Misperceptions of Angular Velocities Influence the Perception of Rigidity in the Kinetic Depth Effect

Fulvio Domini and Corrado Caudek
Cognitive Technology Laboratory, AREA Science Park

Dennis R. Proffitt
University of Virginia

Accuracy in discriminating rigid from nonrigid motion was investigated for orthographic projections of three-dimensional rotating objects. In 3 experiments the hypothesis that magnitudes of angular velocity are misperceived in the kinetic depth effect was tested, and in 4 other experiments the hypothesis that misperceiving angular velocities leads to misperceiving rigidity was tested. The principal findings were (a) the magnitude of perceived angular velocity is derived heuristically as a function of a property of the first-order optic flow called *deformation* and (b) perceptual performance in discriminating rigid from nonrigid motion is accurate in cases when the variability of the deformations of the individual triplets of points of the stimulus displays favors this interpretation and not accurate in other cases.

The human perceptual system is capable of extracting three-dimensional (3-D) information from moving images from which every static pictorial cue to depth has been removed, a phenomenon called the kinetic depth effect (Wallach & O'Connell, 1953). Numerous attempts to reach a theoretical understanding of this phenomenon have been made (Bennett & Hoffman, 1985; Bennett, Hoffman, Nicola, & Prakash, 1989; Koenderink, 1986; Koenderink & Van Doorn, 1975, 1977; Longuet-Higgins & Prazdny, 1980; Prazdny, 1980; Ullman, 1979, 1983, 1984). Mathematical analyses have revealed that the recovery of 3-D properties from projected motions is characterized by an inherent ambiguity: The mapping from two-dimensional (2-D) images to 3-D properties is one to many (i.e., different 3-D motions project to the same 2-D image). Without a priori constraints on the nature of the structure or the motion of the projected objects, the problem of finding a unique 3-D interpretation for a moving image (the so-called *Structure-from-Motion*, or SfM, problem) cannot be solved. One way to overcome this inherent ambiguity is to introduce constraints in the interpretation process in order to restrict the space of possible interpretations to only one solution. The constraints that have been used, are, for example, the rigidity assumption (Ullman, 1979), smoothness of flow field (Hildreth, 1984), rotation at a constant angular velocity (Hoffman & Bennett, 1985, 1986), and so on. The rigidity assumption has been used in many computer vision algorithms; moreover, in some psychological theories (Gibson, 1979; Johansson, 1978; Musatti, 1924) it has been hypothesized that the perceptual recovery of 3-D shape from motion could be based on a rigidity constraint (see Cutting, 1987; Zanforlin, 1988).

In current research the psychological plausibility of the rigidity assumption has been examined by studying human performance in the minimal conditions theoretically necessary for discriminating rigid from nonrigid motion (Braunstein, Hoffman, & Pollick, 1990; Todd & Bressan, 1990). The experiments have been motivated by the theoretical finding that an ideal observer can discriminate rigid from nonrigid 3-D transformations from two orthographic views of four points. The computational model, based on the theorem of Ullan (1977), considers as rigid all the displays for which it is possible to subtract a common component of image rotation in order to keep all the trajectories parallel. If a common component of curl does not exist, the display is considered nonrigid. The results of Braunstein et al. (1990) and Todd and Bressan (1990) support the idea that human observers can perform such discriminations.

Braunstein et al. (1990) found that human observers can discriminate rigid from nonrigid motion when viewing only two views of four points. Their participants were presented with orthographic projections of rigid and nonrigid rotations. The nonrigid stimuli were generated by having each point in the displays rotate about a different axis of rotation; the rigid stimuli were generated by having all points rotate

about the same axis. In both cases, all points in the displays were rotated by the same amount.[1]

We replicated the stimuli used by Braunstein et al. (1990; see our Experiment 4) and computed the variance of the trajectories of the points in the displays.[2] We then performed an analysis of variance (ANOVA) on the stimulus displays by using the variance of the trajectories of each stimulus display as the dependent variable. The independent variables were 3-D rigid versus 3-D nonrigid displays and number of points.

We found that the variance of the trajectories significantly differed for rigid and nonrigid displays, $F(1, 79) = 70.83$, $p < .001$: Mean trajectory variance for nonrigid displays was 70% larger than that for rigid displays. Neither the main effect of number of points on the variance of the trajectories nor the interaction of number of points with 3-D rigid versus nonrigid displays was significant. Braunstein et al. (1990) also reported that in their Experiment 3, performance increased with the number of views. We ran another simulation (with 60 signal trials and 60 noise trials) in which we looked at each frame transition for stimuli with the same parameters as those used by Braunstein et al. in their Experiment 3. Again, we found that the variance of the trajectories was 84% larger for the nonrigid displays, $F(1, 118) = 12.41$, $p < .001$. Moreover, if we consider each frame transition as an independent trial, then it follows that the probability of a correct response would increase with the number of frame transitions, as effectively found by Braunstein et al.

Another demonstration of the ability of human observers to discriminate between rigid and nonrigid motion near the minimum level at which discrimination is theoretically possible has been provided by Todd and Bressan (1990). Their participants were shown displays made up of two line segments rotating in depth. During rotation, the 3-D length of one line segment remained constant, whereas the 3-D length of the other line segment changed. Observers were asked to indicate which line segment was undergoing a nonrigid change in length. In this experiment, the independent variables were the percentage of change in the 3-D length of the nonrigid line segment at each frame transition (1%, 2%, 3%, and 4%) and the number of views in the stimulus displays (2, 4, or 8 views). Todd and Bressan found that observers were able to identify the nonrigid line segments and that accuracy increased with the percentage of change in the 3-D length. No effect of number of frames and no significant interactions were found.

We ran a simulation on the stimulus displays used by Todd and Bressan (1990), and we computed the absolute value of the variation in 2-D length (from the first to the last frame) of each line segment in each stimulus display. Using these values, we performed two analyses. First, we conducted an ANOVA on the stimulus displays, using the 2-D variation in length of the line segments as the dependent variable and rigidity (3-D rigid vs. 3-D nonrigid line segments) as the independent variable. We found that the 2-D variation in length differed significantly for 3-D rigid and 3-D nonrigid line segments, $F(1, 4799) = 190.50, p < .001$: The mean variation in the length of the 3-D nonrigid line

segments was 28% larger than the mean variation in the length of the 3-D rigid line segments. Second, for each experimental condition used by Todd and Bressan, we computed the ratio between the mean 2-D variation in the length of the 3-D rigid line segments and the mean 2-D variation in the length of the 3-D nonrigid line segments. These ratios have been rescaled and plotted in Figure 1 together with the experimental results obtained by Todd and Bressan.

The similarity between the experimental and simulation data, together with the results of the simulation of the stimuli used by Braunstein et al. (1990), suggest that, in both studies, there are alternative explanations for the reported accuracy in discriminating rigid from nonrigid motion. The reason we consider the alternative explanations to be more plausible is directly related to the results of Experiment 4 in this article. In this experiment, we found that stimuli compatible with two orthographic views of a 3-D rigid motion are perceived as nonrigid and that stimuli not compatible with two orthographic views of a 3-D rigid motion are perceived as rigid. We therefore question the view that the discrimination between rigid and nonrigid stimuli is based on a process that checks for the existence of a common component of curl in the optic flow. We suggest that the classification that the perceptual system performs in order to separate perceived rigid from nonrigid stimuli is performed heuristically and depends on the characteristics of the two classes of stimuli that are used in the discrimination task. In the Braunstein et al. stimuli the heuristic analysis may be based on the variance of the trajectories, and in the Todd and Bressan (1990) stimuli it may be based on the 2-D length variation. Variance of trajectories and 2-D length variation are stimulus characteristics that are specific to the type of nonrigidities that were created. We generated a new type of nonrigidity in order to isolate the effect of a first-order property of the optic flow (deformation) that we hypothesized would influence the judgments in the discrimination task.

Discrimination Between Rigid and Nonrigid Motion and Variance of Deformations

Two hypotheses form the basis of this article: (a) Angular rotations can be misperceived because they depend on a first-order variable of the optic flow, the *deformation*, and not, in general, on the simulated rotations. (b) Judgments of rigidity depend on the magnitudes of rotation perceived for the component parts of a moving object; objects are judged

---

[1] Braunstein et al. (1990) simulated the nonrigid displays by rotating each point about a different axis in order to keep their measure of 2-D nonrigidity equal for the rigid and nonrigid stimuli.

[2] The *variance of the trajectories* was defined as the variance of the arctangent (ranging between $+90°$ and $-90°$) of the angular coefficients of the 2-D displacements of the individual points. (Specifically, let $\mathbf{p}_{ij}$ denote the position of the point $\mathbf{p}_i$ in view j. Let $\mathbf{m}_{ijj'}$ denote the angular coefficient of the line connecting the 2-D positions of the point $\mathbf{p}_i$ in the views j and j'. Then $\sigma^2_{jj'}$ is the variance of the arctangent of $\mathbf{m}_{ijj'}$ for all i.)
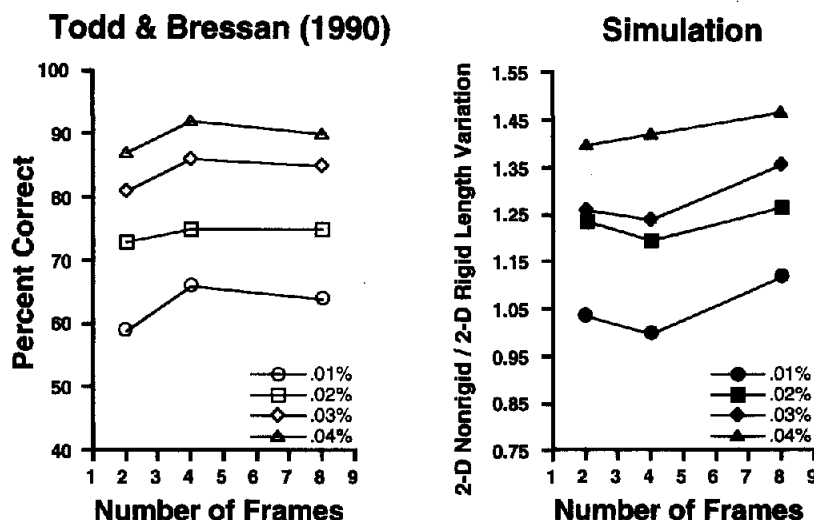
Todd & Bressan (1990)

Simulation

Figure 1. The percentage of correct rigidity discriminations obtained by Todd and Bressan (1990; left panel), and the outcome of a simulation performed on the ratio between the mean two-dimensional (2-D) variations in the length of the rigid and nonrigid stimuli used by Todd and Bressan (right panel), as a function of the number of frames and for four levels of nonrigidity length change.

to move *rigidly* if all their component parts are perceived to rotate by the same amount and to move *nonrigidly* if their component parts are associated with different magnitudes of perceived rotation. The misperceptions of angular rotation predicted by the first hypothesis lead to misperceptions of rigidity.

In order to explain the notion of deformation, let us start by considering a planar patch Π. The orientation of a planar patch in 3-D space can be described in terms of its slant (σ) and tilt (τ). *Slant* is defined as the tangent of the angle between the line of sight (i.e., the z-axis) and the normal to the patch. This angle varies over a range of 90°, and slant is equal to zero if the patch lies perpendicular to the line of sight (i.e., parallel to the x–y plane). *Tilt* is defined as the angle between the projection of the normal to the patch and the x-axis.

Let us consider the optic flow produced by the orthogonal projection of a patch having slant σ and tilt τ and undergoing a generic 3-D rigid motion. The differential of the optic flow can be decomposed into four components (the differential invariants) called *curl, div, def*$_1$ and *def*$_2$. The *curl* component describes a pure rigid rotation in the image plane; the *div* component describes an isotropic contraction or expansion; the *def*$_1$ and *def*$_2$ components describe two orthogonal shears (Koenderink, 1986; Todorovic, 1993). Koenderink and Van Doorn (1986) have shown that only the shear components are informative about the 3-D shape that has produced the projected optic flow. Therefore, we focus our attention on these components. It is easy to show that the square root of the sum of the squared *def*$_1$ and *def*$_2$ is equal to the product of the slant of the planar patch (σ) and its component of angular velocity (ω) parallel to the image plane. We call this quantity *deformation (def;* see

also, Koenderink, 1986; Koenderink & Van Doorn, 1976, 1986):

$$def = \sqrt{def_1^2 + def_2^2} = \sigma\omega. \qquad (1)$$

One can describe *def* by considering the 2-D transformations produced by only two views of a moving object; *def,* therefore, is a property of the first-order optic flow (i.e., the velocity field defined by two views of a moving object). It is important to point out that the first-order optic flow is inherently ambiguous: There is not a one-to-one correspondence between *def* and the two parameters by which it is defined (i.e., the slant, σ, and the rotation component parallel to the image plane, ω). The first-order optic flow, in fact, is compatible with a one-parameter family of solutions for these two parameters, which can be represented by the loci of points of the hyperbola described by Equation 1 (see Figure 2).

For a generic orientation of a moving patch Π, one can compute *def* by measuring the image plane position of the orthogonal projection of three points P_0, P_1, and P_2 that lie on the patch and their instantaneous 2-D velocities V_0, V_1, and V_2, as indicated by Equation 2 (see Figure 3):

$$def = \frac{1}{|\sin \alpha|}$$

$$\sqrt{\left(\frac{|V_1 - V_0|}{\rho_1}\right)^2 + \left(\frac{|V_2 - V_0|}{\rho_2}\right)^2}$$

$$- 2\frac{|V_1 - V_0|}{\rho_1}\frac{|V_2 - V_0|}{\rho_2}\cos(\alpha - \alpha_\Delta), \qquad (2)$$
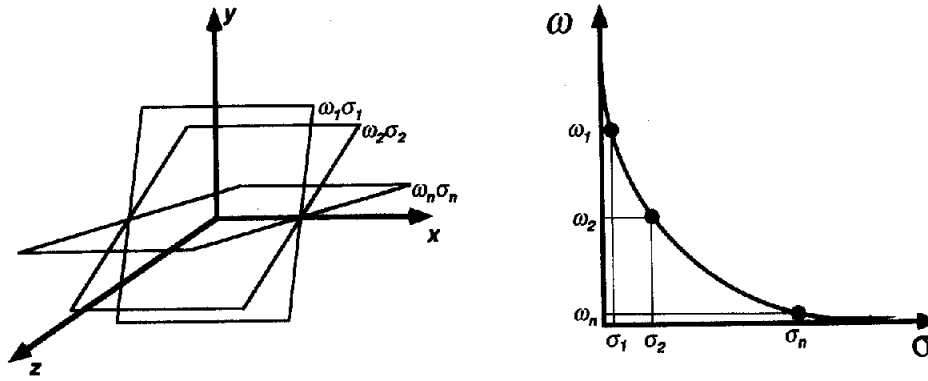
Figure 2. A two-dimensional transformation in the image plane corresponding to a given defor-
mation (*def*) can be produced by different rotations ω of planes with the same tilt but different slant
σ (left panel): $\omega_1\sigma_1$, $\omega_2\sigma_2$, . . ., $\omega_n\sigma_n$. The equilateral hyperbola in the right panel represents the loci
of the (σ, ω) pairs producing the same *def*.

where V_0, V_1, and V_2 are the velocity vectors of the points
P_0', P_1', and P_2', ρ_1 and ρ_2 are the distances of the points
P_1', P_2' from the point P_0', α is the angle between the line
segments $P_0'P_1'$ and $P_0'P_2'$, and α_A is the difference be-
tween the angles of the velocity vectors.[3]

It is important to notice that *def* is influenced by two
different aspects of the 2-D transformations: (a) *configural
properties*, that is, the arrangement of the points in the
image plane (represented in Equation 2 by ρ and α), and (b)
dynamic properties, that is, the velocities associated with
each point in the image plane. The dynamic properties are
characterized by the displacements of the points in the
image plane (configural change) and by the time interval in
which this transformation takes place. Distinguishing be-
tween these two dynamic properties is important because
this distinction has relevance for understanding existing
algorithms that have been proposed to solve the SfM prob-
lem. These algorithms, in fact, can be divided into two
classes: position-based algorithms (Hoffman & Flinch-
baugh, 1982; Huang & Lee, 1989; Longuet-Higgins, 1982;
Ullman, 1979) and velocity-based algorithms (Hoffman,
1982; Koenderink & Van Doorn, 1975, 1976, 1981;
Longuet-Higgins & Prazdny, 1980; Waxman & Wohn,
1987). From the point of view of the computational ap-
proach, both formulations are equivalent: In each case, in
fact, the correct solution is recovered (see Ullman, 1983).
From our point of view, conversely, it is necessary to take
into consideration both time and displacements, because *def*
takes on different values if either of these is manipulated. If
the notion of *def* has a psychological significance, then we
would expect that the manipulation of any of these aspects
of the 2-D transformations would have a perceptual effect
(this issue is addressed in Experiments 1 and 3).

Having characterized the notion of *def*, and having pro-
vided a method for its computation, we can put forward the
two hypotheses that motivated this article. The first hypoth-
esis is that *perceived angular velocity* (ω') *is a monotoni-*

cally increasing function of def (see also Domini, Caudek,
& Gerbino, 1995):

$$\omega' = f_\omega\,(def). \tag{3}$$

It is important to notice that if perceived angular velocities
are derived heuristically from the first-order temporal prop-
erties of the velocity field, then, in general, the perceptual
solution will not be veridical.

The second hypothesis applies to the case in which there
are more than three moving points and states that *the per-
ception of rigidity is influenced by the degree of variability
of the projected deformations*. To clarify this hypothesis, let
us consider a cloud of four points moving rigidly in 3-D
space. Given two views of the points, it is possible to
compute $\binom{4}{3} = 4$ deformations (i.e., the deformations of all
the possible triplets of points). These deformations have the
same magnitude only if the four points lie on a moving
plane. In this case, in fact, all triplets of points in 3-D space
share the same slant and move with the same angular
velocity. As a consequence of this, each triplet of points has
the same *def*, because *def* is equal to the product of σ and ω.
Because we hypothesize that the perceived angular veloci-
ties are a monotonically increasing function of *def*, we
expect this configuration to be perceived as moving rigidly:
Each triplet of points, in fact, should be associated with the

---

[3] The equation for the deformation (Equation 2) was derived
directly from the equation for the optic flow produced by the
projected motion of a planar surface undergoing a generic 3-D
rotation. The velocity components in the x and y directions are
linear functions of the x and y positions on the image plane. If the
positions and the velocities of three points are known, it is possible
to calculate the coefficients of the linear functions, that is, the
gradient of the optic flow. The deformation is the sum of the
squared two components of shear of the optic flow that are directly
derived by the decomposition of the gradient. When we use this
equation to calculate the deformation, the velocities are approxi-
mated by the displacements divided by the time intervals. The
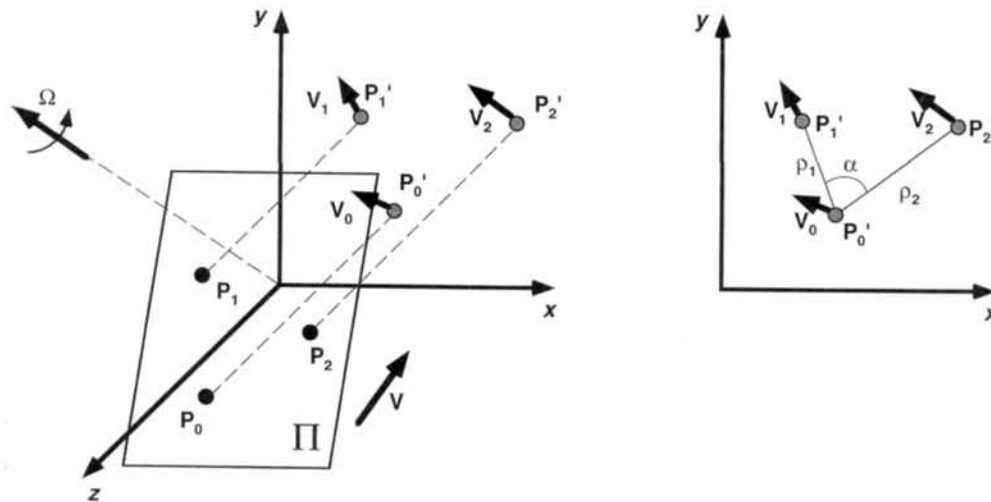smaller the displacements, the more precise the approximation.

Figure 3. Schematic representation of a planar patch Π defined by points P_0, P_1, and P_2 and undergoing generic motions in three-dimensional space (i.e., a rotatory motion Ω and a translatory motion V). P_0', P_1', and P_2' represent the orthogonal projection of P_0, P_1, and P_2 on the x–y image plane, and V_0, V_1, and V_2 represent their instantaneous two-dimensional velocities (left panel). ρ_1 and ρ_2 are the distances of the points P_1' and P_2' from P_0', and α is the angle between the line segments $P_0'P_1'$ and $P_0'P_2'$ (right panel).

same perceived angular velocity. In this case, the variability of the deformations is nil. Conversely, let us consider the case of four points lying within a volume and moving rigidly in 3-D space. In this case, each triplet of points in 3-D space has, in general, a different slant, and as a consequence, each triplet of points will be associated with a different *def* (the product of different slants times the same angular velocity). Because we hypothesize that perceived angular velocity depends on *def*, in this case we expect the configuration to be perceived as moving nonrigidly: In fact, each triplet of points should be associated with a different angular velocity. In conclusion, the greater the variability of the deformations, the greater the probability that the stimulus is perceived as nonrigid. The variance of the deformations of four adjacent points will be considered as the *local measure* of variability of the deformations. If more than four moving points are given, the *global measure* of the variability of the deformations will be the mean of the local variances of the deformations. According to this hypothesis, *stimuli having a small mean local variance of the deformations are considered rigid, whereas those having a large mean local variance of the deformations are considered nonrigid, regardless of whether they are the mathematically correct projection of a rigid structure.* Therefore, a smoothly curved surface rigidly moving in 3-D space is considered rigid, because the local variances of the deformations and their average are close to zero.

Overview of Experiments

In the present research we tested two hypotheses: (a) Magnitudes of angular velocity are misperceived in the kinetic depth effect because they are heuristically derived as a monotonically increasing function of *def* and (b) misperceptions of angular velocities lead to misperceptions of rigidity. Two sets of experiments were designed.

Perception of Angular Velocity in the Kinetic Depth Effect

Experiments 1–3 deal with the problem of perceiving angular velocity. We investigated how perceived angular velocities are influenced by *def* and by the magnitudes of simulated velocity, both in the minimal conditions of three points and two views and in displays having a larger number of points and views. We found that perceived angular velocities were influenced by *def* regardless of the simulated 3-D angular velocity and that perceptual performance did not significantly improve if second-order information was made available in the stimulus displays.

Perception of Rigidity in the Kinetic Depth Effect

Experiments 4–7 deal with the problem of perceiving rigidity. We tested the hypothesis that the perceptual discrimination of rigid from nonrigid motion is influenced by the variability of the deformations (i.e., by the variability of the perceived angular velocities) regardless of whether the stimuli simulate a mathematically correct projection of a rigid structure. We found that perceptual performance was veridical when the variability of the deformations favored this interpretation and nonveridical otherwise.

Perception of Angular Rotation in the Kinetic Depth Effect

Experiment 1

In the present experiment we had two purposes: (a) to manipulate the dynamic (2-D velocities) and configural properties (2-D distances and angular magnitudes among image features) that influence *def* and to measure their effects on the perception of 3-D angular rotation and (b) to investigate the perception of 3-D angular rotation in displays providing only first-order temporal information (2-view displays) and in displays in which second-order temporal information was potentially available (13-view displays).

In this experiment, we simulated the rigid rotation of three points in 3-D space. The points were arranged in 3-D space so that the deformation associated with their 2-D transformations was directly proportional to their projected velocities and inversely proportional to the distances between their projected positions. We hypothesized that the magnitude of perceived 3-D angular velocity is a monotonically increasing function of *def*. Because in this experiment all stimuli had the same temporal duration, a measure of perceived angular velocity is provided by the magnitude of perceived angular rotation. Therefore, we expected that the magnitude of perceived 3-D angular rotation would increase as the projected velocities increased and decrease as the 2-D interpoint distances increased. Moreover, because *def* does not depend on second-order temporal information, we expected that the manipulation of the number of frames in the stimulus displays would not affect perceptual performance.

Method

Participants. Thirty-two University of Trieste undergraduates participated in this experiment. All of them were naive to the purpose of the experiment.

Stimuli. A stimulus display consisted of 2 or 13 views of three high-luminance dots moving on a low-luminance background. The horizontal motions of the dots were computed to simulate an orthographic projection of points undergoing rotation in three dimensions and oscillating about the *y*-axis (see Figure 4). The angle of rotation was randomly chosen in each trial in the interval between 6° and 18°.

To simplify the configural properties of the displays, we positioned the triplets of points (P_0, P_1, and P_2) used in each trial in 3-D space so that they projected in the first view of the apparent motion sequence as the vertices of similar right triangles. The manipulation of the configural properties of the displays was limited to varying the distances among the points. If P_0 is the vertex of the right triangles and the 2-D trajectories of the points are parallel, then Equation 2 reduces to

$$def = \sqrt{\left(\frac{dx_1}{\rho_1}\right)^2 + \left(\frac{dx_2}{\rho_2}\right)^2},\qquad(4)$$

where dx_1 and dx_2 are the displacements of P_1 and P_2, respectively, relative to P_0. If the distances ρ_1 (between P_0 and P_1) and ρ_2 (between P_0 and P_2) are manipulated by means of a multiplicative factor m_s (i.e., $\rho_1 = m_s \rho_{01}$, and $\rho_2 = m_s \rho_{02}$, where ρ_{01} and
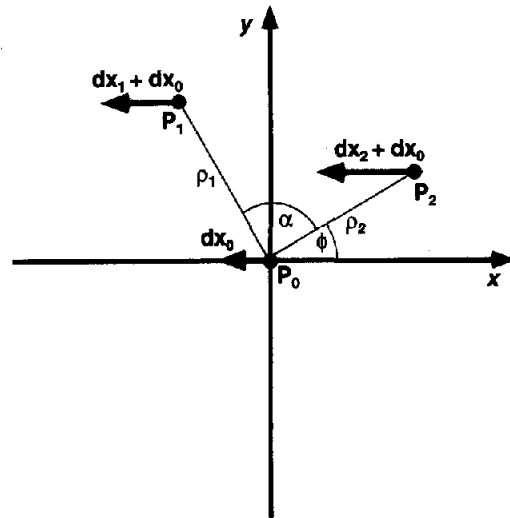

Figure 4. Schematic representation of the stimulus used in Experiment 1. Three points P_0, P_1, and P_2 were depicted on the projection plane with P_0 in the center of the system of coordinate axes. P_0 was one of the vertices of a right triangle having P_1 and P_2 as the other two vertices. The size of the catheti of the right triangle was equal to ρ_1 and ρ_2. dx_0 represents the two-dimensional (2-D) velocity of the point P_0, and dx_1 and dx_2 represent the 2-D velocities of the points P_1 and P_2 relative to P_0. ϕ is the angle between the cathetus P_0P_2 and the *x*-axis.

ρ_{02} are arbitrary), and if the projected relative displacements dx_1 and dx_2 of the points P_1 and P_2, respectively, are manipulated by means of a multiplicative factor m_v (i.e., $dx_1 = dx_2 = m_v dx_a$, where dx_a is arbitrary), then Equation 4 reduces to

$$def = \frac{m_v}{m_s}\sqrt{\frac{dx_a^2}{\rho_{01}^2} + \frac{dx_a^2}{\rho_{02}^2}} = k\frac{m_v}{m_s}.\qquad(5)$$

In this way, it is easy to see that *def* is directly proportional to the multiplicative factor m_v and inversely proportional to the multiplicative factor m_s.

As indicated in Figure 4, the point P_0 was presented in the center of the display with a projected displacement dx_0 equal to 0.27° of visual angle. The angle ϕ between the *x*-axis and the line passing through P_0 and P_2 was randomly chosen in each trial in the interval between 0° and 360°. The distances between P_0 and P_1 (ρ_1) and between P_0 and P_2 (ρ_2) were $\rho_1 = m_s \cdot 1.33°$ of visual angle and $\rho_2 = m_s \cdot 1.06°$ of visual angle, respectively. The 2-D displacements of the points P_1 and P_2 (dx_1 and dx_2) were $dx_1 = dx_2 = dx_0 + (.16 \cdot m_v)$. The multiplicative factors m_v and m_s could take on the values of 1, 2, and 4. In each stimulus, the *z* value of every point was chosen in such a way as to obtain the appropriate 2-D displacement for its specific angle of rotation and for the specific magnitudes of the factors m_v and m_s. The stimulus onset asynchrony (SOA) between views was 400 ms for the 2-view displays and 33 ms for the 13-view displays.

An icon representing two line segments forming an angle was shown in the upper part of the terminal screen. Movement of a mouse connected to an IRIS Workstation varied the represented angular magnitudes. The angular magnitude represented by the icon before the participants' adjustments was randomly selected on each trial.

Apparatus. The stimuli were presented on a high-resolution

color monitor (1,280 by 1,024 addressable locations) under the control of a Silicon Graphics IRIS Workstation. The screen had a refresh rate of 60 Hz and was approximately photometrically linearized. An anti-aliasing procedure was used: For point-light locations falling on a pixel boundary, the screen luminance was proportionally adjusted in the relevant addressable locations. The graphics buffer was 8 bits deep (256 gray levels). Participants viewed the displays through a reduction screen that reduced the field of view to a circular area with a diameter that subtended a visual angle of about 9.6°. The eye-to-screen distance was 1.1 m.

Design. Three variables were studied in this experiment: (a) interpoint distance (m_s), (b) projected relative velocity of the points within each oscillation cycle (m_v), and (c) number of frames. The first two were within-subjects variables, whereas the third was a between-subjects variable. Each participant viewed, in random order, one presentation of each of the nine combinations of the within-subjects variables. We presented nine additional trials at the beginning of each experimental session in order to familiarize the participants with the stimulus displays.

Procedure. All participants were tested individually in one session. They were instructed to judge the magnitude of angular rotation about the y-axis for each stimulus display. In each trial, participants provided their judgments by manipulating the icon present in the upper part of the terminal screen. Vision was monocular, head motion was not restricted, and eye movements were permitted. During the experiment, the experimental room was dark. No restriction was placed on viewing time. No feedback was given until after the experiment was completed.

Results and Discussion

Mean perceived rotation for each level of the two factors m_s and m_v is presented in Figure 5. A repeated measures ANOVA indicated a significant effect of the factor m_s (interpoint distance), $F(2, 56) = 22.53, p < .001$, and of the factor m_v (projected velocity), $F(2, 56) = 35.04, p < .001$. As expected, the magnitude of perceived rotation decreased as the multiplicative factor of interpoint distance increased (on average, by 36%) and increased as the multiplicative
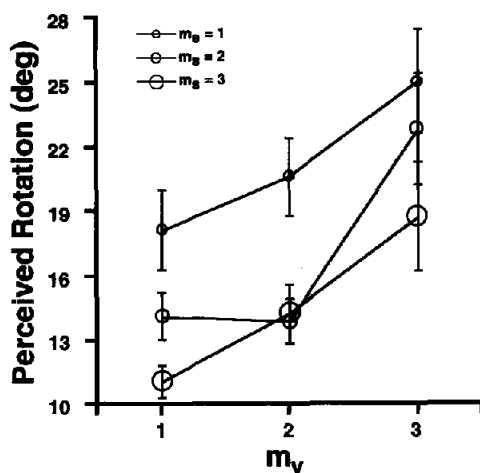

Figure 5. The average judged rotation in Experiment 1: m_s is the multiplier factor for the three levels of interpoint distances; m_v is the multiplier factor for the three levels of projected displacements. (Vertical bars represent 1 SE.) deg = degree.

factor of the projected velocity increased (on average, by 43%). The manipulation of number of frames did not significantly affect performance, $F(1, 28) = 2.35, ns$. None of the interactions reached significance.

We performed a second repeated measures ANOVA on the judgments of perceived rotation by using as independent variables the simulated *def* and the number of frames. We calculated *def* by considering the relative displacements dx_1 and dx_2 between the first and the last frames, respectively. The effect of *def* was significant, $F(2, 56) = 22.53, p < .001$: As expected, the magnitude of perceived rotation increased as *def* increased (see Figure 6). On the other hand, the manipulation of the number of frames did not significantly affect performance, $F(1, 30) = 2.41, ns$. The interaction between these two variables did not reach significance.

In conclusion, it is important to notice two aspects of the present results. Performance in the 13-frame condition was not significantly different from performance in the 2-frame condition. This finding supports our hypothesis that only the first-order properties of the stimulus displays are used, even if second-order temporal information is available (a finding also reported by Braunstein et al., 1990; Liter, Braunstein, & Hoffman, 1994; Todd & Bressan, 1990). Second, even though 2-frame sequence stimulus displays are compatible with a one-parameter family of solutions (there are infinite couples of different slants and angular rotations that could have produced the simulated 2-D transformations), participants' judgments revealed a coherent pattern of responses. This finding is consistent with the predictions of the model described in the introduction: The magnitude of perceived angular velocity depends on a specific monotonically increasing function of *def*.

Experiment 2

In Experiment 1 we tested the hypothesis that perceived angular velocity is a monotonically increasing function of *def* under the minimal conditions for computing *def* from an SfM display (i.e., two views of three points). The displays, therefore, were very impoverished. In Experiment 2 we tested the psychological relevance of *def* in less impoverished conditions, that is, by considering the orthogonal projection of planar surfaces made up of random dots rather than the projection of just three points.

Method

Participants. Ten University of Trieste undergraduates participated in this experiment. All of them were naive to the purpose of the experiment.

Stimuli. The stimuli were high-luminance dots moving on a low-luminance background. The horizontal motions of the dots were computed to simulate an orthographic projection of a planar surface undergoing rotation in three dimensions and oscillating about the y-axis. Each stimulus display was contained within a circular "window" with a diameter of 5.3° of visual angle so that changes in the projected contours of the simulated surfaces would not be visible. Changes in projected contours, in fact, are a poten-
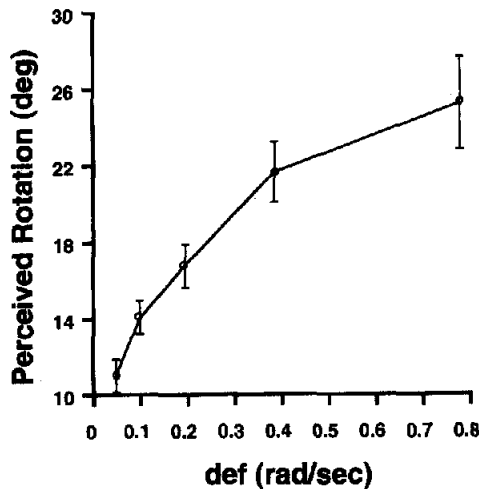
Figure 6. The average judged rotation in Experiment 1 as a function of *def* (deformation). (Vertical bars represent 1 SE.) deg = degree.

tially useful source of information about 3-D object shape (e.g., Cortese & Andersen, 1991; Pollick, 1989). The dots were randomly distributed with uniform probability density over the projection plane (not evenly distributed over the simulated surfaces). We manipulated dot lifetime so as to keep the number of dots constant in each frame of the stimulus display. In every stimulus display there were 116 dots. We manipulated two variables: the magnitude of angular velocity and *def.* We manipulated angular velocity by keeping the duration of each oscillation cycle constant for each stimulus display and by varying the magnitude of angular rotation. The simulated angles of rotation could take on the values of 8°, 14°, 20°, 26°, and 32°. For each stimulus display, the slant of the simulated planar surfaces was calculated in such a way as to produce the appropriate deformations for each simulated angular rotation. Simulated slant varied between 0.219 and 5.625 (i.e., tan 12.3° and tan 79.9°). Because *def* changes during rotation, a summary measure of this quantity can be provided by the sum of the instantaneous *def* calculated for each successive couple of frames *(integral-def)*. In this experiment, *integral-def* could take on the values of 0.30, 0.53, 0.75, 0.98, and 1.20 rad/s. We added a translatory motion component to every dot to equate the maximum 2-D velocity in each stimulus display (3.16 deg/s). One oscillation cycle took 60 frames or 3 s. As in Experiment 1, an icon representing two line segments forming an angle was shown in the upper part of the terminal screen.

Apparatus. The apparatus was the same as in Experiment 1.

Design. Two variables were studied in this experiment: (a) *def* and (b) magnitude of angular rotation within each oscillation cycle. Both were within-subjects variables. Each participant viewed, in random order, one presentation of each of the 25 different stimuli. Eight additional trials were presented at the beginning of each experimental session in order to familiarize the participants with the stimulus displays.

Procedure. The experimenter was careful to explain to participants that a translatory motion component was added to the horizontal projected motions in order to equate the maximum projected velocities for each stimulus display. The participants were instructed to judge the magnitude of angular rotation about the *y*-axis. Otherwise, the procedure and instructions were the same as in Experiment 1.

Results and Discussion

Mean perceived rotations for each simulated rotation and for each *def* are presented in Figure 7. As can be seen in the figure, the magnitudes of perceived rotation were not influenced by the magnitudes of simulated rotation, $F(4, 36) = 0.97$, *ns*, but increased as *def* increased, $F(4, 36) = 36.69$, $p < .001$. The interaction between these two variables was not significant.

In conclusion, the magnitudes of perceived rotation were a monotonically increasing function of *def* and were not affected by the magnitudes of simulated rotation. These findings, therefore, replicated those of the first experiment.

Experiment 3

In the introduction we pointed out that *def* is influenced by two aspects of the 2-D transformations: (a) configural properties (i.e., the arrangement of the points in the image plane) and (b) dynamic properties (i.e., the velocities associated with each point in the image plane). The dynamic properties are characterized by both the displacements of the points in the image plane (configural change) and by the time interval (t) in which this transformation takes place. Because in the previous experiments all stimulus displays were presented for the same temporal duration, those experiments do not clarify whether the perceptual responses were influenced by the *displacement fields* (configural change independent from time) or by the *velocity fields* (configural change in the interval of time).

We designed the present experiment to distinguish the perceptual effects of the displacement fields from those of the velocity fields. We hypothesized that the perceived angular velocities are a monotonically increasing function of *def*: $\omega' = f(def)$. If we assume that $\alpha' = t \cdot f(def)$, then we should also expect that the magnitude of perceived angular rotation (α') is influenced by both of the aspects
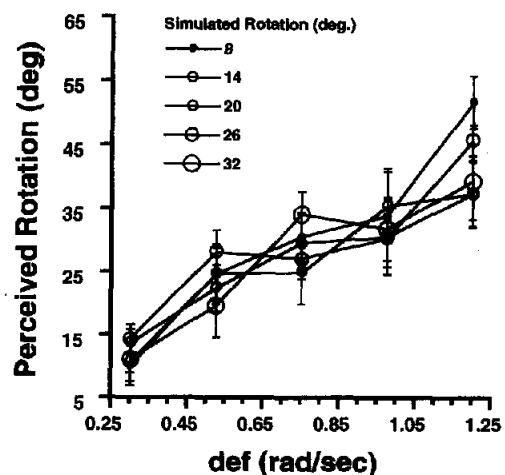

Figure 7. The average judged rotation in Experiment 2 as a function of *def* (deformation) for each simulated rotation. (Vertical bars represent 1 SE.) deg = degree.

that influence *def*: (a) the magnitude of the displacement fields and (b) their temporal duration.

Method

Participants. Sixteen University of Trieste undergraduates participated in this experiment. All of them were naive to the purpose of the experiment.

Stimuli. The stimuli were high-luminance dots moving on a low-luminance background. The horizontal motions of the dots were computed as simulating an orthographic projection of a planar surface undergoing rotation in three dimensions and oscillating about the *y*-axis. The stimulus displays simulated two displacement fields corresponding to angular rotations (α) of 12° and 18° (see Figure 8). Before rotation, the simulated surfaces were oriented so that their slant was equal to tan ($70° - \alpha/2$) and their tilt was equal to 0°. In Figure 8 are presented the top views of the planar surfaces before and after each of the two simulated rotations. In this particular case (rotation about the *y*-axis of a planar surface parallel to the *y*-axis), the displacement fields are a linear function of *x*. The magnitudes of the displacements as a function of *x* are represented in the top graphs of the figure. The steepness of these linear functions represents the gradients of the displacement fields. In the case of this specific stimulus display, the gradient of the velocity field is equal to *def* and can be computed from the ratio between the gradient of the displacement fields and

Figure 8. Schematic representation of the stimulus displays used in Experiment 3. The figure depicts the top views (*x–z*) of the simulated planar surfaces before (black circles) and after (white circles) each of the two simulated rotations (12° and 18°). The orientations of the simulated surfaces before rotation are equal to 64° and 61°, and their tilt is equal to 0°. d_{max} represents the magnitude of the maximum displacement. In the top graphs the magnitudes of the displacements $d(x)$ are represented as a function of *x*.

the temporal duration of the stimulus displays. In this experiment, we manipulated the temporal duration of the stimulus displays by varying the number of frames in each apparent motion sequence and by keeping the SOA constant. By manipulating the temporal duration of the stimulus displays we kept the magnitudes of *def* constant for the two displacement fields. Three *integral-defs* (gradients of the velocity field) were simulated for each displacement field: 0.45, 0.78, and 2.09 rad/s. For the 12° displacement fields, 6, 15, and 24 frames were used; for the 18° displacement fields, 9, 22, and 36 frames were used. The relations between the temporal duration of the stimulus displays, the number of frames, the displacement fields, and *def* are reported in Table 1.

Each stimulus display was contained within a rectangular "window" of constant size, the height and length of which measured, respectively, 1.34° and 2.67° of visual angle. The dots were randomly distributed with uniform probability density over the projection plane. Each frame of the stimulus display contained approximately 50 dots. Frame rate was equal to 15 Hz. As in the previous experiments, an icon representing two line segments forming an angle was shown in the upper part of the terminal screen.

Apparatus. The apparatus was the same as in Experiment 1.

Design. Two variables were manipulated in this experiment: (a) the gradient of the displacement field and (b) the gradient of the velocity field (*def*). The manipulation of both variables was within subjects. Each participant viewed three presentations, in random order, of the six combinations of the within-subjects variables. Six additional trials were presented at the beginning of each experimental session in order to familiarize the participants with the stimulus displays.

Procedure. The participants were instructed to judge the magnitude of angular rotation about the *y*-axis. Otherwise, procedure and instructions were the same as in the previous experiments.

Results and Discussion

The present experiment was based on the assumption that the magnitudes of perceived angular rotation (α') are equal to the product of the temporal duration of the stimulus display (*t*) and a monotonically increasing function of *def*: $\alpha' = t \cdot f(def)$. We performed a first analysis, therefore, to provide evidence supporting this assumption. In this analysis we studied the effects of the temporal duration (*t*) on the perceptual response for stimulus displays with the same *def*. We expected that if the stimulus displays had the same *def* but different temporal durations, then the magnitudes of perceived angular rotation should be directly proportional to the temporal duration of the stimulus displays. Because the temporal duration of the 18° displacement field was 50% larger than the temporal duration of the 12° displacement field, we expected the magnitude of perceived rotation to increase correspondingly by 50%. To test this hypothesis, we computed the ratio between the magnitudes of perceived angular rotation for stimuli with the same *def* but different temporal durations. If our hypothesis is correct, then we would expect that

$$\frac{\alpha'(def_i, 1.5t_i)}{\alpha'(def_i, t_i)} = \frac{1.5t_i f(def_i)}{t_i f(def_i)} = 1.5. \qquad (6)$$

A repeated measures ANOVA was performed with the three *defs* as independent variables and the ratio between the
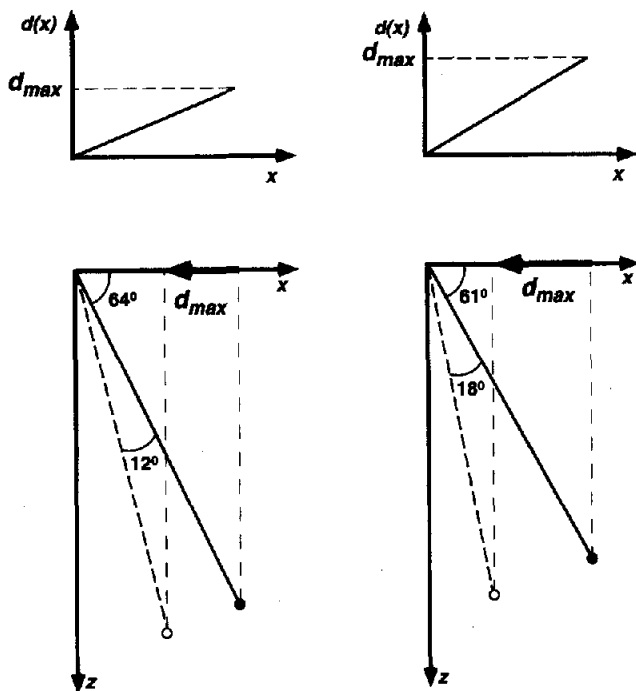
Table 1
Temporal Duration (in Milliseconds) of Stimulus Displays as a Function of Angular Rotation and Deformation (in Radians/Second)

	Displacement field				
	12°			18°	
Deformation	Duration	No. of frames		Duration	No. of frames
0.45	400	6	1.5 ·	400	9
0.78	1,000	15	1.5 ·	1,000	22
2.09	1,600	24	1.5 ·	1,600	36

magnitudes of the perceived angular rotation for stimuli with the same *def* but different time durations as the dependent variable. This analysis indicated that the dependent variable was not significantly influenced by *def*, $F(2, 30)$ = 0.40, *ns*. No other effects or interactions were significant. Post hoc comparisons indicated that, for all three simulated deformations, the ratios between the magnitudes of perceived rotation for stimuli with the same *def* but different time durations were not significantly different from the expected value of 1.5: $t_0(15)$ = -0.01, *ns*; $t_1(15)$ = 1.19, *ns*; $t_2(15)$ = 1.26, *ns*. The results of this analysis are consistent with the assumption that the magnitudes of perceived angular rotation are modulated by the temporal duration of the stimulus displays, as indicated by Equation 6, and therefore allow us to conclude that (within the present experimental parameters) it is reasonable to estimate the magnitude of perceived angular velocity by computing the ratio between the magnitude of perceived angular rotation and the temporal duration of the stimulus displays. Having tested the assumption at the basis of the present experiment, we analyzed the experimental data according to the variables that had been manipulated in the experimental design. Mean perceived rotations for each gradient of the displacement field and for each gradient of the velocity field are presented in Figure 9. A repeated measures ANOVA performed on gradient of the displacement field and gradient of the velocity field indicated that both the gradient of the displacement field, $F(1, 15)$ = 183.23, $p < .001$, and the gradient of the velocity field, $F(2, 30)$ = 57.66, $p < .001$, significantly affected the magnitudes of perceived rotation. Also significant was the interaction between these two variables, $F(2, 30)$ = 6.97, $p < .01$. No other effects or interactions reached significance.

The important thing to notice in these results is that the magnitudes of the perceived angular rotation did not depend only on the magnitudes of the projected displacements (the perceived angular magnitudes increased as the gradient of the displacement field increased) but were also influenced by the time interval within which the 2-D transformations took place (the perceived angular magnitudes decreased as the gradient of the velocity field increased). The significant interaction was due to the fact that the gradient of the displacement field had a smaller effect for larger values of the gradient of the velocity field. These findings, therefore, are consistent with the hypothesis that motivated the present experiment.

One can also analyze the present data, however, by considering the relation between the magnitudes of perceived angular velocity and the gradient of the velocity field (*def*). An estimation of the perceived angular velocities can be provided by the ratio between the magnitudes of perceived angular rotation and the temporal durations of the stimulus displays. The estimated mean angular velocities for each experimental condition are plotted in Figure 10. A repeated measures ANOVA with gradient of the displacement field and gradient of the velocity field as independent variables and estimated perceived angular velocity as the dependent variable indicated a significant effect of the gradient of the velocity field, $F(2, 30)$ = 34.38, $p < .001$, whereas the effect of the gradient of the displacement field was not statistically significant. No other effects or interactions reached significance.

These results indicate that perceived angular velocities are a monotonically increasing function of the gradient of the velocity field and that they are not affected by the gradient of the displacement field. This result is consistent with the model described in the introduction because *def* is the gradient of the velocity field. An implication of the present experiment is that an adequate characterization of
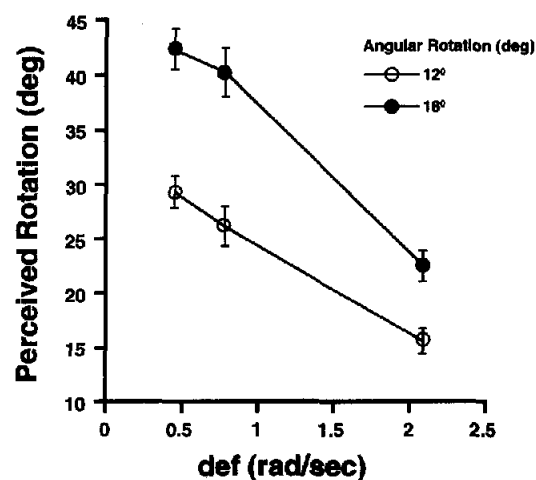

Figure 9. The average judged rotation in Experiment 3 as a function of the gradient of the velocity field (*def*, or deformation) for the two displacement fields corresponding to the simulated rotations of 12° and 18°. (Vertical bars represent 1 SE.)

the input for the perceptual derivation of structure from motion cannot be limited to the description of the gradient of the displacement fields but must also take into account the temporal duration within which the 2-D transformations occur.

Perceived Rigidity in the Kinetic Depth Effect

In the previous three experiments we investigated the perception of angular rotation in the kinetic depth effect. We found that the magnitudes of perceived angular rotation are systematically misperceived and that these misperceptions are consistent with a heuristic analysis based on the first-order properties of the optic flow. We designed the following four experiments to test the hypothesis that the misperceptions of angular rotation lead to misperceptions of rigidity in the kinetic depth effect.

Experiment 4

In the present experiment we had two objectives: (a) to replicate the results obtained by Braunstein et al. (1990) in order to ensure a correct implementation of their stimulus displays and (b) to examine the effect of the variance of the deformation on the discrimination between rigid and non-rigid 3-D motion of a configuration of four dots when the variance of the trajectories has been controlled. We hypothesized that the perception of rigidity depends on the variance of the deformations: The larger the variance of def, the greater the probability that the stimulus will be perceived as nonrigid, whether it simulates a rigid transformation or not.

Method

Participants. Eight University of Trieste undergraduates participated in this experiment. All of them were naive to the purpose of the experiment.
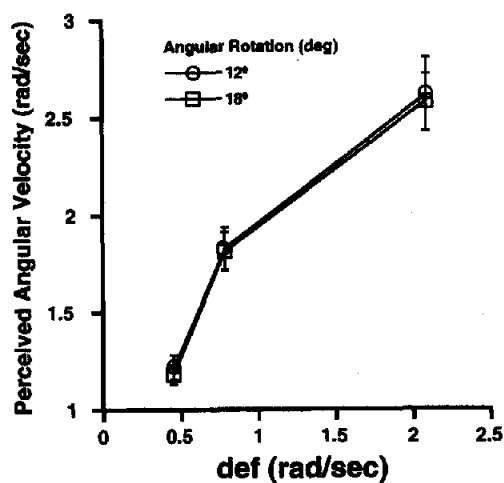

Figure 10. The estimated averaged angular velocities in Experiment 3 as a function of the gradient of the velocity field (def, or deformation) for the two displacement fields corresponding to the simulated rotations of 12° and 18°. (Vertical bars represent 1 SE.)

Stimuli. The stimuli consisted of two views of four high-luminance dots moving on a low-luminance background. The motions of the dots were computed as simulating an orthographic projection of points undergoing rotation in three dimensions. Initial point positions were randomly selected within the volume of a sphere having a diameter of 5.3° of visual angle. There were 34 possible axes of rotations. We generated the axes by using the same procedure as Braunstein et al. (1990): 34 points evenly spaced were placed on the surface of a sphere, and the axes of rotation were defined by connecting each of these points to the center of the sphere. The slant of the axes varied between 45° and 90°, and their tilt varied between 0° and 180°. There were two experimental conditions. In the *replication* condition, we simulated the stimulus displays by using the parameters described by Braunstein et al. (1990) in their Experiment 1. For rigid stimuli, all points rotated about the same axis (randomly selected from the set of potential axes previously defined). For nonrigid stimuli, each point rotated about a different axis randomly selected from the same set. For both rigid and nonrigid stimuli, the angle of rotation was selected at random in each trial from the interval between 6° and 18° and was the same for all points in the display.

In the *test* condition, we created the rigid stimuli in the same way as in the replication condition, whereas we created the nonrigid stimuli by having a single axis of rotation for all points of each display (randomly selected from the set of potential axes previously defined) but a different magnitude of rotation for each point (the angle of rotation for each point was selected at random from the interval between 6° and 18°). We calculated the deformations of the four possible triplets of points for each stimulus display and computed their variance. Furthermore, in the test condition we selected the rigid and nonrigid stimuli so as to keep the variance of the deformations for the rigid stimuli at least three times larger than the variance for the nonrigid stimuli. In the replication condition, the variance of the deformations was not controlled: An ANOVA performed on the stimulus displays indicated that rigid and nonrigid stimuli did not significantly differ with reference to this variable.

Three constraints were imposed on the stimulus displays: (a) nearest neighbor correspondence, (b) minimum 2-D motion, and (c) minimum 3-D spacing. The constraints were the same as those of Braunstein et al. (1990, p. 207). We computed a measure of 2-D nonrigidity for both rigid and nonrigid displays to determine whether the 2-D projections of nonrigid displays were less rigid than the 2-D projections of the rigid displays; a measure of 3-D nonrigidity for both test and control conditions was computed as well (for details see Braunstein et al., 1990, p. 207). We also computed the variance of the projected velocities, the minimum–maximum velocity ratio, the minimum–maximum velocity difference, and the variance of trajectories of individual points. We conducted six ANOVAs on the stimulus displays by using these measures as the dependent variables. The independent variables were condition (replication vs. test) and rigidity (3-D rigid vs. 3-D nonrigid displays). (In the case of the 3-D nonrigidity measure, the analysis was restricted to the nonrigid displays, with condition as the only independent variable.) The 2-D nonrigidity did not differ significantly for the rigid and nonrigid displays, but there was a significant difference for the replication and test displays: 2-D nonrigidity was 26.6% smaller in the displays of the test condition, $F(1, 316) = 8.26$, $p < .05$. The 3-D nonrigidity did not differ significantly in the replication and test conditions. The results of the other analyses are discussed below with reference to the experimental results.

Apparatus. The apparatus was the same as in Experiment 1.

Design. Two variables were studied in this experiment: (a)

condition (replication vs. test) and (b) rigidity (3-D rigid vs. 3-D nonrigid displays). Condition was a between-subjects variable, and rigidity was a within-subjects variable.

Procedure. Each participant took part in a practice session followed by two experimental sessions. Each session consisted of a random sequence of 80 trials containing 40 "rigid" trials and 40 "nonrigid" trials. The participants were instructed to discriminate between rigid and nonrigid motion. They were told that points move rigidly if all their 3-D interpoint distances remain constant over time. They were also told that within each experimental session, half of the stimuli simulated a rigid transformation, and the other half simulated a nonrigid transformation. In each trial participants indicated their choice with a keypress. In the replication condition, they were given feedback after each incorrect response.[4] No feedback was used for the test condition. Otherwise, the procedure and instructions were the same as in Experiment 1.

Results and Discussion

Like Braunstein et al. (1990), we analyzed the experimental data by means of a signal detection paradigm in which the rigid trials were considered as signal trials. A d' measure was computed for each participant and stimulus condition. As in Braunstein et al., d' was based on 160 trials, half of which were signal trials. We calculated the significance of d' by using Marascuilo's (1970) one-signal significance test. The results of this analysis are reported in Table 2.

All eight d's were significantly different from zero. The results for the replication condition fully replicated those obtained by Braunstein et al. (1990). Average d' for their participants was 0.64, and for our participants it was 0.731. In this experimental condition, participants veridically discriminated rigid from nonrigid motion at the minimum level at which this discrimination is theoretically possible: two views and four points. However, when we controlled for the variance of the trajectories of the individual points (by using a different procedure for creating the nonrigid stimuli) and manipulated the variance of the deformations, performance changed dramatically: In the test condition, observers judged as "nonrigid" the displays with the largest variance of the deformations (which simulated the projection of rigid transformations), and they judged as "rigid" the displays with the smallest variance of the deformations (which were

Table 2
d' Scores in Experiment 4

Participant	d'
Replication condition	
Al.	0.539*
Lu.	1.063*
Ri.	0.650*
Ra.	0.671*
Test condition	
An.	−1.306*
El.	−1.649*
Ka.	−1.659*
Ma.	−0.474*

* $p < .05$.

incompatible with the projection of rigid transformations). This result supports the hypothesis that motivated the present experiment: The discrimination between rigid and nonrigid motion is not necessarily veridical but depends on the variance of deformations of the triplets of points of the stimulus displays.

Let us consider alternative interpretations that could possibly account for the results of the present experiment in terms of 2-D variables other than *def.* In our stimuli, the variance of velocities was significantly larger (80%) for nonrigid displays, $F(1, 316) = 45.32, p < .05$, and significantly larger (84%) in the test condition, $F(1, 316) = 48.03, p < .05$. The significant interaction, $F(1, 316) = 60.97, p < .05$, between these two variables is due to the fact that the difference between rigid and nonrigid stimuli was larger in the test condition than in the replication condition. The variance of the velocities, therefore, does not provide an alternative explanation for the results obtained in the test condition, because it is implausible that participants considered as "rigid" the displays having the largest variance of 2-D velocities.

The minimum–maximum velocity ratio was significantly larger (23%) in the test condition, $F(1, 316) = 15.05, p < .05$, but did not differ significantly for rigid and nonrigid displays. The minimum–maximum velocity ratio, therefore, cannot explain the results obtained in the test condition. The minimum–maximum velocity difference was significantly larger (38%) in the nonrigid displays and significantly larger (40%) in the test condition, $F(1, 316) = 46.89, p < .05$. The interaction between these two variables was due to the fact that the minimum–maximum velocity difference was 83% larger for the nonrigid displays in the test condition and 5% larger for the rigid displays in the replication condition. The minimum–maximum velocity difference, therefore, cannot explain the results obtained in the present experiment because it is not plausible that in the test condition, participants judged as "rigid" the displays having the larger velocity differences.

The variance of trajectories of individual points was significantly larger in the nonrigid displays (28%), $F(1, 316) = 9.43, p < .01$, and was also larger in the replication condition (43%), $F(1, 316) = 5.42, p < .05$. An analysis performed only on the stimulus displays of the test condition, however, indicated that rigid and nonrigid stimuli did not differ significantly. The variance of trajectories, therefore, could explain the results obtained in the replication condition (the displays having the largest variance of the trajectories, in fact, were judged as "nonrigid") but not those obtained in the test condition.

In conclusion, the results of the present experiment are consistent with the hypothesis that in the replication condition, participants judged as "rigid" the displays having the smallest variance of trajectories of the individual points,

---

[4] Feedback was used in this condition because Braunstein et al. (1990) reported that only 1 out of 3 participants in this experimental condition was able to successfully discriminate rigid from nonrigid motion without feedback whereas all of their participants were able to perform this task if given feedback.

whereas in the test condition, participants judged as "rigid" the displays having the smallest variance of the deformations (in the replication condition, rigid and nonrigid displays did not significantly differ in terms of the variance of deformations). None of the other 2-D variables that had been considered (i.e., the 2-D nonrigidity, the variance of the projected velocities, the minimum–maximum velocity ratio, the minimum–maximum velocity difference, and the variance of trajectories of individual points) is sufficient to explain the results obtained in the test condition.

Experiment 5

In Experiment 4 we found that the variance of the deformations of the triplets of points of a stimulus display influenced perceived rigidity: Stimuli with a lower variance of the deformations were judged as "rigid" more often than were those with a higher variance of the deformations, whether the stimulus displays simulated a rigid transformation or not. The variability of the deformations takes on its minimum value when all the triplets of points share the same *def*. This condition occurs when all points of a stimulus display are simulated as lying on a moving planar surface: In this case the variance of the deformations of each group of four adjacent points is equal to zero. We could hypothesize, therefore, that the optic flow produced by the 3-D motion of points lying on a planar surface should give rise to the perception of a rigid transformation more so than should the optic flow produced by the motion of a cloud of points in 3-D space. The continuum between the spatial arrangement of points lying on a planar surface (at one extreme) or within a volume (at the other extreme) can be quantified in terms of the R^2 of the regression plane that comes as close to the points as possible. In this experiment we tested the hypothesis that observers are more likely to judge as "rigid" the stimuli with high R^2 (as previously defined) than they are those with low R^2. With this manipulation our intent was to extend the findings of the previous experiment by using a larger number of views and by manipulating the number of points.

Method

Participants. Five University of Trieste undergraduates participated in this experiment. All of them were naive to the purpose of the experiment.

Stimuli. A stimulus display consisted of 4, 7, or 10 high-luminance dots moving on a low-luminance background. The motions of the dots were computed as simulating an orthographic projection of points undergoing rotation in three dimensions. The initial point positions were selected randomly within the volume of a sphere with a diameter measuring 5.32° of visual angle. The slant (relative to the viewing direction) and the tilt of the axis of rotation could take on any value randomly selected, respectively, from the intervals between 70° and 90° and between 0° and 180°.

We calculated the R^2 of the regression plane that came as close as possible to the points of each stimulus display by considering the z values as the dependent variable and the x and y values as the independent variables. We generated the stimuli by manipulating the R^2 of the best-fitting plane for the simulated dots in the

stimulus displays. There were two experimental conditions. In the high-R^2 condition, the R^2 was forced to be at least .9, while in the low-R^2 condition, R^2 was smaller than .2 (for nonrigid displays, these values refer only to the initial moment of the oscillation sequence because the dots were simulated as rotating by different amounts in 3-D space). In the case of the nonrigid displays in the high-R^2 condition, at the beginning of each motion sequence all points were simulated as lying on the surface of a plane. In both conditions, the motions of the dots were simulated so that the tilt of the best-fitting plane was parallel to the tilt of the axis of rotation, whereas the slant of the plane could take on any value between 20° and 40°. For the rigid displays, the angle of rotation was randomly selected from the interval between 6° and 26°. For the nonrigid displays, the angle of rotation of each point in the display was randomly selected from the interval between 6° and 26°. In both rigid and nonrigid displays all points rotated about the same axis of rotation. One oscillation cycle took 100 frames or 1.4 s.

The same constraints as in Experiment 4 were imposed on the stimulus displays. We equated the 2-D nonrigidity for all stimulus displays by adding or subtracting a common component of rotation to all the points of a display. As in the previous experiment, for each display we computed the variance of deformations, the variance of 2-D velocities, the minimum–maximum velocity ratio, the minimum–maximum velocity difference, the variance of trajectories of the individual points, and the variance of the deformations. We performed six ANOVAs on the stimulus displays by using these measures as the dependent variables. The independent variables were condition (low R^2 vs. high R^2), rigidity (3-D rigid vs. 3-D nonrigid displays), and number of points (4, 7, and 10). We computed the global variance of the deformations by taking the mean of the local variance of the deformations for every local quadruple of points. For N points we considered N local quadruples formed by the ith point of the configuration and the three points closest to the ith point. The mean of the local variance of the deformations was more than five times larger in the low-R^2 condition, $F(1, 228) = 67.15, p < .001$. For this dependent variable there were no other significant effects or interactions.

Apparatus. The apparatus was the same as in Experiment 1.

Design. Three variables were studied in this experiment: (a) condition (low R^2 vs. high R^2), (b) rigidity (3-D rigid vs. 3-D nonrigid displays), and (c) number of points (4, 7, and 10). All were within-subjects variables.

Procedure. Each participant took part in three sessions. Each session began with 10 practice trials followed by a random sequence of 80 trials consisting of 20 "rigid" trials and 20 "nonrigid" trials for each experimental condition (low R^2 vs. high R^2). The number of points of the displays was kept constant within each experimental session. The order of presentation of the sessions, which presented stimuli made of 4, 7, and 10 points, was randomized for each participant. The participants were instructed to discriminate between rigid and nonrigid motion. Otherwise, the procedure and instructions were the same as in Experiment 4.

Results and Discussion

The results are reported in Figure 11. A repeated measures ANOVA indicated that neither the number of points nor the simulated rigidity significantly affected performance. The only significant effect was due to the condition variable, $F(1, 12) = 29, p < .001$. As we expected, the stimuli in the high-R^2 condition were judged as "rigid" more often than were those in the low-R^2 condition (74% vs.
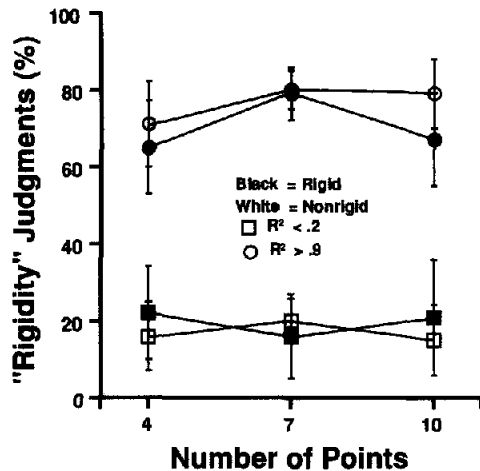
Figure 11. Percentages of "rigid" judgments in Experiment 5 for rigid (black) and nonrigid (white) displays and for the low-R^2 (squares) and high-R^2 (circles) conditions.

18%, respectively), regardless of whether they were the mathematically correct projection of a rigid structure.

As in the previous experiment, we checked for alternative interpretations of the experimental results by considering the other 2-D variables that we computed for the stimulus displays. The variance of 2-D velocities was significantly larger (222%) for the stimuli in the high-R^2 condition, $F(1, 228) = 207.78, p < .01$, and increased significantly in the displays with larger numbers of points. This variable, therefore, cannot account for the results of the present experiment because it is not plausible that participants judged as "rigid" the displays having the largest variance of velocities and because performance was not affected by the manipulation of the number of points. Both the minimum–maximum velocity ratio and the minimum–maximum velocity difference were significantly influenced by the variables of condition and number of points: These variables were larger in the high-R^2 condition, $F(1, 228) = 31.55, p < .01$, and $F(1, 228) = 257.55, p < .01$, respectively, and increased with the number of points, $F(1, 228) = 132.77, p < .01$, and $F(1, 228) = 63.21, p < .01$, respectively. The interaction between number of points and rigidity also was significant for the variable of minimum–maximum velocity ratio, $F(1, 228) = 6.13, p < .01$. These properties of the stimulus display, however, cannot explain the present results because the judgments of the participants were not influenced by either rigidity and number of points or by their interaction. Finally, the variance of the trajectories of individual points was 30% larger in the rigid displays, $F(1, 228) = 4.53, p < .01$. Again, this variable cannot explain the present results because we did not find any effect that was due to the variable of rigidity.

In conclusion, we found that the discrimination between rigid and nonrigid motion was significantly affected by the manipulation of the global variance of the deformations (which we manipulated in this experiment by varying the 3-D simulated spatial arrangement of the dots in the stim-

ulus displays), regardless of whether the stimuli were the mathematically correct projection of a rigid structure. Moreover, we found that the manipulation of number of points did not significantly affect performance. These results, therefore, replicate the findings of the previous experiment and extend their generality to displays with larger numbers of points and views.

Experiment 6

Our objective in this experiment was to extend the previous findings by considering stimulus displays that provide a compelling impression of two moving planar surfaces. Experiments 1–3 showed that the magnitudes of perceived rotation depend on *def* rather than on the magnitudes of simulated rotation. These findings allow us to hypothesize that if two planar surfaces generate the same deformations, then they should also be perceived to rotate by the same amount. In turn, this should produce the perception of a global rigid motion (even for 2-D transformations not compatible with the projection of a 3-D rigid motion). Conversely, if two planar surfaces generate different deformations, then they should be perceived to rotate by different amounts. By the same token, this should produce the perception of a global nonrigid motion (even for 2-D transformations compatible with the projection of a 3-D rigid motion).

Method

Participants. Ten University of Trieste undergraduates participated in this experiment. All of them were naive to the purpose of the experiment.

Stimuli. The stimuli were high-luminance dots moving on a low-luminance background. The motions of the dots were computed to simulate an orthographic projection of two intersecting transparent planar surfaces undergoing rotation in three dimensions and oscillating about a fixed axis. There were 34 possible axes of rotation. We generated the axes by using the same procedure as Braunstein et al. (1990), and we replicated those used in Experiment 4. In each trial, both planes were simulated as oscillating about the same axis. In different trials, a different axis was randomly chosen from the previously defined set of potential axes of rotation. For rigid displays, both planar surfaces were simulated as rotating through the same angle, randomly chosen in each trial from the interval between 8° and 18°. For nonrigid displays, one planar surface was simulated as rotating through an angle randomly chosen in each trial from the interval between 8° and 12°, whereas the other planar surface was simulated as rotating through an angle randomly chosen in each trial from the interval between 14° and 18°. The tilt of each planar surface was simulated at random in each trial in the interval between 0° and 180° with the constraint that the tilts of the two simulated surfaces differ by at least 30°. Once the tilt and the angle of rotation were chosen for a stimulus display, we used an iterative procedure to determine the slant necessary to produce the *def* appropriate for each experimental condition. There were four kinds of stimuli: (a) *same-def rigid*, in which a single *def* was randomly chosen from the interval between 1 and 1.2 rad/s for the two surfaces and the stimuli simulated a global rigid transformation; (b) *same-def nonrigid*, in which a single *def* was randomly chosen from the interval between

1 and 1.2 rad/s for both surfaces and the stimuli simulated a global nonrigid transformation; (c) *different-def rigid*, in which the *def* of the two surfaces took on the values of 0.22 and 1.11 rad/s and the stimuli simulated a global rigid transformation; and (d) *different-def nonrigid*, in which the *def* of the two surfaces took on the values of 0.22 and 1.11 rad/s and the stimuli simulated a global nonrigid transformation.

Each stimulus display was contained within a circular "window" with a diameter of 6.65° of visual angle. The dots were randomly distributed with uniform probability density over the projection plane and were randomly assigned with the same probability to the two simulated planar surfaces. Each frame of the stimulus display contained 160 dots. Dot density was kept constant. One oscillation cycle took 120 frames or 3.6 s.

Apparatus. The apparatus was the same as that in Experiment 1.

Design. Two variables were studied in this experiment: (a) rigidity (3-D rigid vs. 3-D nonrigid displays) and (b) condition (same *def* vs. different *def*). Both were within-subjects variables. Each participant viewed 20 presentations in random order of the four combinations of the within-subjects variables. Sixteen additional trials were presented at the beginning of each experimental session in order to familiarize the participants with the stimulus displays.

Procedure. The participants were instructed to discriminate between rigid and nonrigid motion. Otherwise, the procedure and instructions were the same as in Experiment 4.

Results and Discussion

Percentages of "rigid" judgments for the *same-def* and *different-def* conditions and for rigid and nonrigid stimuli are presented in Figure 12. The first thing to notice in these results is that observers were unable to correctly discriminate between rigid and nonrigid motion. Observers provided accurate judgments in only 47% of the cases. As we hypothesized, however, the perception of rigidity was strongly determined by the manipulation of *def*. *Different-def-rigid* stimuli were judged as "nonrigid" in 85% of the cases, even

Figure 12. Percentages of "rigid" judgments in Experiment 6 for both rigid (black) and nonrigid (white) displays and for the *same-def* and *different-def* conditions (*def* = deformation).

though these displays were mathematically correct projections of rigid transformations. *Same-def-nonrigid* stimuli were judged as "rigid" in 87% of the cases, even though these displays were projections of nonrigid transformations. In both of these conditions, the perceptual response was nonveridical. *Same-def-rigid* stimuli were judged as "rigid" in 90% of the cases, whereas *different-def-nonrigid* stimuli were judged as "nonrigid" in 75% of the cases. In these last two conditions the perceptual response was veridical.

A repeated measures ANOVA was performed on the frequencies of rigidity judgments provided by each participant in each experimental condition. The condition variable (*same-def* vs. *different-def*) reached significance, $F(1, 9) = 120.57$, $p < .001$, whereas the rigidity variable (3-D rigid vs. 3-D nonrigid displays) did not, $F(1, 9) = 2.19$, *ns*. Also nonsignificant was the interaction between these two variables. These results are consistent with the two hypotheses that motivated the present research: (a) For each simulated surface, the magnitude of perceived angular rotation depends on *def* and (b) a moving structure gives rise to the perception of rigid motion only if its component parts are perceived to rotate by the same amount.

In the previous discussion of this experiment's data, we made reference only to the deformations produced by the triplets of points belonging to the two simulated surfaces. In each stimulus display, however, there are also other deformations besides those produced by the two simulated surfaces (i.e., the deformations of the triplets made up of two points belonging to one surface and one point belonging to the other surface). One might wonder, therefore, why the perceptual system disregards the deformations of the triplets of points not belonging to each one of the two simulated surfaces.

To answer this question, let us consider the following simulation. We placed 15 points on each of the two simulated surfaces in the four experimental conditions, and we calculated *def* for each possible triplet of points. The number of possible triplets is $\binom{30}{3}$, or 4,060. The outcomes of this simulation, for the four experimental conditions, are reported in Figure 13. The figure indicates that in the *same-def* condition, the probability distribution of the deformations has a single peak (corresponding to the single *def* of the two simulated surfaces), whereas in the *different-def* condition, the probability distribution has two peaks. Apart from these peaks, all other deformations are associated with negligible probability values. By considering the relative probability value of each deformation, therefore, one could derive a single angle of rotation for the stimuli of the *same-def* condition and two different angles of rotation for the stimuli of the *different-def* condition. The judgments of rigidity could follow from this analysis: The stimulus displays could be judged as "rigid" if the probability distribution specifies a single *def* (i.e., a single magnitude of rotation); the stimulus displays could be judged "nonrigid" if the probability distribution specifies two different deformations (i.e., two magnitudes of rotation).

In conclusion, the previous considerations indicate that the model described in the introduction can also be applied to stimulus displays having a large number of dots (see also,
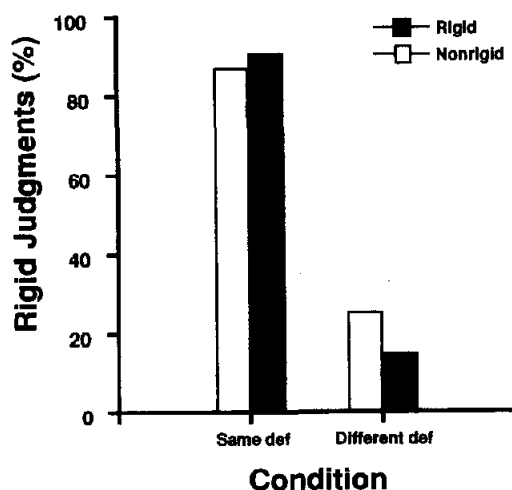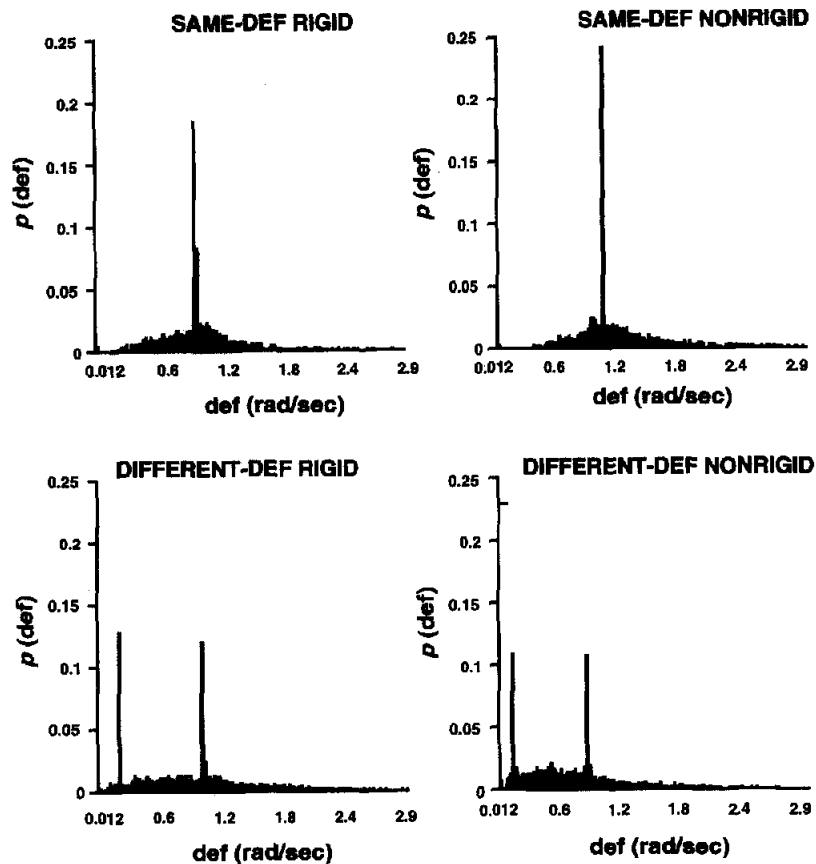
Figure 13. Probability distributions for simulated data generated with the same parameters that were used for the four stimulus types of Experiment 6. The plots represent the probability $p(def)$ for each possible triplet of the 15 points used in the simulation (def = deformation).

Domini, Caudek, & Proffitt, 1994). Moreover, the present experimental findings indicate that, as in the previous experiments, the perception of rigidity depends on the manipulation of the deformations of the triplets of points of the stimuli, regardless of whether the displays simulate a mathematically correct projection of a rigid transformation.

Experiment 7

In Experiment 6 we tried to control for the presence of artifactual sources of information (see Sperling, Landy, Dosher, & Perkins, 1989) by eliminating all pictorial depth cues. Even in the case of those displays, however, the possibility still remains that observers' performance could have been produced, without their detecting rigid or nonrigid motion in 3-D space, by their simply responding to the difference in the projected velocities of the two simulated surfaces (if we assume that the segregation of the stimulus features could take place independently from the SfM process). Observers' performance could have been produced by their simply responding "rigid" to the stimuli in which similar maximum projected velocities were associated with the two simulated surfaces and responding "nonrigid" in the other cases. Our purpose in Experiment 7 was to test for this

possibility by equating the maximum projected velocities of the simulated surfaces in each stimulus display.

Method

Participants. Eight University of Trieste undergraduates participated in this experiment. All of them were naive to the purpose of the experiment.

Stimuli. The stimuli were high-luminance dots moving on a low-luminance background. The horizontal motions of the dots were computed as simulating an orthographic projection of two surfaces oscillating about the y-axis. Each stimulus display was contained within a rectangular "window," the height and length of which subtended, respectively, 1.4° and 9° of visual angle. The two planar surfaces were simulated as meeting at a vertical edge coinciding with the axis of rotation. The position of the axis of rotation evenly bisected the length of the stimulus's window. The dots were randomly distributed with uniform probability density over the projection plane. Average dot density was equal to 4 dots/cm². Dot lifetime was manipulated so as to keep average dot density constant in each frame of the stimulus display. Within every oscillation cycle all dots of the stimulus display were randomly replaced by other dots placed in different positions of the stimulus's window. We defined four kinds of stimuli: (a) *same-def rigid*, in which the two surfaces were simulated as having a mean

slant equal to tan 60° and underwent the same angular rotation (10°); (b) *same-def nonrigid*, in which the two surfaces were simulated with different slants (tan 10° and tan 60°) and underwent different angular rotations (38.7° and 10°); (c) *different-def rigid*, in which the two surfaces were simulated with different slants (tan 10° and tan 60°) and underwent the same angular rotation (10°); and (d) *different-def nonrigid*, in which the two surfaces were simulated with different slants (tan 10° and tan 60°) and underwent different angular rotations (2.5° and 10°). The deformation of each surface in all of the conditions is calculated as the mean slant of the surface multiplied by the angular velocity (see Equation 1). We added a translatory motion component (parallel to the simulated surfaces) to the horizontal motions of the dots to equate the maximum projected velocities for the surfaces of the simulated dihedral angle. Maximum projected velocity for each simulated surface was equal to 3.43 deg/s. One oscillation cycle took 30 frames or 1.3 s.

Apparatus. The apparatus was the same as in Experiment 1.

Design. Two variables were studied in this experiment: (a) condition (*same-def* vs. *different-def*) and (b) rigidity (3-D rigid vs. 3-D nonrigid displays). Both were within-subjects variables. Each participant viewed nine presentations, in random order, of each combination of the within-subjects variables. Sixteen additional trials were presented at the beginning of each experimental session in order to familiarize the participants with the stimulus displays.

Procedure. The participants were instructed to discriminate rigid from nonrigid motion. Rigid stimuli were defined as the stimuli in which the angle between the two simulated surfaces remained constant during rotation; nonrigid stimuli were defined as the stimuli in which the angle between the two simulated surfaces changed during rotation. The experimenter was careful to explain that a translatory motion component was added to the horizontal projected motion in order to equate the maximum projected velocities of each simulated surface. Otherwise, the procedure and instructions were the same as in Experiment 4.

Results and Discussion

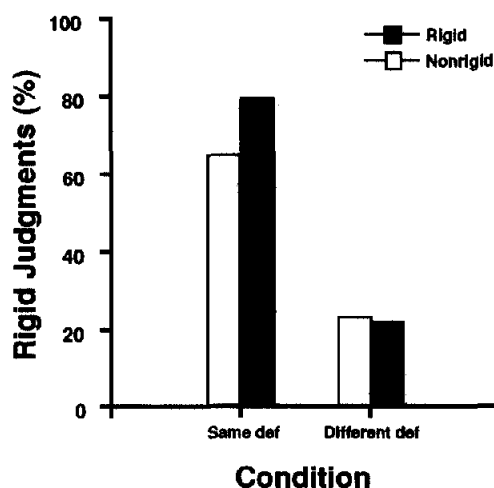The results are presented in Figure 14 and fully replicated those of the previous experiment. As in Experiment 6,

Figure 14. Percentages of "rigid" judgments in Experiment 7 for both rigid (black) and nonrigid (white) displays and for the *same-def* and *different-def* conditions (*def* = deformation).

observers were unable to correctly discriminate between rigid and nonrigid transformations. Overall, the judgments of the participants were veridical in only 47% of the cases. The perception of rigidity was strongly influenced by the manipulation of *def*, even though the same maximum projected velocity was associated with each of the two simulated surfaces in the stimulus displays. *Different-def-rigid* stimuli were judged "nonrigid" in 78% of the cases, even though these displays were mathematically correct projections of rigid transformations. *Same-def-nonrigid* stimuli were judged "rigid" in 65% of the cases, even though these displays were projections of nonrigid transformations. In both of these conditions, the perceptual response was non-veridical. *Same-def-rigid* stimuli were judged "rigid" in 80% of the cases, whereas *different-def-nonrigid* stimuli were judged "nonrigid" in 77% of the cases. In both of these conditions, the perceptual response was veridical.

A repeated measures ANOVA performed on the frequencies of rigidity judgments provided by each participant in each experimental condition indicated that the condition variable reached significance, $F(1, 7) = 40.92$, $p < .001$, whereas the rigidity variable and the interaction of these two variables were not significant. In conclusion, the results of the present experiment fully replicated those obtained in Experiment 6, even though we controlled for a potential 2-D confound that had been present in that experiment.

General Discussion

In Experiments 1–3, we examined the perception of angular velocity in SfM displays. We hypothesized that when *def* is pitted against 3-D angular velocity, perceived 3-D angular velocity covaries with *def* rather than with the simulated angular velocity. Our findings support this hypothesis. Moreover, we found no difference in the perceived angular magnitudes for two-view and multiview displays, which is consistent with the hypothesis that the perceptual parameters are heuristically derived from the first-order properties of the optic flow.

Experiments 4–7 were designed to investigate the perception of rigidity in the kinetic depth effect. We hypothesized that the perception of rigidity depends on the magnitudes of angular rotation perceived for the component parts of a moving object. If the component parts of an object are perceived to move by the same amount, then the object gives rise to the perception of an overall rigid motion; conversely, if the component parts of an object are perceived to move by different amounts, then the object gives rise to the perception of an overall nonrigid motion. From this hypothesis it follows that misperceptions of angular rotation (revealed by the first three experiments) should lead to misperceptions of rigidity.

We found that the manipulation of *def* influences not only the magnitudes of perceived angular velocity but also the perception of rigidity in the kinetic depth effect. In four experiments we showed that the perceptual discrimination between displays in which all points were either moving rigidly or rotating by different magnitudes about the same

axis (a) was strongly determined by the manipulation of the variability of the deformations of the triplets of image features in the stimulus displays and (b) was inconsistent with a mathematically correct analysis of the stimulus displays. When the variability of the deformations was low, observers judged the stimuli as undergoing rigid motion; conversely, when the variability of the deformations was high, observers judged the stimuli as undergoing nonrigid motion. This result was obtained both in the minimal conditions theoretically sufficient for discriminating rigid from nonrigid motion (Experiment 4) and with displays providing a larger number of points and views (Experiments 5–7).

The present findings cannot be accounted for by current approaches proposed to deal with the problem of the perceptual discrimination of rigid from nonrigid motion. According to the algorithm proposed by Ullman (1979), it is sufficient to compare the curl components of the image features. This approach, however, cannot explain the results of Experiments 4–7, in which observers consistently judged as "rigid" the orthographic projections of nonrigid moving objects and as "nonrigid" the projections of rigidly moving objects. A similar argument can also be applied to the *affine approach* (Koenderink & Van Doorn, 1991; Norman & Todd, 1992, 1993; Todd & Bressan, 1990). According to the affine approach, the discrimination between rigid and nonrigid motion is restricted to the discrimination between linear and nonlinear transformations along the line of sight. This approach, therefore, cannot explain why nonlinear transformations (our nonrigid stimuli) should have been judged "rigid" more often than linear transformations (our rigid stimuli).

We propose that the discrimination between rigid and nonrigid motion is performed by the perceptual system in a heuristic manner. The heuristic approach to the SfM problem was first proposed by Braunstein (1972, 1994). We suggest that such an approach can be applied to our findings as well as to the Braunstein et al. (1990) data. When the nonrigid displays are created by making each point rotate about a different axis of rotation, the variable that covaries with the 3-D rigidity is the variance of the trajectories. The heuristic used by the perceptual system, in this case, is to consider as nonrigid the displays with higher variance of the trajectories. In order to control this variable, we created another type of nonrigidity by rotating each point about the same axis with different angular velocities. Furthermore, we selected the rigid and nonrigid stimuli in order to keep the variance of the deformations for the rigid stimuli at least three times larger than the variance for the nonrigid stimuli. In our view, observers compute the magnitudes of perceived 3-D angular velocity by considering the first-order information of the optic flow according to the heuristic in which greater deformations are associated with greater angular velocities, independently of the simulated angular velocities. The model described in the introduction postulates an analysis performed independently for each triplet (or subset of triplets) of image features. Once the angular velocity of the individual triplets has been computed, a process of integration takes place. The judgment of rigidity is based on the outcome of this process. In this framework, the judg-

ment of rigidity is not an either–or judgment but instead could be based on the computation of the variability (e.g., the variance) of the angular velocities associated with the individual triplets or clusters of image features.

Conclusions

The perception of angular velocity and the perceptual discrimination of rigid from nonrigid motion may be based on an analysis of first-order temporal relations (i.e., the information provided by two views), even if additional views are available. For two-view and multiview displays, one can predict perceptual performance by considering solely the deformations of the optic flow. In three experiments, the main determinant of the perceived magnitudes of angular velocity in SfM displays was *def* rather than the magnitude of the simulated angular velocities. Moreover, in four experiments, stimuli with a low variability of the deformations of the individual triplets of image features tended to be judged as "rigid" more often than did those with a high variability of the deformations, regardless of whether they simulated a rigid transformation. These findings challenge the hypothesis that a rigidity constraint is embedded in the perceptual analysis of depth from motion.

References

Bennett, B. M., & Hoffman, D. D. (1985). The computation of structure from fixed axis motion: Non-rigid structures. *Biological Cybernetics, 51,* 293–300.

Bennett, B. M., Hoffman, D. D., Nicola, J. E., & Prakash, C. (1989). Structure from two orthographic views of rigid motion. *Journal of the Optical Society of America, A6,* 1052–1069.

Braunstein, M. L. (1972). Perception of rotation in depth: A process model. *Psychological Review, 79,* 510–524.

Braunstein, M. L. (1994). Decoding principles, heuristics and inference in visual perception. In G. Jansson, S. S. Bergstrom, & W. Epstein (Eds.), *Perceiving events and objects.* (pp. 436–446). Hillsdale, NJ: Erlbaum.

Braunstein, M. L., Hoffman, D. D., & Pollick, F. E. (1990). Discriminating rigid from non-rigid motion: Minimum points and views. *Perception & Psychophysics, 47,* 205–214.

Cortese, J. M., & Andersen, G. J. (1991). Recovery of 3-D shape from deforming contours. *Perception & Psychophysics, 49,* 315–327.

Cutting, J. E. (1987). Perception and information. *Annual Psychological Review, 38,* 61–90.

Domini, F., Caudek, C., & Gerbino, W. (1995). Perception of surface attitude in structure from motion displays. *Investigative Ophthalmology & Visual Science, 36,* 1694. (Abstract)

Domini, F., Caudek, C., & Proffitt, D. (1994). Discriminating rigid from non-rigid motion in the Kinetic Depth Effect. *Perception, 23*(Suppl.), 56a.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Hildreth, E. C. (1984). *The measurement of visual motion.* Cambridge, MA: MIT Press.

Hoffman, D. D. (1982). Inferring local surface orientation from motion fields. *Journal of the Optical Society of America, A72,* 888–892.

Hoffman, D. D., & Bennett, B. M. (1985). Inferring the relative

3-D positions of two moving points. *Journal of the Optical Society of America, A2,* 350–353.

Hoffman, D. D., & Bennett, B. M. (1986). The computation of structure from fixed-axis motion: Rigid structures. *Biological Cybernetics, 54,* 71–83.

Hoffman, D. D., & Flinchbaugh, B. E. (1982). The interpretation of biological motion. *Biological Cybernetics, 42,* 195–204.

Huang, T., & Lee, C. (1989). Motion and structure from orthographic projections. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 11,* 536–540.

Johansson, G. (1978). About the geometry underlying spontaneous visual decoding of the optical message. In E. Leeuwenberg & H. Buffart, *Formal theories of visual perception* (pp. 265–276). Chichester, England: Wiley.

Koenderink, J. J. (1986). Optic flow. *Vision Research, 26,* 161–180.

Koenderink, J. J., & Van Doorn, A. J. (1975). Invariant properties of the motion parallax field due to the movement of rigid bodies relative to an observer. *Optica Acta, 22,* 773–791.

Koenderink, J. J., & Van Doorn, A. J. (1976). Local structure of movement parallax of the plane. *Journal of the Optical Society of America, A66,* 717–723.

Koenderink, J. J., & Van Doorn, A. J. (1977). How an ambulant observer can construct a model of the environment from the geometrical structure of the visual flow. In G. Hauske & F. Butenand (Eds.), *Kibernetik* (pp. 224–247). Munich, Germany: Oldenberg.

Koenderink, J. J., & Van Doorn, A. J. (1981). Exterospecific component of the motion parallax field. *Journal of the Optical Society of America, A71,* 953–957.

Koenderink, J. J., & Van Doorn, A. J. (1986). Depth and shape from differential perspective in the presence of bending deformations. *Journal of the Optical Society of America, A3,* 242–249.

Koenderink, J. J., & Van Doorn, A. J. (1991). Affine structure from motion. *Journal of the Optical Society of America, A8,* 377–385.

Liter, C. J., Braunstein, M. L., & Hoffman, D. D. (1994). Inferring structure from motion in two-view and multiview displays. *Perception, 22,* 1441–1465.

Longuet-Higgins, H. C. (1982). The role of vertical dimension in stereoscopic vision. *Perception, 11,* 377–386.

Longuet-Higgins, H. C., & Prazdny, K. (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society of London,* Series B, *208,* 385–397.

Marascuilo, L. A. (1970). Extensions of the significance test for one-parameter signal detection hypotheses. *Psychometrika, 35,* 237–243.

Musatti, C. L. (1924). Sui fenomeni stereocinetici [On stereokinetic phenomena]. *Archivio Italiano di Psicologia, 6,* 205–219.

Norman, J. F., & Todd, J. T. (1992). The visual perception of 3-D dimensional form. In G. A. Carpenter & S. Grossberg (Eds.), *Neural networks for vision and image processing* (pp. 93–110), Cambridge, MA: MIT Press.

Norman, J. F., & Todd, J. T. (1993). The perceptual analysis of structure from motion for rotary objects undergoing affine stretching transformations. *Perception & Psychophysics, 53,* 279–291.

Pollick, F. E. (1989). Shape perception from dynamic occluding contours. *Investigative Ophalmology & Visual Science, 30*(Suppl. 3), 251.

Prazdny, K. (1980). Egomotion and relative depth map from optical flow. *Biological Cybernetics, 36,* 87–102.

Sperling, G., Landy, M. S., Dosher, B. A., & Perkins, M. E. (1989). Kinetic depth effect and identification of shape. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 826–840.

Todd, J. T., & Bressan, P. (1990). The perception of 3-dimensional affine structure from minimal apparent motion sequences. *Perception & Psychophysics, 48,* 419–430.

Todorovic, D. (1993). Analysis of two- and three-dimensional rigid and non-rigid motions in the stereokinetic effect. *Journal of the Optical Society of America, A10,* 804–826.

Ullman, S. (1977). *The interpretation of visual motion.* Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge.

Ullman, S. (1979). *The interpretation of visual motion.* Cambridge, MA: MIT Press.

Ullman, S. (1983). The interpretation of structure from motion. *Proceedings of the Royal Society of London,* Series B, *203,* 405–426.

Ullman, S. (1984). Rigidity and misperceived motion. *Perception, 13,* 219–220.

Wallach, H., & O'Connell, D. N. (1953). The kinetic depth effect. *Journal of Experimental Psychology, 45,* 205–217.

Waxman, A., & Wohn, K. (1987). Contour evolution, neighbourhood def, and image flow: Textured surfaces in motion. In W. Richards & S. Ullman (Eds.), *Image understanding 1985–86* (pp. 72–98). Norwood, NJ: Ablex.

Zanforlin, M. (1988). The height of a stereokinetic cone: A quantitative determination of a 3-D effect from 2-D moving patterns without a "rigidity assumption." *Psychological Research, 50,* 162–172.