

Human and Machine Vision, A. Rosenfeld et. al, ed., Academic Press, 1983, pp. 237-257;
Also available in Proc. The First Annual National Conference on Artificial Intelligence, 1980, pp. 4-6.

Mapping Image Properties into Shape Constraints: Skewed Symmetry, Affine-Transformable Patterns, and the Shape-from-Texture Paradigm

Takeo Kanade
John R. Kender

Carnegie-Mellon University
Pittsburgh, Pennsylvania

Abstract

In this paper we demonstrate two new approaches to deriving three-dimensional surface orientation information ("shape") from two-dimensional image cues. The two approaches are the method of affine-transformable patterns and the shape-from-texture paradigm. They are introduced by a specific application common to both: the concept of skewed symmetry. Skewed symmetry is shown to constrain the relationship of observed distortions in a known object regularity to a small subset of possible underlying surface orientations. Besides this constraint, valuable in its own right, the two methods are shown to generate other surface constraints as well. Some applications are presented of skewed symmetry to line drawing analysis, to the use of gravity in shape understanding, and to global shape recovery.

I. Introduction

Certain image properties, such as parallelism, symmetries, and repeated patterns, provide cues for perceiving 3-D shape from a 2-D picture. This paper demonstrates how we can map these image properties into 3-D shape constraints by associating appropriate assumptions with them and by using appropriate computational and representational tools.

We begin with the exploration of how one specific image property, “skewed symmetry,” can be defined and formulated to serve as a cue to the determination of surface orientations. Then we will discuss the issue from two new, broader viewpoints. One is the class of affine-transformable patterns. It has various interesting properties, and includes skewed symmetry as a special case. The other is the computational paradigm of shape-from-texture. Skewed symmetry is derived in a second, independent way, as an instance of the application of the paradigm. **Also**, it is proven that the same skewed-symmetry constraint can arise from greatly different image conditions.

This paper further claims that the ideas and techniques presented here are applicable to many other properties under a general framework of the shape-from-texture paradigm with the underlying meta-heuristic of non-accidental image properties.

2. Skewed Symmetry

In this section we assume the standard orthographic projection from scene to image, and a knowledge of the gradient space (see Mackworth, 1973).

2.1. Definition, Assumption and Constraints

Symmetry in a 2-D picture has an axis for which the opposite sides are reflective; in other words, the symmetrical properties are found along the transverse lines perpendicular to the symmetry axis. The concept *skewed symmetry* was introduced by Kanade (1979) by relaxing this condition a little. It means a class of 2-D shapes in which the symmetry is found along lines not necessarily perpendicular to the axis, but at a fixed angle to it. Formally, such shapes can be defined as 2-D affine transforms of real symmetries. Figures 1a-c show a few examples.¹

Stevens (1980) presents a number of psychological experiments which suggest that human observers can perceive surface orientations from figures with this property. This is probably because such qualitative symmetry in the image is often due to real symmetry in the scene. Thus let us associate the following assumption with this image property:

A skewed symmetry depicts a real symmetry viewed from some unknown view angle.

Note that the converse of this assumption is always true under orthographic projection.

We can transform this assumption into constraints in the gradient space. **As** shown in Figure 1, a skewed symmetry defines two directions: let us call them the skewed-symmetry axis and the skewed-transverse axis, and denote

¹The mouse hole example of Figure 1c is due to K. Stevens (1980).




Figure 1. Skewed symmetry.

their directional angles in the picture by α and β , respectively (Figure 1d). Let $G = (p, q)$ be the gradient of the plane which includes the skewed symmetry. The 3-11 vectors on the plane corresponding to the directions α and β are

$$(\cos\alpha, \sin\alpha, -p\cos\alpha - q\sin\alpha) \text{ and } (\cos\beta, \sin\beta, -p\cos\beta - q\sin\beta).$$

The assumption demands that these two vectors be perpendicular; their inner product vanishes:

$$\cos(\alpha - \beta) + (p\cos\alpha + q\sin\alpha)(p\cos\beta + q\sin\beta) = 0. \quad (1)$$

By rotating the p - q coordinates into the p' - q' coordinates so that the new p' - q' axes are the bisectors of the angle made by the skewed symmetry and skewed-transverse axes, it is easy to show that

$$p'^2 \cos^2\left(\frac{\alpha - \beta}{2}\right) - q'^2 \sin^2\left(\frac{\alpha - \beta}{2}\right) = -\cos(\alpha - \beta) \quad (2)$$

where

$$\begin{aligned} p' &= p \cos\lambda + q \sin\lambda \\ q' &= -p \sin\lambda + q \cos\lambda \\ \lambda &= (\alpha + \beta)/2. \end{aligned}$$

Thus, the (p, q) 's are on the hyperbola shown in Figure 2. That is, the skewed symmetry defined by α and β in the picture can be a projection of a real symmetry if and only if the gradient is on this hyperbola. The skewed symmetry

thus imposes a one-dimensional family of constraints on the underlying surface orientation (p, q) . As we will see in Section 5, other constraints can be exploited for the unique determination of surface orientation.

The tips or vertices G_T and G_T' of the hyperbola represent special orientations with interesting properties. First, since they are closest to the origin of the gradient space, and since the distance from the origin to a gradient represents the magnitude of the surface slant, G_T and G_T' correspond to the least slanted orientations that can produce the skewed symmetry in the picture from a real symmetry in the scene.

Second, since they are on the line (the axis of the hyperbola) which bisects the obtuse angle made by α and β , they correspond to the orientations for which the rates of depth change along the directions of α and β in the picture are the same. In other words, the apparent ratio of length to width of the object in the picture represents the real ratio in the scene (see Kanade [1979] for the proof.)

2.2. Rationale and Justification

Skewed symmetry has straightforward applications to scenes containing objects that have been manufactured, whether naturally or artificially. Many constructed items exhibit symmetry, occasionally about many axes.

Some symmetries are introduced due to economies of the manufacturing process: an object is often composed of identically formed component parts (fibers, cells, bricks, etc.). The symmetries result from the three-dimensional tessellation of the components into the whole. Often the tessellation is effectively two-dimensional, in laminae (cloth, honeycombs, walls, etc.), and the application of the skewed symmetry method is straightforward. Further, the requirement for a close symmetric packing of the components occasionally imposes a local symmetry on the individual components, too. The method can then be applied to individual parts (such as the bricks themselves). Notice the method does *not* assume 3-D symmetry of the whole object; what is assumed is *local* 2-D symmetry.

A further source of symmetry is the bilateral symmetry that results from biological manufacture (growth). It not only contributes symmetric objects to the environment: it may also be responsible for an imitative esthetic bias in human manufacture. If the extent of a bilaterally symmetric pattern into the third dimension is not too great (a face, a leaf, an airplane), the skewed symmetry method can be approximately applied also.

3. Affine-Transformable Patterns

In texture analysis we often consider small patterns (texels=texture elements) whose repetition constitutes "texture." Suppose we have a pair of texel




Figure 2. The hyperbola determined by a skewed symmetry defined by α and β

patterns in which one is a 2-D affine transform of the other; we call them a pair of **affine-transformable** patterns. Let us assume:

A pair of affine-transformable patterns in the picture are projections of similar patterns in the 3-D space (i.e., they can be overlapped by scale change, rotation, and translation).

Note that, as in the case of **skewed** symmetry, the converse of this assumption is always true under orthographic projection. The above assumption can be schematized by Figure 3. Consider two **texel** patterns P_1 and P_2 in the picture, and place the origins of the x - y coordinates at their centers, respectively. The transform from P_2 to P_1 can be now expressed by a regular 2×2 matrix $\mathbf{A} = (a_{ij})$. P_1 and P_2 are projections of patterns P'_1 and P'_2 which are drawn on the 3-D surfaces. We assume that P'_1 and P'_2 are small enough so that we can regard them as being drawn on **small** planes. Let us denote the gradients of those small planes by $G_1 = (p_1, q_1)$ and $G_2 = (p_2, q_2)$, respectively; i.e., P'_1 is drawn on a plane $-z = p_1x + q_1y$ and P'_2 on $-z = p_2x + q_2y$.

Now, **our** assumption amounts to saying that P_1 is transformable from




Figure 3. A schematic diagram showing the assumption about the affine-transformable patterns.

P_2 by a scalar scale factor σ and a rotation matrix:

$$R = \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix}.$$

(We can omit the translation from our consideration, since for each pattern the origin of the coordinates is placed at its gravity center, which is preserved under the affine transform.) Thinking about a pattern drawn on a small plane, $-z = px + qy$, is equivalent to viewing the pattern from directly overhead; that is, rotating the x - y - z coordinates so that the normal vector of the plane is along the new z -axis (line of sight). For this purpose we rotate the coordinates first by ϕ around the y -axis and then by θ around the x' -axis (Figure 4). We have




Figure 4. Rotation of the x - y - z coordinates.

the following relations among ϕ , θ , p , and q :

$$\begin{aligned} \sin\phi &= p/\sqrt{p^2+1}, & \cos\phi &= 1/\sqrt{p^2+1}, \\ \sin\theta &= q/\sqrt{p^2+q^2+1}, & \cos\theta &= \sqrt{p^2+1}/\sqrt{p^2+q^2+1}. \end{aligned} \quad (3)$$

Further, let Γ denote the angle of slant of the pattern, i.e., the angle between the old and the new z axes. Then

$$\cos\Gamma = 1/\sqrt{p^2+q^2+1}. \quad (4)$$

The plane which was represented as $-z = px + qy$ in the old coordinates is, of course, now represented as $-z' = 0$ in the new coordinates.

Let us denote the angles of the coordinate rotations to obtain P_1 and P_2 in Figure 3 by (ϕ_1, θ_1) and (ϕ_2, θ_2) , respectively. The 2-D mapping from P_i (x' - y' plane) to P_j (x - y plane) can be conveniently represented by the following 2×2 matrix T , which is actually a submatrix of the usual 3-D rotation matrix:

$$T_i = \begin{pmatrix} \cos\phi_i & -\sin\phi_i \sin\theta_i \\ 0 & \cos\theta_i \end{pmatrix}$$

Now, in order for the schematic diagram of Figure 3 to hold, what relationships have to be satisfied among the matrix $A = (a_{ij})$, the gradients $G_i = (p_i, q_i)$ for $i = 1, 2$, the angles (ϕ_i, θ_i) for $i = 1, 2$, the scale factor σ , and the matrix R ? We equate the two transforms that start from P_2' to reach P_1 : one following the diagram counter-clockwise, $P_2' \rightarrow P_2 \rightarrow P_1$, and the other clockwise, $P_2' \rightarrow P_1$. We obtain

$$AT_2 = T_1\sigma R$$

That is,

$$\begin{aligned}
 a_{11}\cos\phi_2 &= \sigma(\cos\alpha\cos\phi_1 - \sin\alpha\sin\phi_1\sin\theta_1) \\
 a_{12}\cos\theta_2 - a_{11}\sin\phi_2\sin\theta_2 &= -\sigma(\sin\alpha\cos\phi_1 + \cos\alpha\sin\phi_1\sin\theta_1) \\
 a_{21}\cos\phi_2 &= \sigma\sin\alpha\cos\theta_1 \\
 a_{22}\cos\theta_2 - a_{21}\sin\phi_2\sin\theta_2 &= \sigma\cos\alpha\cos\theta_1.
 \end{aligned} \tag{5}$$

By eliminating σ and α and substituting for $\sin\phi_i$, $\cos\phi_i$, $\sin\theta_i$, and $\cos\theta_i$ from (3), we have the following equations in p_1 , q_1 , p_2 , and q_2 :

$$\begin{aligned}
 &\sqrt{p_2^2 + q_2^2 + 1}(a_{11}(p_1^2 + 1) + a_{21}p_1q_1) \\
 &\quad = \sqrt{p_1^2 + q_1^2 + 1}(a_{22}(p_2^2 + 1) - a_{21}p_2q_2) \\
 &\quad (-a_{12}(p_2^2 + 1) + a_{11}p_2q_2)(p_1^2 + 1) - (a_{22}(p_2^2 + 1) - a_{21}p_2q_2)p_1q_1 \\
 &\quad = a_{21}\sqrt{p_1^2 + q_1^2 + 1}\sqrt{p_2^2 + q_2^2 + 1}.
 \end{aligned} \tag{6}$$

We thus find that the assumption of affine-transformable patterns yields the constraint represented by (6) on surface orientations. The constraint is determined solely by the matrix $A = (a_{ij})$, which is determined by the relation between P_2 and P_1 observable in the picture without knowing either the original patterns (P_1' and P_2') or their relationships (σ and R) in the 3-D space.

In order to have an idea about the degree of the constraint represented by (6), if we assume that the orientation of P_2' is known (i.e., $G_2 = (p_2, q_2)$ is known), then (6) gives two simultaneous equations for $G_1 = (p_1, q_1)$. The system appears to be of degree 4, but it can be shown that there are only two solutions; they are of the form (p_0, q_0) and $(-p_0, -q_0)$, which are symmetrical around the origin of the gradient space (see the Appendix).

From (5) we can also derive the following relationship:

$$\frac{\det(A)}{\sigma^2} = \frac{\sqrt{p_2^2 + q_2^2 + 1}}{\sqrt{p_1^2 + q_1^2 + 1}} - \frac{\cos\Gamma_1}{\cos\Gamma_2}. \tag{7}$$

This means that the ratio of cosines of the slant angles of the patterns is equal to the ratio $\det(A)/\sigma^2$. If we assume $\sigma = 1$ (the original patterns are of the same size) or that σ is known, (7) shows that we can order the texel patterns according to the magnitude of slant, Γ_i or $\sqrt{p_i^2 + q_i^2}$, using the values of $\det(A)$.

²This indicates that $\det(A)$ should be positive. But if it is negative, then we can assume that P_1' and P_2' are mirrored patterns, and put $R = \begin{pmatrix} \cos\alpha & \sin\alpha \\ \sin\alpha & -\cos\alpha \end{pmatrix}$.

3.1. Skewed Symmetry from Affine-transformable Patterns

The affine transform from P_2 to P_1 is more intuitively understood by how a pair of perpendicular unit-length vectors (typically along the x and y coordinate axes) are mapped into their transformed vectors. As shown in Figure 5, two angles (α and β) and two lengths (τ and ρ) can characterize the transform. The components of the transformation matrix $\mathbf{A} = (a_{ij})$ are represented by

$$\begin{aligned} a_{11} &= \tau \cos \alpha & a_{12} &= \rho \cos \beta \\ a_{21} &= \tau \sin \alpha & a_{22} &= \rho \sin \beta. \end{aligned} \quad (8)$$

Suppose, for simplicity, the orientation of P_2 in Figure 3 is known to be $(p_2, q_2) = (0, 0)$. This simplifies equation (6) to

$$\begin{aligned} a_{11}(p_1^2 + 1) + a_{21}p_1q_1 &= a_{22}\sqrt{p_1^2 + q_1^2 + 1} \\ -a_{12}(p_1^2 + 1) - a_{22}p_1q_1 &= a_{21}\sqrt{p_1^2 + q_1^2 + 1}. \end{aligned} \quad (9)$$

If we assume that α , β , τ , and ρ are known, then (p_1, q_1) has two possible solutions. This is essentially the case which Ikeuchi (1980b) investigated in his shape recovery method by assuming a known standard pattern, even though he used the constraint only partially.

Let us consider the case where α and β are known, but τ and ρ are not. One can substitute a_{ij} in (9) by (8), and eliminate τ and ρ . Then we obtain

$$(p_1 \cos \alpha + q_1 \sin \alpha)(p_1 \cos \beta + q_1 \sin \beta) + \cos(\alpha - \beta) = 0$$

which reduces to the same as the hyperbola (1). This can be interpreted as follows.

As was noted in the previous subsection, a pair of affine-transformable patterns impose the constraints (6) between their surface orientations, in which, if one is fixed, the other has only two possible orientations. However, if we loosen the transform in such a way that the angular (rotational) correspondence (α and β) is known while the length relationship is not known (or arbitrary), then the one-dimensional constraint of the skewed-symmetry hyperbola is obtained.




Figure 5. An affine transform (without translation) as characterized by two angles and two lengths.

4. The Shape-from-texture Paradigm

This section derives the same skewed-symmetry constraints from a second theory, different from that of the affine-transformable patterns. The shape-from-texture paradigm is a method of relating image texture properties to scene object properties, by explicitly incorporating assumptions about the imaging phenomenon into a computational framework. The paradigm is briefly presented here, but a fuller discussion can be found in (Kender, 1980).

The paradigm has two major portions. In the first, a given image textural property is “normalized” to give a general class of surface orientation constraints. In the second, the normalized values are used in conjunction with assumed scene relations to refine the constraints. If there are sufficiently many textural elements (“texels”) in the image to be normalized, and if enough assumptions are made about their scene counterparts, then the underlying surface’s orientation can be specified uniquely. Somewhat more weakly, only two texels are required, and only one assumption (equality of scene textural objects, or some other simple relation), to generate a well-behaved one-dimensional family of possible surface orientations. The method of skewed symmetry – the use of qualitative symmetries in the image to create a perspective distorted right angle – is an example of such a weak method.

The first step in the paradigm is the normalization of a given texel property. The goal is to create a normalized texture property map (NTPM), which is a representational and computational tool relating image properties to scene properties. The NTPM summarizes the many different conditions that may have occurred in the scene leading to the formation of the given textural element. In general, the NTPM of a certain property is a scalar-valued function of two variables. The two input variables describe the postulated surface orientation in the scene (top-bottom and left-right slants: (p, q) when we use the gradient space). The NTPM for a horizontal unit line length in the image summarizes the lengths of lines that would have been necessary in 3-D space under various orientations: at surface orientation (p, q) , it would have to be $\sqrt{p^2 + q^2 + 1}$.

More specifically, the NTPM is formed by selecting a texel and a texel property, back-projecting the texel through the known imaging geometry onto all conceivable surface orientations, and measuring the texel property there. The representation chosen for the two-dimensional space of orientations is important; we will, however, only use the gradient space here.

In the second phase of the paradigm, the NTPM is refined in the following way. Texels usually have various orientations in the image, and there are many different texel types. Each texel generates its own image-scene relationships, summarized in its NTPM. If, however, assumptions can be made to relate one texel to another, then their NTPMs can also be related; in most cases only a few scene surface orientations can satisfy *both* texels’ requirements. Some examples of the assumptions that relate texels are: both lie in the same

plane, both are equal in textural measure (length, area, etc.), one is k times the other in measure, etc. Relating texels in the manner forces more stringent demands on the scene. If enough relations are invoked, the orientation of the local surface supporting two or more related texels can be very precisely determined.

4.1. Skewed Symmetry from the Paradigm Applied to Slope

What we now show is that the skewed symmetry method is a special case of the shape-from-texture paradigm; it can be derived from considerations of texel slope.

To normalize the slope of a texel, it is back-projected onto a plane with the postulated orientation. The back-projected texel now has a new shape on this new surface. Its exact value, however, depends upon the coordinate system on this surface plane. Many coordinate systems are possible; we chose here a coordinate system whose x-axis lies along the gradient direction. The normalized slope is then the angle that the back-projected texel makes with respect to the surface coordinate system x-axis. The calculation is a bit involved, especially under perspective, which requires a knowledge of both the location of the center of focus and the length of the focal distance.

Using the construction in Figure 6, together with several lemmas relating surfaces in perspective to their local vanishing lines, slope is normalized as follows. Assume a slope is parallel to the p-axis; the image and gradient space can always be rotated into such a position. (If rotation is necessary, the resulting NTPM can be de-rotated into the original position using the standard two-by-two orthonormal matrix.) Also assume that the slope is somewhere along the line $y = y_s$, where the unit of measurement in the image is equal to one focal length. The normalized value of the slope is equal to the tangent of the 3-D space angle η , whose base (of length R) is parallel to the surface plane, and is in the direction of the gradient. R is determined from the focal distance, and from the point of the nearest approach of the vanishing line of the plane. This line has equation $px + qy = 1$ (or $G \cdot P = 1$) and its nearest approach is $G/\|G\|^2$. The distance d is given by the intersection of the line $y = y_s$ with the vanishing line. Then, the normalized slope value – the Normalized Texture Property Map – is given by

$$\frac{q - y_s(p^2 + q^2)}{p \sqrt{1 + p^2 + q^2}} \quad (10)$$

This normalized value can be exploited in several ways. Most important is the result that is obtained when one has two slopes in the image that are assumed to arise from equal slopes in the scene. Under this assumption, their normalized property maps can be equated. The resulting constraint, surprisingly, is a simple straight line in the gradient space. It is intimately related to the vanishing




Figure 6. Back-projecting an image slope onto a plane with gradient (p, q) .

point formed by the intersection of the extensions of the two image slopes (Kender, 1980).

The constraint equations resulting from assuming that the two slopes arose from perpendicular lines in the scene is, however, enormously complex. It unfortunately does not appear to have many tractable forms or special cases.

Under orthography, nearly everything simplifies. The normalized slope of a texel becomes

$$\frac{q}{p\sqrt{1+p^2+q^2}}. \quad (11)$$

It is independent of y_s ; in effect, all slopes are at the focal point.

Considering two image slopes to have arisen from parallel lines in the scene has a trivial solution. If the image slopes are parallel, the entire gradient space is a solution. If they are not, there is no solution at all. This corresponds to the projective geometry theorem that under orthography, parallels are taken into parallels regardless of surface orientation.

In the case where the scene slopes are assumed to be perpendicular, we again get a simplification, but this time a useful one. Not only is the solution tractable, it is the skewed symmetry method of Section 2. We derive it as follows

Consider Figure 7. Note that under orthography, texels can be translated arbitrarily, since the focal length is infinite and the focal point is effectively everywhere; there is no information in image position. Given the angle that the two texels form, rotate the gradient space so that the positive p -axis bisects the angle. Call this adjustment angle A ; we will use it to de-adjust our results into the original position after they have been computed.




Figure 7. Two image texels assumed to be perpendicular in the scene.

Let the angle that is bisected be 2δ . The normalized value of either slope is obtained directly from the standard normalized slope formula, corrected for the displacement of $+\delta$ and $-\delta$ respectively. That is, for the slope at the positive θ orientation, instead of formula (11), we use the formula under the substitution $p \cos \delta + q \sin \delta$ for p , $-p \sin \delta + q \cos \delta$ for q . We proceed similarly for the slope at $-\theta$. Note that the factor $\sqrt{1+p^2+q^2}$ is invariant under this transformation (it is the length of the normal vector of the surface).

The fact that the normalized slopes are assumed to be perpendicular in the scene allows us to set one of the normalized values equal to the negative reciprocal of the other. The resultant equation becomes

$$p^2 \cos^2 \delta - q^2 \sin^2 \delta = \sin^2 \delta - \cos^2 \delta = -\cos 2\delta. \quad (12)$$

This is exactly the hyperbola in Section 2 with $2\delta = |\alpha - \beta|$.

4.2. Skewed Symmetry from the Paradigm Applied to Length and Angle

The paradigm is similarly applicable to other texture measures. Using texel length as the property to be normalized, we find that under perspective, lengths must lie on the same line in order for the resultant equations to be simpler than the fourth order. If they are collinear, again the resultant gradient space constraint is a simple straight line.

Under orthography and the assumption that image lengths have arisen from equal scene lengths, the constraint equation is again a hyperbola – the skewed-symmetn hyperbola, somewhat offset. In fact, the geometric construction in Figure 8 shows that the assumption of equal length can be made equivalent to skewed symmetry.

First, a triangle is formed by translating one or the other of the lengths so that they meet at a common endpoint. Under orthography, such translations do not affect the resulting constraints. Connecting the remaining endpoints creates a triangle which must be isosceles in the scene. Further, under orthography,




Figure 8. Assuming lengths are equal generates the skewed symmetry constraint.

midpoints of lines are preserved (the midpoint of the base of the scene triangle is imaged as the midpoint of the base of the image triangle). The line connecting the vertex and this midpoint has the property that, in the scene, it must form a right angle with the base. Its distortion to something other than a right angle in the image – the induced angle 2δ – is precisely the distortion which characterizes skewed symmetry. Therefore, the same methods apply.

One other case is worth mentioning. Suppose the image has two angles such that one leg of the first is parallel to one leg of the second. See Figure 9. In this case, again the constraint is equivalent to skewed symmetry, as the construction shows. Choosing one of the angles, extend its non-parallel leg until it intersects both legs of the other angle. (If it cannot do so, then first translate the angle before extending.) The resulting triangle *must* be isosceles in the scene, since the angles are assumed equal in the scene. However, this is the same situation encountered above with the construction involving lengths. Therefore, the altitude from the midpoint of the base (here, the midpoint of the parallel side) to the vertex must form a right angle. Again, the distortion observed in the image is the skewed symmetry distortion.

5. Applications of Skewed Symmetry and Affine-transformable Patterns

5.1. Quantitative Shape Recovery from Line Drawings

Given the line drawing of Figure 10a, we usually perceive a right-angled parallelepiped. The Huffman-Clowes-Waltz labeling scheme for the trihedral




Figure 9. Assuming angles are equal generates the skewed symmetry constraint.




Figure 10. (a) A line drawing of a block; (b) Huffman-Clowes-Waltz labeling; (c) constraints in the gradient space.

world gives the labeling shown in Figure 10b, which signifies that the three edges meeting at the central FORK vertex are all convex, i.e., the object is a convex corner of a block. However, it does *not* specify a particular quantitative shape. In fact, the labeling indicates only that the gradients of the three surfaces should be placed in the gradient space so as to form the triangle shown in Figure 10c. The edges of the triangle should be perpendicular to the picture edges separating the corresponding regions, but the location and size of the triangle are arbitrary in the gradient space. Therefore, the object is not necessarily right-angled.

We can use skewed symmetry here to provide additional constraints. The three regions are skewed-symmetrical with the axes shown in Figure 11a. The hyperbolas corresponding to these regions are shown in Figure 11b. Thus the problem is now how to place the triangle of Figure 10c in Figure 11b so that each vertex is on the corresponding hyperbola. Kanade (1979) proves that the combination of locations shown in Figure 11b is the only possibility, and that the resultant shape is a right-angled block.

It is interesting to note that if we apply the same procedure to the line drawing of Figure 12, we find that there is no way for all the three regions to satisfy the skewed symmetry assumptions. That is, at least one of them has to be non-symmetrical (skewed) in the 3-D space; in other words, the object *cannot* be right angled, but should be rhomboid (a prism). Remember that Figure 10a *can* be either right-angled or rhomboid, but it is usually perceived as right-angled.

Figure 13 demonstrates how the above procedure results in the interpretation of the drawing as a trapezoidal block in this case.

5.2. Skewed Symmetry under Gravity

One principal influence toward symmetry seems to be an object's structural necessity to oppose the gravitational field. Objects that must support




Figure 11. (a) **Axes** of the skewed symmetry of the regions of Figure 10a; (b) corresponding hyperbolas and allocations of the gradients.

themselves tend to have structural members aligned **parallel** to the direction of force, that is, vertically. Such members are mutually **parallel** – a type of symmetry. The base of such an object is often perpendicular to gravity to distribute weight and provide balance. Together, then, the base and structural members provide a local symmetry frame that can also be exploited by the skewed-symmetry method. One can show that in this last case it is usually possible to specify surface orientation uniquely.

We will assume that the direction of the gravity field is known, say the top-to-bottom lines in the image frame are assumed to be true projections of a line of gravity force. The gradient space is also considered to be aligned in the direction of the gravity field; $-q$ is also “down.”

Under such conditions, suppose we do find a **portion** of the image that is assumed to be a vertical, symmetric surface: say, a building face as in Figure




Figure 12. A line drawing of a rhomboid: this *cannot* be a right-angled block. Notice that Figure 10a can be a rhomboid.




Figure 13. Shape recovery of a trapezoid block: (a) axes; (b) gradient allocations.

14. Using skewed symmetry (or even direct observation), it is not hard to obtain an angle in the image that corresponds to a right angle in the scene. Suppose one of the legs of the angle is parallel to the known gravity field as in Figure 14. The skewed-symmetry method generates the following constraint hyperbola:

$$p = -(q + 1/q)\cot\gamma. \quad (13)$$

This constraint is somewhat interesting: it expresses p (left-right slant) as a function of q (top-bottom slant). The value of q itself is easily obtained.

If gravity points in the $-z$ direction, the ground plane must have as its orientation $(0, q_g)$, for a value of q_g determinable through sensing. Since all




Figure 14. Assumptions about gravity can uniquely specify surface orientations.

vertical planes are perpendicular to the ground plane, all vertical planes must have the orientation $(p_v, -1/q_g)$, for variable p_v . (A quick check shows that the dot product of the corresponding normals is zero: $(0, q_g, 1) \cdot (p_v, -1/q_g, 1) = 0$.) Note that the value of q for any vertical plane is fixed at $-1/q_g$. Thus, in our example, p is also determined: it is $-(q_g + 1/q_g)\cot\gamma$. Since q_g is a constant, p varies simply with y . Figure 14 shows the constraints graphically.

5.3. Shape Recovery of an Object with Many Patterns Stamped

Consider the problem of recovering the shape from a picture of a ball with a number of patterns stamped on it (see Figure 15). For each pair of texel patterns, if they are affine-transformable, we compute a transformation matrix A . Thus we obtain many constraints on the gradients of texels. From these, however, we cannot uniquely determine the surface orientation of each texel.

We need more assumptions or data. We will suppose we know the gradients of some particular texels, and assume that the surface is smooth (together, maybe, with an assumption of global concavity or convexity). Then a relaxation or cooperative technique similar to the one for shape-from shading (Woodham, 1977; Ikeuchi, 1980a) will allow us to determine consistent assignments of gradients to the texels which satisfy those many constraints. Notice that we need not assume that the original pattern is known, nor that the patterns are stamped in a particular manner. Even other patterns can be mixed together with them.

One of the plausible methods of determining the gradient of one particular texel is to use equation (7). Assuming $\sigma = 1$, we order the texels by the magnitude of $\sqrt{p_i^2 + q_i^2}$, and assign $(p, q) = (0, 0)$ (the orientation that is directly facing the viewer) to the least slanted texel. This is analogous to a similar hypothesis in shape from shading. That is, we tend to assign to the brightest point the orientation directly facing the light source, even though under the assumptions of parallel lights and a matte surface, one can only say that the brightest pixels have the minimum incident angle of light, not necessarily 0°.

6. Conclusion

The assumptions we used for skewed symmetry, affine-transformable patterns, and texture analysis can be generalized as:

Properties observable in the picture are not accidental, but are projections of some preferred Corresponding 3-D properties.

This provides a useful meta-heuristic for exploiting image properties: we can call it the meta-heuristic of *non-accidental image properties*. It can be regarded as a generalization of general view directions, often used in the **blocks** world, to exclude the cases of accidental line alignments.





Figure 15. A picture of a ball with a number of  s stamped.

Instances that can **fall** within this meta-heuristic include: parallel lines in the picture vs. **parallel** lines in the scene, texture gradients due to distance, **and** sets of lines convergent to a vanishing point.

The most essential point of our technique is that we relate certain image properties to certain 3-D space properties, and that **we** map the relationships into convenient representations of shape constraints. We explicitly incorporate assumptions based either on the meta-heuristic or on *a priori* knowledge of the world. The shape-from-texture paradigm provides a computational framework for our technique. In most of our discussion we assumed orthography. Similar (though more involved **and** less intuitive) results can be obtained under perspective projection.

Acknowledgment

This research was sponsored by the Defense Advanced Research Projects Agency (DOD) under ARPA Order No. 3597, and monitored by the Air Force

Avionics Laboratory under Contract F33615-78-C-1551. The views and conclusions in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Appendix

Proof that (6) has two symmetrical *solutions*:

We will try to solve (6) for p_1 and q_1 , assuming that p_2, q_2 , and $A = (a_{ij})$ are known. We assume $\det(A) > 0$. Let us put $\gamma = \sqrt{p_1^2 + q_1^2 + 1}$. Then (6) can be rewritten as

$$\begin{aligned} Ca_{11}(p_1^2 + 1) + Ca_{21}p_1q_1 &= B\gamma \\ A(p_1^2 + 1) - Bp_1q_1 &= Ca_{21}\gamma \end{aligned} \quad (14)$$

where

$$\begin{aligned} A &= a_{11}p_2q_2 - a_{12}(p_2^2 + 1) \\ B &= a_{22}(p_2^2 + 1) - a_{21}p_2q_2 \\ C &= \sqrt{p_2^2 + q_2^2 + 1}. \end{aligned}$$

We can derive a quadratic equation on γ from (14):

$$f(\gamma) = DE\gamma^2 - (D^2 + E^2 + F^2)\gamma + DE = 0, \quad (15)$$

where

$$\begin{aligned} D &= C(Ba_{11} + Aa_{21}) = C \det(A)(p_2^2 + 1) > 0 \\ E &= B^2 + (Ca_{21})^2 > 0 \\ F &= -C^2a_{11}a_{21} + AB. \end{aligned}$$

The discriminant of (15) is

$$\begin{aligned} \text{disc} &= (D^2 + E^2 + F^2)^2 - 4(DE)^2 \\ &= F^4 + 2F^2(D^2 + E^2) + (D - E)^2 \\ &\geq 0. \end{aligned}$$

Thus, $f(\gamma)$ has real roots. Now, notice that $\gamma \geq 1$ and thus we are interested in the root greater than or equal to 1. Let us check the sign of $f(1)$ multiplied by the coefficient of γ^2 :

$$\begin{aligned}
 f(1)DE &= (2DE - (D^2 + E^2 + F^2))DE \\
 &= -(F^2 + (D - E)^2)DE \\
 &\leq 0.
 \end{aligned}$$

This means that one and only root of $f(y)$ is greater than or equal to 1. Let us denote this root by y_0 . By substituting y_0 into (14), we can solve it as a simultaneous quadratic equation on p_1 and q_1 , and know that (p_1, q_1) has two solutions in the form of (p_0, q_0) and $(-p_0, -q_0)$, which are symmetrical to the gradient space origin.

References

- Ikeuchi, K. Numerical shape from shading and occluding contours in a single view. MIT AI Memo 566, 1980. (a)
- Ikeuchi, K. Shape from regular patterns (an example of constraint propagation in vision). MIT AI Memo 567, 1980. (b)
- Kanade, T. Recovery of the 3-dimensional shape of an object from a single view. *Artificial Intelligence*, 1981, **17**, 409-460.
- Kender, J. R. Shape from texture. Doctoral dissertation, Carnegie-Mellon University, 1980.
- Mackworth, A. K. Interpreting pictures of polyhedral scenes. *Artificial Intelligence*, 1973, **4**, 121-137.
- Stevens, K. A. Surface perception from local analysis of texture and contour. MIT AI Memo 512, 1980.
- Woodham, R.J. A cooperative algorithm for determining surface orientation from a single view. Proceedings, 5th International Joint Conference on Artificial Intelligence, 1977, 635-641.

