985). Environment-centered and viewer-centered perception of surface orientation. *Computer Vision, Graphics, and Image Processing*, *31*, 248–260.
- Shepard, R. N. (1978). The mental image. *American Psychologist*, *33*, 125–137.
- Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review*, *91*, 417–447.
- Shiffrar, M. M., & Shepard, R. N. (1991). Comparison of cube rotations around axes inclined relative to the environment or to the cube. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 44–54.
- Stevens, K. A. (1981). The information content of texture gradients. *Biological Cybernetics*, *42*, 95–105.
- Stevens, K. A. (1983a). Slant-tilt: The visual encoding of surface orientation. *Biological Cybernetics*, *46*, 183–195.
- Stevens, K. A. (1983b). Surface tilt (the direction of slant): A neglected psychophysical variable. *Perception & Psychophysics*, *33*, 241–250.
- Stewart, I., & Golubitsky, M. (1992). *Fearful symmetry: Is God a geometer?* Cambridge, MA: Blackwell.
- Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology*, *13*, 407–433.
- Wertheimer, M. (1950). Laws of organization in perceptual forms. In W. D. Ellis (Ed.), *A source book of Gestalt psychology* (pp. 71–88). New York: Humanities Press. (Original work published 1923)
- Wertheimer, M. (1959). *Productive Thinking*. New York: Harper.
- Weyl, H. (1952). *Symmetry*. Princeton, NJ: Princeton University Press.
- Zerroug, M., & Nevatia, R. (1994). *Segmentation and recovery of SHGCs from a real intensity image*. Paper presented at the European Conference on Computer Vision, Stockholm, Sweden.

Appendix

Orientations of the Disks in Experiment 1

Below are spherical coordinates in degrees for the orientations of the 13 lines of intersection and the 78 pairs of disks in Experiment 1. For the lines of intersection, (0, 0) refers to a vertical line. A positive number in the first coordinate refers to a rotation around the x axis of the environment, with the top coming forward. A positive number in the second coordinate refers to a subsequent rotation around the y (vertical) axis of the environment, with the

front moving from left to right.

For the disks, (0, 0) refers to a horizontal disk (i.e., with a vertical surface normal). The rotations corresponding to coordinates are the same as for the lines of intersection. In the experiment, all six possible pairs of disk orientations were formed from the four disk orientations corresponding to each line of intersection.

Intersection: (0, 0) Disks: (90, -45) (90, 0) (90, 45) (90, 90)	Intersection: (90, -45) Disks: (135, -135) (0, 0) (135, 45) (90, 45)	Intersection: (45, -45) Disks: (120, -100) (135, -45) (120, 10) (90, 45)
Intersection: (90, 90) Disks: (135, 0) (0, 0) (45, 0) (90, 0)	Intersection: (90, 45) Disks: (135, -45) (0, 0) (135, 135) (90, 135)	Intersection: (45, 45) Disks: (120, -10) (135, 45) (120, 100) (90, 135)
Intersection: (90, 0) Disks: (135, -90) (0, 0) (135, 90) (90, 90)	Intersection: (45, -90) Disks: (120, -145) (135, -90) (120, -35) (90, 0)	Intersection: (135, 45) Disks: (60, -10) (45, 45) (60, 100) (90, 135)
Intersection: (45, 0) Disks: (120, -55) (135, 0) (120, 55) (90, 90)	Intersection: (45, 90) Disks: (120, 35) (135, 90) (120, 145) (90, 0)	Intersection: (135, -45) Disks: (60, -100) (45, -45) (60, 10) (90, 45)
Intersection: (-45, 0) Disks: (60, -55) (45, 0) (60, 55) (90, 90)		