

Recognizing Partially Occluded Parts

JERRY L. TURNER, STUDENT MEMBER, IEEE, TREVOR N. MUDGE, SENIOR MEMBER, IEEE,
AND RICHARD A. VOLZ, MEMBER, IEEE

Abstract—The problem of recognizing an object from a partially occluded boundary image is considered, and the concept of saliency of a boundary segment is introduced. Saliency measures the extent to which the boundary segment distinguishes the object to which it belongs from other objects which might be present. An algorithm is presented which optimally determines the saliency of boundary segments of one object with respect to those of a set of other objects. An efficient template matching algorithm using templates weighted by boundary segment saliency is then presented and employed to recognize partially occluded parts. The results of these experiments illustrate the effectiveness of the new technique.

Index Terms—Hough transform, least squares, occluded parts, saliency, weighted template matching.

I. INTRODUCTION

THE problem of recognizing partially occluded parts is of considerable interest in the field of industrial automation. While it is possible to employ shakers, conveyor belts, and custom machinery to separate, palletize, or otherwise prearrange the parts for easy recognition, a vision system which can recognize the parts even though they may be partially occluded and in random positions is much more flexible. Attempting to match a model, or template, of the boundary of the part with the boundary of the image has been suggested for this problem [12], [13]. If a position of the template can be found for which a large number of its boundary pixels matches those of the image, it may be concluded that the part has been found. Unfortunately, in its basic form, this procedure requires an excessive amount of computation and can produce a number of false matches resulting in unreliable recognition. This paper presents an efficient method of template matching which is based on matching subtemplates, and by which the subtemplates are optimally weighted to distinguish those of the object being sought from those derived from other objects which might be present.

In a typical industrial scene, the parts may be intermixed, partially occluded, and of unknown pose (position and orientation); however, the types of parts that will be present are almost always known *a priori*. We assume that there is available a database describing the geometry of all the parts that could conceivably appear in the scene. This database could be con-

Manuscript received December 30, 1983; revised July 6, 1984. This work was supported in part by the Zimmer Foundation, and in part by the Air Force Office of Scientific Research under Contract F49620-82-C-0089.

The authors are with the Department of Electrical Engineering and Computer Science, Robot Systems Division, Center for Robotics and Integrated Manufacturing, University of Michigan, Ann Arbor, MI 48109.

structed during a “learning” phase, or it could be constructed by a computer aided design (CAD) system as a byproduct of the design process. The latter method is likely to predominate in the future as CAD systems become more sophisticated and capable of being integrated into a “factory-of-the-future” environment. Indeed, work in this direction has already begun; for example, silhouettes and grasppoints for parts can automatically be extracted from CAD data and organized into a suitable database for part handling [4], [5], [18], [19].

From the database of part geometries, boundary templates for every stable position of every part are determined, and a set of subtemplates which most differentiate the parts is identified. The subtemplates in this set are called *salient* features. Salient features depend on the set of parts being considered. To illustrate this point, consider Fig. 1. Fig. 1(a) shows the salient features when parts *A* and *B* comprise the set of parts, while Fig. 1(b) shows the salient features when parts *A* and *C* comprise the set of parts (salient features are shown as heavy lines). Those features of *A*’s boundary which most distinguish it are clearly dependent on the set of parts, i.e., on the context in which *A* could potentially be viewed. To make this concept more precise, a measure of *saliency* is introduced to denote the distinctiveness of a subtemplate. The saliency of a subtemplate corresponds to the optimal (in the least-squares sense) weighting mentioned above and is defined on the closed interval [0, 1] (see Section IV). By giving saliency a continuous range, the pose of a part that is partially occluded can be accurately determined if the visible portion of the part matches some set of subtemplates with enough combined saliency.

Preliminary versions of some of the concepts used in the approach to partially occluded part recognition developed in this paper were first reported in [16] and [17]. For the most part, the approach is most suitable for fairly flat parts of known scale. However, work in progress suggests that the technique is more general and may be extended to three dimensions, in which case the concept of saliency still applies, but the subtemplates are no longer boundary segments, rather surface subregions. Even with the above restrictions, the technique applies to a wide range of industrial applications.

This paper is organized as follows. Section II gives the background for the techniques to be employed, including a brief summary of the Hough transform which motivated portions of this work. Section III develops an efficient template matching scheme based on an extension of the Hough transform that uses subtemplates. Section IV addresses the problem of determining the saliency of subtemplates for a set of parts. Section V presents some experimental results which illustrate the ef-

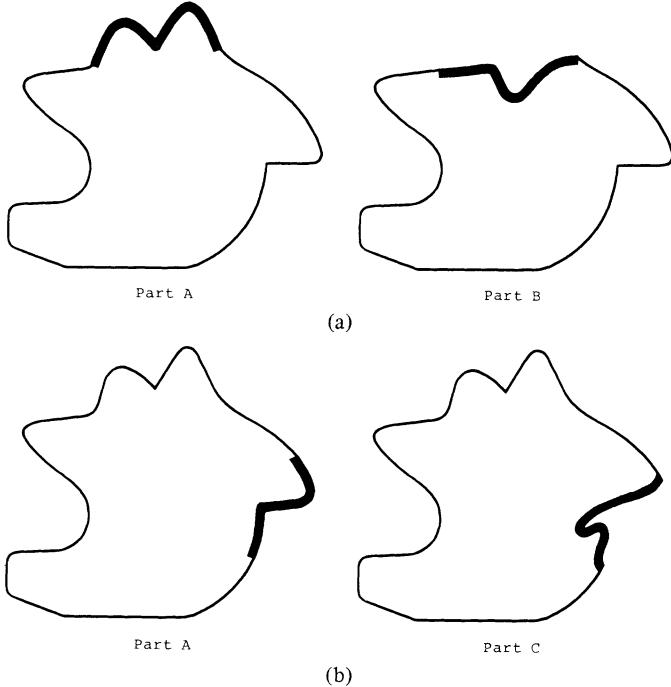


Fig. 1. Context dependency of salient features.

fectiveness of using saliency for recognizing partially occluded parts. Section VI concludes with a discussion on further directions for research, in particular, proposals for improving recognition times.

II. BACKGROUND

The approach to recognizing partially occluded parts developed in this paper works with the shape of their boundaries. In the context of the intended application—recognizing and locating fairly flat machined parts—this choice retains that aspect of the parts which conveys a significant amount of information useful for recognition; at the same time, it allows us to work with a relatively compact boundary representation of the scene.

The scene in which the parts appear will be referred to as the application scene. It is digitized into a two-dimensional array of pixels, $I(1 \leq x \leq M, 1 \leq y \leq N: x \text{ and } y \text{ integer})$, which can take on values from a set of gray levels. The boundary representation of the parts in the application scene is called the boundary image. Although its machine representation is usually in a compact form, such as a linked list, it can be thought of as a function $B(x, y)$ which is defined on the same domain as I and which has the value one at boundary pixels, and zero elsewhere. Boundary pixels and their slope angles¹ are extracted from the application scene so that only the strongest boundary pixels which maintain continuity remain [7]. The output is the boundary image. It is the boundary image against which the templates are matched to locate the part in the application scene. For this discussion, a template $T(x, y)$ can also be thought of as a binary valued function defined similarly to B , but with a smaller domain ($1 \leq x \leq m, 1 \leq y \leq n: m \ll M$ and $n \ll N$). Furthermore, template pixels will be considered

¹A pixel's slope angle is the angle of the tangent to the boundary at that pixel.

to be those pixels that are, strictly speaking, boundary pixels of the template, i.e., where $T(x, y)$ has the value of one.

As mentioned earlier, template matching is performed using an extension of the Hough transform. The following subsection gives a brief summary of the Hough transform and its relationship to normal template matching.

A. The Hough Transform and Template Matching

The Hough transform was originally designed to identify colinear subsets of points within larger sets of points [8]. Subsequently, a number of significant improvements has been made to the original transform. In particular, Merlin and Farber generalized it to identify arbitrary sets of points [10]. We begin by describing how this version of the transform can be used to identify the position of “best match” of a template $T(x, y)$ to a boundary image $B(x, y)$.

Assume that $B(x, y)$ is produced from the application scene in the manner outlined above. Associated with each pixel in B is an accumulator whose purpose is to record template matches. This two-dimensional array of accumulators is assumed to be initially zero. A reference point is selected that is fixed with respect to T ; we used T ’s centroid because it was useful in a related project where a robot was required to grasp the part, but any appropriate reference point may be used. A template $T^*(x, y)$ is formed by rotating T 180° about its centroid. The centroid of T^* is placed at a boundary pixel in B and every accumulator that coincides with a pixel of T^* is incremented by one (see Fig. 2). This procedure is repeated for each boundary pixel in B . When all the boundary pixels have been visited, the accumulator with the largest value identifies the position of the centroid of template T that gives the best match.

Sklansky demonstrated in [15] that the function produced in the accumulator array by the above procedure is the convolution of B with T^* which is simply the cross correlation or template matching T with B . Since the Hough transform is equivalent to template matching, it performs the function of a matched filter and, therefore, is the optimum filter when B is corrupted by additive white Gaussian noise under a wide variety of criteria: signal-to-noise, likelihood ratio, and inverse probability [6, ch. 7]. However, none of these criteria implies that it is optimal for recognizing one shape in the presence of others (see the discussion on normalized correlation in Rosenfeld and Kak [14]), suggesting that template matching is by no means sufficient for recognizing partially occluded parts.

In [15], it is also suggested that the Hough transform could be improved if local properties of the boundary pixels were used to restrict the set of accumulators that are incremented. Less computation would be needed and fewer false matches would occur, i.e., the “ Q ” of the filter would be increased. Ballard [2] presents an improved version of the Hough transform in which the slope of the boundary pixels is used to restrict the set of accumulators that are incremented. This is achieved by performing the Hough transform in the following way. Create a set of vectors, $V = \{v_i, i = 1, \dots, |T|\}$, from each pixel t_i on the boundary of T to T ’s centroid. When the vectors in V are all originated at the same point, they trace

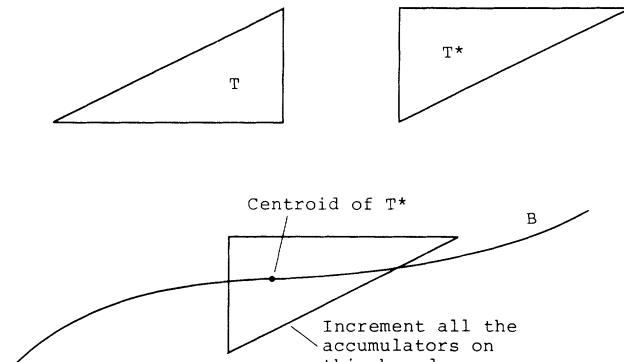
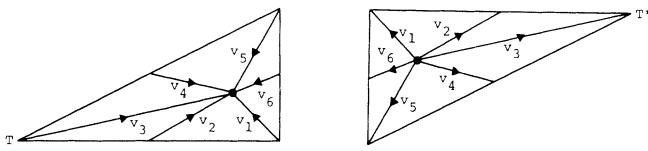


Fig. 2. Hough transform.



$$V = \{v_1, v_2, v_3, v_4, v_5, v_6\}$$

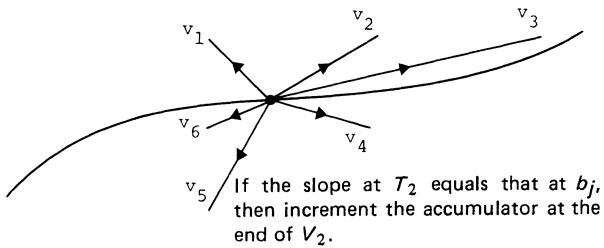
Fig. 3. Creating T^* from V .

Fig. 4. Slope restricted Hough transform.

the rotated template T^* (see Fig. 3). Therefore, the above procedure of Merlin and Farber can be performed equivalently by originating the vectors at a boundary image pixel and incrementing the accumulators pointed to by each of the vectors; this step is then repeated for every pixel on the boundary image. The equivalence between the Hough transform and template matching can be seen if it is considered that originating a vector v_i at a boundary pixel b_j is, in effect, proposing b_j as a match for the template pixel t_i , and voting for that match by incrementing the accumulator pointed to by v_i . Restrictions can be applied by simply using subsets of V at each boundary point. In Ballard's version of the Hough transform, the slope of the boundary pixel is used to select only those v_i whose corresponding t_i have the same slope² (see Fig. 4). This approach has already shown some success in an industrial application [1]. In the next section, the idea of local restriction is taken one stage further by requiring that a pixel's neighboring pixels satisfy a certain configuration or subtemplate.

The computational complexity of the Hough transform (and hence template matching) can increase rapidly if it is necessary

²In the original description of the algorithm, the perpendicular to the slope—the gradient—was used.

to deal with variations in orientation and scale. For example, the above procedure for the Hough transform assumes a fixed orientation of T^* (or T). To account for variations in orientation, the complete procedure must be repeated for every orientation to be distinguished. In other words, if it is required to distinguish between orientations that are 1° apart, the procedure must be repeated 360 times, resulting in 360 accumulator arrays. The best match would then be identified by the accumulator with the largest value in all of the 360 arrays. A similar complexity arises if the scale of the parts is not known in advance—all sizes of T to be distinguished must be tried.

The next section presents a technique based on an extension of the Hough transform which uses subtemplates. This technique reduces the complexity that arises from the need to consider variations in orientation; however, variations in scale are not dealt with. In addition, and more importantly, subtemplates provide convenient boundary features that can receive salient weightings (see Fig. 1).

III. SUBTEMPLATE MATCHING

Let T be a template of the boundary of a part. A set τ of overlapping boundary subtemplates $\{\tau_i\}$ are created from template T . Each subtemplate τ_i has an associated vector v_i that points to the location of the centroid of T . The Hough transform can now be modified by requiring that a segment of the boundary image centered on a particular pixel b_j match τ_i before v_i can be originated at b_j . In addition, one can introduce a coefficient of match between τ_i and the boundary segment (see below) and use this to allow the accumulator pointed to by v_i to be incremented by a fractional value. In our experiments, each τ_i was 20 pixels in length and $|T|/2$ subtemplates were created for each template T , i.e., each template pixel was part of 10 subtemplates. The choice of the length and the number of subtemplates is application dependent.

In the process of matching the subtemplates with the boundary, it is important to place the subtemplate in the correct orientation. This can be achieved efficiently if the subtemplate and the boundary image segments are represented in their *intrinsic* coordinate system in which the angle of slope θ and the arc length along the boundary s act as coordinates for pixels [3], [12], [13]. This will be referred to as the θ — s representation of a boundary. The θ — s representation allows subtemplates and boundary image segments to be characterized by functions of the form $\theta(s)$. An important property of these intrinsic functions is that a change in orientation θ_c of a boundary in x — y space corresponds to simply adding θ_c to $\theta(s)$ in θ — s space. The slope θ is determined, as mentioned before, during boundary extraction, while s is determined by measuring the distance in pixels along the boundary image (see Section II). The values of the arc length s for a boundary are dependent to some extent on the orientation of the boundary because of the different amounts of quantization noise introduced into the boundary image at different orientations. This noise also causes an accumulated error in the determination of the true arc length. In our experiments, the effects of quantization noise were minimized by fitting a simple third-order interpolating polynomial to the points on the curve. The s values were extracted from the interpolated boundary.

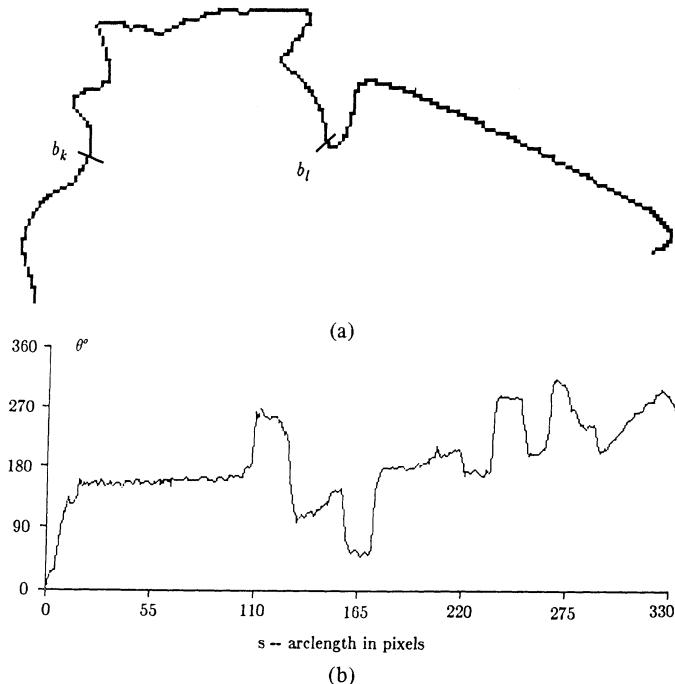
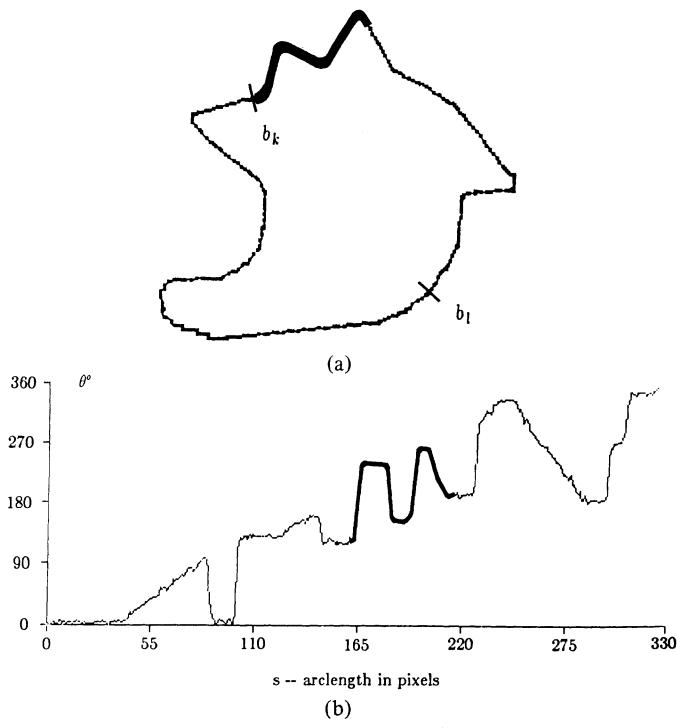
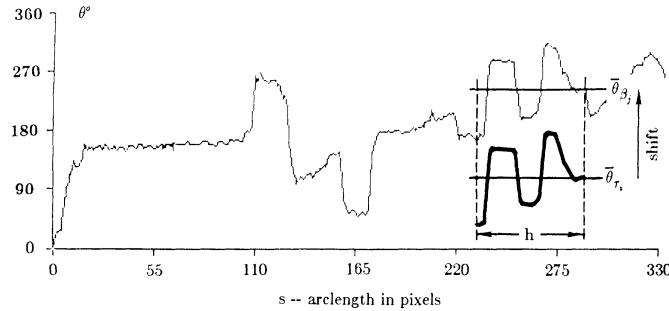
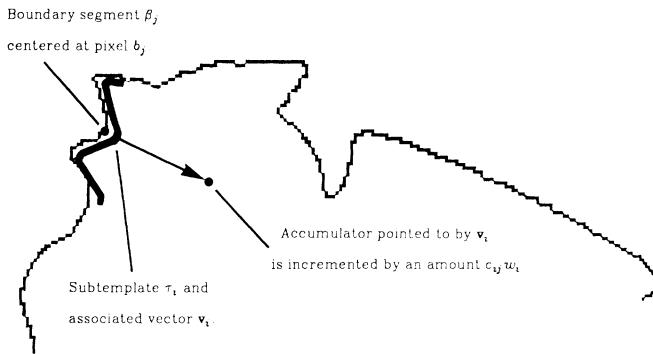
Fig. 5. A boundary image and its $\theta - s$ representation.Fig. 6. A template and its $\theta - s$ representation.

Fig. 5(a) shows a boundary image containing a partially occluded part between pixels b_k and b_l . Fig. 5(b) displays its $\theta - s$ representation when the boundary is traced counter-clockwise. Fig. 6(a) shows the complete template of the partially occluded part in Fig. 5; a subtemplate is shown with a heavy line. Fig. 6(b) shows the $\theta - s$ representation of the template.

Matching starts by choosing a subtemplate τ_i from the template boundary and attempting to match it against an equal length segment of the boundary image β_j centered at pixel b_j . The attempted match is performed in $\theta - s$ space at regularly

spaced boundary image pixels, $b_j, j = 1, \dots, |\beta|$, where $|\beta|$ is the number of boundary segments. In our experiments, the spacing was every other boundary image pixel; as with subtemplate length and number, spacing is application dependent. Strictly speaking, since the matching is performed in $\theta - s$ space, the $\theta - s$ graph of τ_i , $\theta_{\tau_i}(s)$ is matched against the $\theta - s$ graph of β_j , $\theta_{\beta_j}(s)$. Conceptually, this can be thought of as moving $\theta_{\tau_i}(s)$ along the s axis direction in the $\theta - s$ graph of the boundary image $\theta_B(s)$ and performing a match at regularly spaced points (see Fig. 7). The matching operation is repeated for a “representative” set of subtemplates for the part. In the

Fig. 7. Matching in θ - s space.Fig. 8. Matching in x - y space.

experiments reported in Section V, the representative set is taken to be all the subtemplates. However, preliminary results mentioned in the Conclusion suggest various strategies for selecting this representative set that significantly improve recognition times.

As noted earlier, it is important to get the subtemplates in the correct orientation before matching. This was done by taking a least-squares fit in θ - s space of the subtemplate to the boundary segment. Consider the quantity

$$\gamma_{ij} = \sum_{p=1}^h [\theta_{\beta_j}(s_p) - (\theta_{\tau_i}(s_p) + \theta_{ij})]^2$$

where h is the number of pixels in τ_i and β_j (recall, in our experiments $h = 20$), and θ_{ij} is a single variable used to adjust the fit. Let θ_{ij}^* represent the value of θ_{ij} at which γ_{ij} takes on a minimum. It is easy to show that θ_{ij}^* is equal to the difference between the average slope of the boundary segment and the average slope of the subtemplate, i.e., $\theta_{ij}^* = \bar{\theta}_{\beta_j} - \bar{\theta}_{\tau_i}$. Thus, the minimum value of γ_{ij} is achieved by simply shifting the slopes of the subtemplate points by θ_{ij}^* before they are subtracted from the slopes of the boundary segment points (see Fig. 7).

The shift in θ corresponds to an angular rotation of the subtemplate to the average orientation of the boundary image segment. Fig. 8 shows the corresponding match between the subtemplate and the boundary image segment in x - y space.

A measure of the mismatch between the rotated version of subtemplate τ_i and the boundary image segment β_j is given by the minimum value of γ_{ij} . Denote this value as γ_{ij}^* , then

$$\gamma_{ij}^* = \sum_{p=1}^h [\theta_{\beta_j}(s_p) - (\theta_{\tau_i}(s_p) + \theta_{ij}^*)]^2.$$

Using γ_{ij}^* , we can define a *matching coefficient* which gives a measure of the match between a subtemplate and a boundary image segment. Let

$$c_{ij} = \frac{1}{1 + \gamma_{ij}^*}.$$

The matching coefficient c_{ij} is used to adjust the contribution added to the accumulator during subtemplate matching, and it is defined in the above manner so that values close to 0 imply a very poor match and values close to 1 imply a very good match ($0 < c_{ij} \leq 1$).

In order to locate the potential centroid for template T (and hence the part that T represents), the vector v_i , associated with subtemplate τ_i , is rotated by the angle θ_{ij}^* and is originated at the center of the boundary image segment, as shown in Fig. 8. The accumulator pointed to by v_i is then incremented by the quantity $c_{ij}w_i$ where w_i is the saliency of τ_i mentioned in the Introduction (see also, the next section). The result of matching a subtemplate τ_i with a boundary segment β_j is thus characterized by two factors: the similarity of τ_i and β_j (measured by c_{ij}), and the significance of finding a match for τ_i in recognizing the part from which τ_i is derived (measured by w_i).

In addition, the angular difference θ_{ij}^* is appended to a linked list associated with the same pixel as the accumulator pointed to by v_i . At the end of all the subtemplate matching, the list of angles associated with the pixel which is judged to be at T 's centroid is averaged to estimate the orientation of T . As before, T 's centroid is judged to be the pixel whose associated accumulator contains the greatest number.

The above method for subtemplate matching provides a simple technique for performing a least-squares fit. In addition,

the method does not rely on long connected boundary images; if the boundary image is fragmented, the subtemplates are matched to each fragment. However, the value chosen for h (the number of pixels in a subtemplate) does determine a lower limit on the size of useful fragments.

IV. DETERMINING THE SALIENCY OF SUBTEMPLATES

As mentioned in Section I, by properly weighting subtemplates, emphasis can be given to those features which distinguish an object—the salient features. Thus, partially occluded parts can be recognized if features with enough saliency are exposed. This section develops an algorithm by which salient weights can be automatically assigned to the subtemplates.

The algorithm works off line with the database of part geometries, specifically the subtemplates derived from the part geometries. It has the advantage that a new part geometry can be introduced into the database and new weights can readily be assigned to all the subtemplates affected by the addition of the new part. The algorithm optimally adjusts the weights of subtemplates so that different templates correlate poorly with one another.

First, for simplicity, only two templates A and B will be considered. They may be derived from different parts or from different stable positions of the same part. Assume that template A consists of $|\tau|$ subtemplates τ_i and that each subtemplate is assigned a positive weight w_i . Assume, furthermore, that each τ_i is matched with each of the $|\beta|$ boundary segments β_j of B . Although A and B are both templates, it is convenient for this explanation to consider A to be a template and B a boundary image. If subtemplate τ_i of A matches a boundary segment β_j of B , a symbolic record of the match $c_{ij}w_i$ is stored at the accumulator pointed to by v_i . Let s be a unique index into the accumulator array defined by $s = (N - 1)^* \text{row_index} + \text{column_index}$ for an $M \times N$ array. The symbolic value stored at accumulator s is a variable w_i and a numerical value c_{ij} , the matching coefficient, derived through the subtemplate matching process described in the previous section. Fig. 9 illustrates the contributions that accumulator s might receive from matching subtemplates with boundary segments. Since template A has $|\tau|$ subtemplates, and B has $|\beta|$ subtemplates, $|\tau| |\beta|$ matches are attempted.

After all the subtemplate matching is completed, accumulator s contains a linear function in $|\tau|$ variables w_i , $i = 1, \dots, |\tau|$. Let this function be a_s , then

$$a_s = \sum_{i=1}^{|\tau|} \sum_{j=1}^{|\beta|} \delta_{ij}^s c_{ij} w_i$$

where δ_{ij}^s is a selector function that takes the value 1 if s receives a contribution from a match between subtemplate τ_i and boundary segment β_j ; δ_{ij}^s takes the value 0 otherwise. The above equation can be written in matrix form as

$$\mathbf{a} = \mathbf{Dw}$$

where the vector \mathbf{w} has the elements w_i , and \mathbf{D} is an $(M \times N) \times |\tau|$

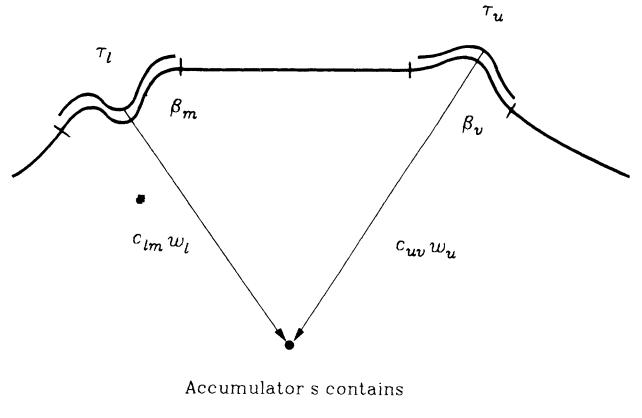


Fig. 9. Computing saliency.

$|\tau|$ matrix with nonnegative elements given by

$$d_{si} = \sum_{j=1}^{|\beta|} \delta_{ij}^s c_{ij}.$$

In practice, D is very sparse.

As noted above, the objective of the weighting algorithm is to adjust the weights w_i so as to minimize the correlation of template A with boundary B . The approach taken is to minimize the sum of the squares of all of the a_s 's. This results in the subtemplates that are most unlike others—the most salient ones—receiving the greatest weights. The problem may be stated as follows:

$$\text{Minimize: } \sum_{s=1}^{M \times N} a_s^2 = \mathbf{a}^t \mathbf{a} = \mathbf{w}^t \mathbf{D}^t \mathbf{D} \mathbf{w}$$

$$\text{Subject to: } \mathbf{1}^t \mathbf{w} = 1 \quad \text{and } \mathbf{w} \geq \mathbf{0}$$

where $\mathbf{1}$ is a $|\tau|$ -dimensional vector of 1's and $\mathbf{0}$ is a $|\tau|$ -dimensional vector of 0's. Given \mathbf{w} , a nonzero vector of weights, $k\mathbf{w}$ ($k > 0$) will produce the same ordering in the accumulator array and hence the same solution to the part location problem. Accordingly, the constraint $\mathbf{1}^t \mathbf{w} = 1$ is used to provide a normalization of the weights and to avoid the degenerate solution $\mathbf{w} = \mathbf{0}$.

The constraint $\mathbf{1}^t \mathbf{w} = 1$ is a hyperplane normal to the vector $\mathbf{1}$ and at a distance of $|\tau|^{-1/2}$ from the origin. The surfaces of constant value of the function $\mathbf{w}^t \mathbf{D}^t \mathbf{D} \mathbf{w}$ are hyperellipsoids. The constrained solution occurs for the value which makes the corresponding hyperellipse tangent to the constraining hyperplane. When the ratio of maximum-to-minimum eigenvalues of $\mathbf{D}^t \mathbf{D}$ is large (as is often the case in this application) and the major axis of the hyperellipsoid is skewed so as to be nearly parallel to the hyperplane, large negative and large positive weights can occur. Fig. 10 shows a two-dimensional example in which positive and negative weights result— S is the solution point. In terms of the matched filter interpretation of Section II, this amounts to creating a very high “ Q ” filter for the part. In other words, using the weights to recognize the part in an application scene results in the contents of the array of accumulators forming a sharp peak at the location of the part's centroid, if it is present. This narrow bandwidth is obtained, in part, by the

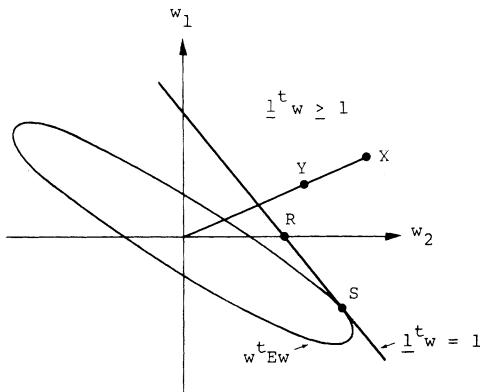


Fig. 10. Minimal solutions in two dimensions.

weighted contributions to the accumulators cancelling one another everywhere except close to the centroid. If positive and negative weights are allowed, there is a danger of creating neighboring positive and negative peaks which should cancel but do not due to noise and quantization (the weights are determined from noiseless images derived from the database of part geometries). The constraint $w \geq 0$ "detunes" the filter and prevents this from occurring. In terms of Fig. 10, the solution point is moved to R .

It is convenient to restate the minimization problem in a slightly different form by letting E denote the product $D^t D$. E is then symmetric and, although the D matrix is typically an extremely large sparse $(M \times N) \times |\tau|$ matrix (consider $M = N = 256$), E is a much smaller $|\tau| \times |\tau|$ matrix. The above minimization problem becomes

$$\text{Minimize: } \sum_{s=1}^{M \times N} a_s^2 = \mathbf{a}^t \mathbf{a} = \mathbf{w}^t E \mathbf{w}$$

$$\text{Subject to: } \underline{1}^t \mathbf{w} = 1 \quad \text{and } \mathbf{w} \geq 0.$$

This can be transformed to a standard least-squares problem and solved accordingly (see [9]). Details can be found in the Appendix.

In the above discussion, weights were chosen for template A , so that its correlation with the boundary of B would be minimized everywhere. More than two parts can be taken into consideration. For example, if parts A , B , and C are to appear in a scene and it is required to distinguish part A , then it is necessary to minimize both $\sum_{s \in \{B\}} a_s^2$ and $\sum_{s \in \{C\}} a_s^2$, where the sets $\{B\}$ and $\{C\}$ denote the accumulators that receive contributions from parts B and C , respectively. Assume that $\{B\}$ and $\{C\}$ are disjoint (*disjoint accumulators* assumption), then the correlation matrix for A correlated with both B and C , E_{BC}^A , is related to the E matrices for A correlated with B and A correlated with C by the relation

$$\mathbf{w}^t E_{BC}^A \mathbf{w} = \sum_s a_s^2 = \sum_{s \in \{B\}} a_s^2 + \sum_{s \in \{C\}} a_s^2 = \mathbf{w}^t (E_B^A + E_C^A) \mathbf{w}.$$

If it is required to recognize the correct orientation of A , then the autocorrelation of A should also be minimized. (The constraint $\underline{1}^t \mathbf{w} = 1$ ensures that the accumulator corresponding to the correct match always has a value of one.) The associated autocorrelation matrix is denoted by E_A^A . The correlation

matrix for identifying A in the presence of B and C is, given the disjoint accumulators assumption, computed as follows:

$$E_{ABC}^A = E_A^A + E_B^A + E_C^A.$$

The correlation matrix E_{ABC}^A can then be used in the weighting algorithm to obtain values for the weights w_i of the subtemplates of A .

The solution to the weights gives precision to the term "saliency." The weights may be interpreted as the saliency of the subtemplate of A with respect to parts B and C . Furthermore, the weights distinguish, in a precise least-squares sense, each subtemplate of A from B , C , and other subtemplates of A , allowing the recognition of A in the presence of B and/or C even if A is partially occluded by B and/or C . The degree of occlusion possible depends on how much saliency is visible.

In summary, the pairwise correlation matrices E_i^i are determined from the database of part geometries by computing the D matrices using the subtemplate matching process of Section III. They can then be added according to the particular part(s) required to be recognized. An $E_{\text{set of parts}}^i$ matrix will result for each part i to be recognized. A separate set of weights can then be computed, as outlined above, for use in the recognition of each part. This can amount to a considerable amount of computation. However, in the type of environment envisioned for this recognition algorithm—a CAD-driven manufacturing cell, for instance—the computation is performed as part of a set-up phase and reused over a much longer period of time.

Adding a new part to the database amounts to finding the weighting matrix E_j^{new} between the new part and every other member j of the database and adding them to the $E_{\text{set of parts}}^i$, then recomputing the sets of salient weights. Removing an old part from consideration can be done in a similar fashion by first subtracting E_j^{old} from the $E_{\text{set of parts}}^i$.

Finally, it should be noted that the $E_{\text{sets of parts}}^i$ matrices are constructed on the disjoint accumulator assumption. This is equivalent to ignoring the possibility of accidental contributions to an accumulator that can occur in practice and result in faulty recognition. In fact, this rarely occurs while the assumption greatly simplifies the calculation of saliency.

V. EXPERIMENTAL RESULTS

Fig. 11 illustrates the results of an experiment to test the recognition technique. Fig. 11(a) shows three keys, two of which (key A and key C) are very similar in appearance. In Fig. 11(b), key B has been placed on key A , occluding key A ; key C remains unoccluded. An attempt was made to locate key A in the application scene depicted in Fig. 11(b) using the slope restricted Hough transform [2]. Fig. 11(c) illustrates how the location of key A in Fig. 11(b) was incorrectly determined to be at the location of key C . Table I shows the largest peaks in the accumulator array and their location after the above experiment. They are ranked according to the value of their contents—the value shown is actually the true value normalized by the total weight of all the subtemplates from key A scaled by 1000. The accumulators form a 256×256 array, and the accumulator locations shown are of the form *abscissa, ordinate*.

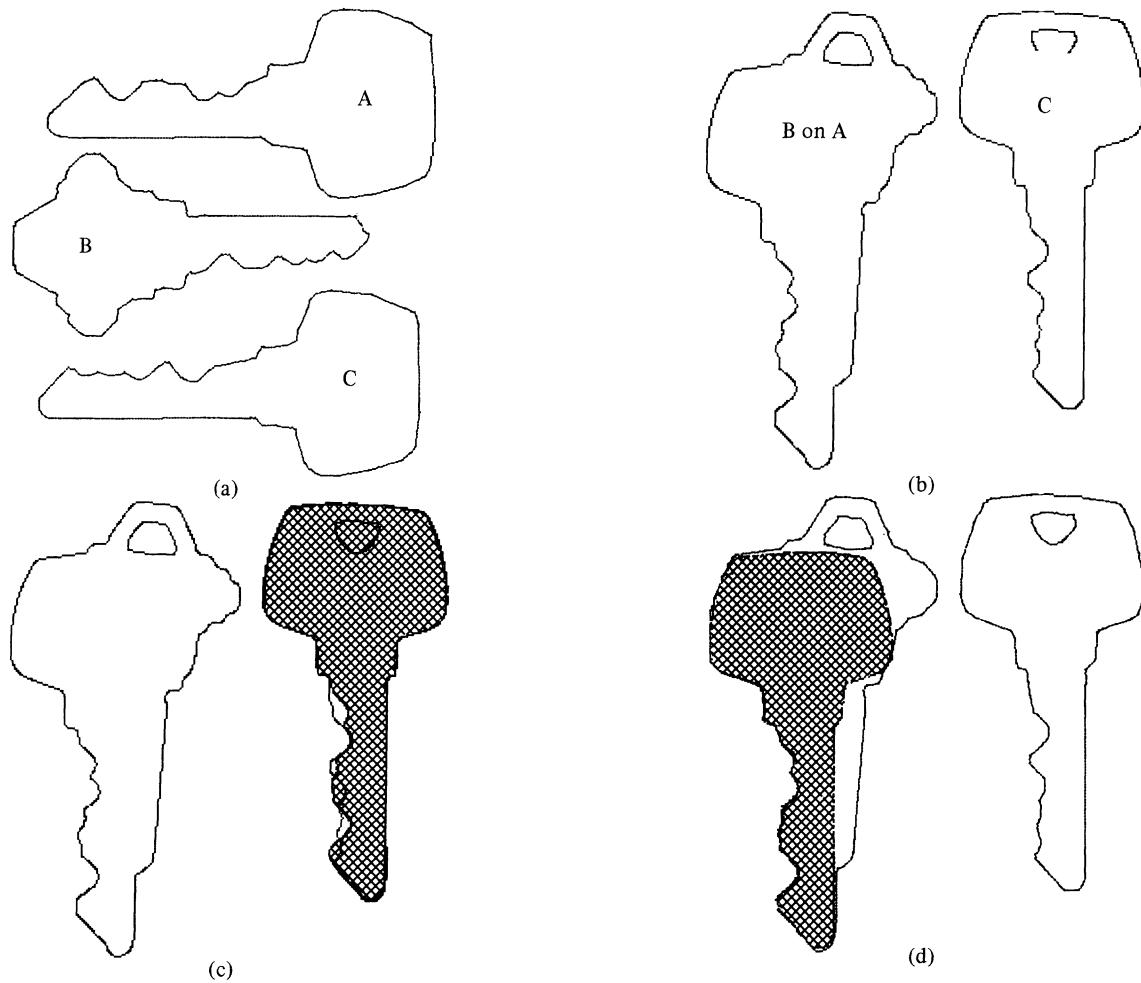


Fig. 11. The three keys experiment.

TABLE I
ACCUMULATOR ARRAY FOR SLOPE RESTRICTED HOUGH TRANSFORM

Rank	Value	Accumulator Coordinates
100	101	(188,145)
99	82	(188,143)
98	82	(188,144)
97	82	(188,146)
96	74	(188,133)
95	72	(188,128)
•	•	
36	46	(188,165)
35	44	(56,118) ●
34	44	(187,144)

The peak that corresponds to the centroid of key *A*—shown next to a bullet—is at (56, 118) and is ranked 65 places below the maximum which is the centroid of key *C* at (188, 145).

Next, an attempt was made to locate key *A* in the application scene depicted in Fig. 11(b) using the weighted subtemplate method. This attempt was successful and Fig. 11(d) illustrates how key *A* was correctly located. Table II shows the contents of the accumulators; the maximum peak is at the location of *A*'s centroid, (56, 118).

The slope restricted Hough transform failed in this experi-

TABLE II
ACCUMULATOR ARRAY FOR THE WEIGHTED SUBTEMPLATE METHOD

Rank	Value	Accumulator Coordinates
100	89	(56,118)
99	84	(57,118)
98	80	(56,117)
97	65	(57,117)
96	62	(58,118)
95	59	(187,147)
94	50	(59,118)
93	42	(55,118)
92	40	(58,116)

ment because it placed equal weight on all template points that were matched. There was a considerable number of points of equal slope on the template, especially on the straight portions of the boundary; thus, a large number of false matches occurred between the points of the template and the boundary image, resulting in false peaks in the accumulator array and the subsequent incorrect location of the part. Distinct weights cannot be meaningfully assigned to points because they are all alike. In contrast, distinct weights can be assigned to subtemplates because they have enough structure to be distinguishable.

A more complex experiment was performed with a set of nine parts. Fig. 12 shows templates derived from the parts. Fig. 13 illustrates the weights associated with the subtemplates

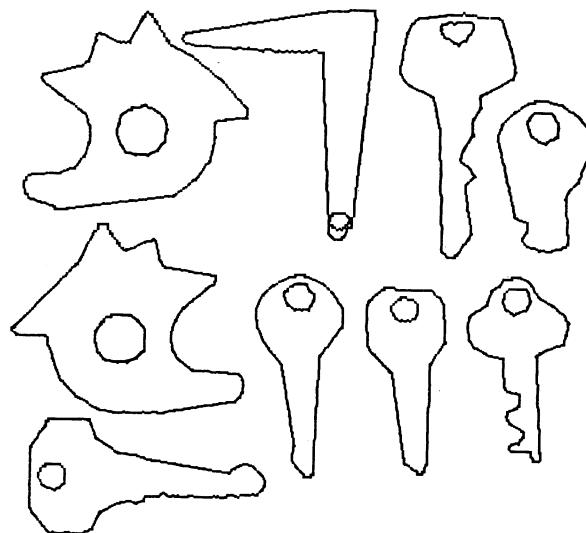


Fig. 12. Templates derived from a set of nine parts.

of three representative parts from the set of nine parts. At left, in the figure, the weight of a subtemplate is represented graphically by a vertical line attached to the center of the subtemplate; the length of the line is proportional to the weight of the subtemplate. At right, a histogram of the weights is shown. Fig. 14(a) shows the boundary extracted from an application scene in which the parts were piled in an arbitrary manner. Fig. 14(b) shows the result of trying to locate the shaded part using the slope restricted Hough transform. Its position and orientation were incorrectly identified. Again, since the restricted Hough transform placed equal weight on each point of the template, there was a considerable number of false matches and a false peak in the accumulation array with a consequent false location of the part. Fig. 14(c) shows the result of using the weighted subtemplate method to locate the part.

In other experiments, with the pile in Fig. 14(a), it was found that typical times for locating individual parts using the slope restricted Hough transform were in the range 4-4.5 CPU min using a VAX 11/780. In several cases, an incorrect match was made. The typical times for the weighted subtemplate method were in the range 3-3.5 CPU min. In contrast, no incorrect matches were made. Preliminary results mentioned in the Conclusion suggest recognition time can be considerably shortened if the recognition process is based primarily on the more salient subtemplates. Both algorithms were written in C. Not much was attempted in the way of optimizing the implementation of the algorithms, apart from representing the accumulator arrays—very sparse arrays—as hash tables. This significantly reduced storage requirements and made possible a speedup in the search for the location of the maximum peak. The time to calculate the weights for a part was approximately 20 CPU min.

It is possible to get false matches using the weighted subtemplate method. They typically occur if the part to be located is almost totally occluded, or if the exposed portion does not distinguish the part from other parts. The latter situation is not unlike the confusion a human observer would experience when presented with views of occluded parts in which two or more parts had all their distinguishing features hidden from view.

VI. CONCLUSION

A method for recognizing partially occluded parts has been presented. There are two principal components to the method. The first is a subtemplate matching scheme derived from earlier Hough transform techniques. This provides the features (subtemplates) for the second component which is a weighting scheme. The weighting scheme derives weights of the subtemplates that are a measure of their distinctiveness or saliency. Since the saliency of a part's subtemplate is entirely dependent on the parts with which it can appear (refer again to Fig. 1), deriving the weights requires *a priori* knowledge about the set of parts that can appear in an application scene. The saliency weights are an embodiment of this knowledge and account for the added recognition capability of the method over earlier template matching methods. It is this important reliance on *a priori* information that distinguishes the weighted subtemplate method from other template matching methods that use weighting schemes (see [2, sect. 5] for a good summary of these methods). However, the added recognition capability is bought at the expense of having to know ahead of time what the application scene may contain. Clearly, there are applications where this may not be possible. However, in our principal area of concern—recognizing manufactured parts—the set of parts likely to appear can usually be determined.

The time complexity of the present method is $O(h|\tau| |\beta|)$. There are several techniques which may reduce recognition times that warrant further investigation. An obvious one, when several parts are being sought, is to eliminate segments of the boundary as soon as they are determined to belong to a part.

Another is to implement a coarse to fine search in matching the subtemplates to the boundary segments. The subtemplate matching is performed between subtemplates and boundary images that have been resampled at coarse intervals in $\theta - s$ space, i.e., less frequently than every pixel. A coarsely sampled subtemplate is matched to a coarsely sampled boundary segment and, if the match is poor, the matching coefficient is set to zero. On the other hand, if the match is satisfactory, matching is attempted again but with subtemplates and boundaries more finely sampled. A sequence of increasingly more finely resampled templates and boundaries is used in matching and,

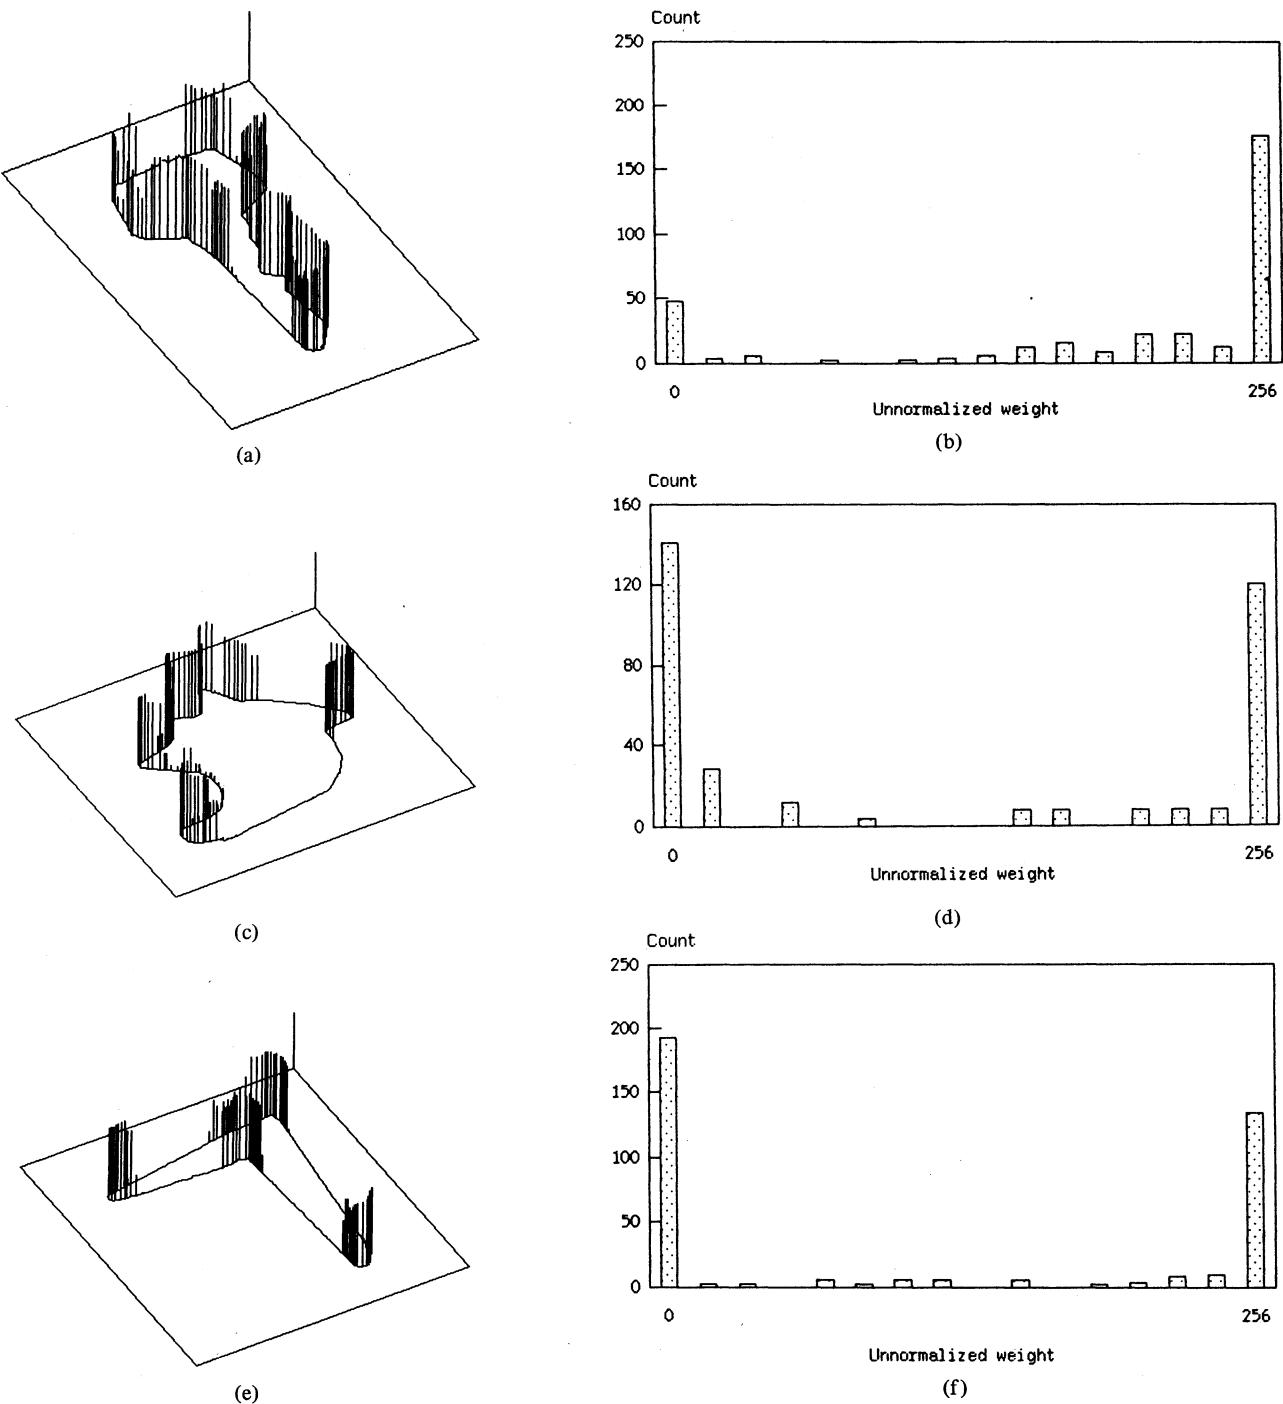


Fig. 13. Weights associated with representative parts.

if they continue to show satisfactory matches, a situation is reached where the normal subtemplate is matched to the normal boundary, and the matching coefficient can be determined. The potential for speedup lies in the fact that most matches are very poor and that this can be recognized early in the above sequence when the number of points in the matching is relatively small. This technique requires a suitable definition of satisfactory and poor matches.

A technique which has perhaps the greatest potential to reduce recognition times takes advantage of the saliency of the subtemplates. A strategy is adopted where, after matching only a few of the more salient subtemplates, the contents of affected accumulators are compared to a prescribed "detection"

threshold. If the contents of one of these accumulators exceed the threshold, it is taken as a candidate for the associated template's centroid; the entire template is rotated to the angle of match and an attempt is made to match the entire template to the boundary. Work on this technique is still preliminary but it has shown to provide a considerable speedup—approximately 40 CPU seconds for a match. A quantitative method of choosing a "detection" threshold has yet to be determined. Such a threshold is needed also to determine if a part is absent. The idea is based on the observation that a large part of the recognition time is spent on matching the more common low-weight subtemplates.

Finally, recognition times can be reduced if more than one

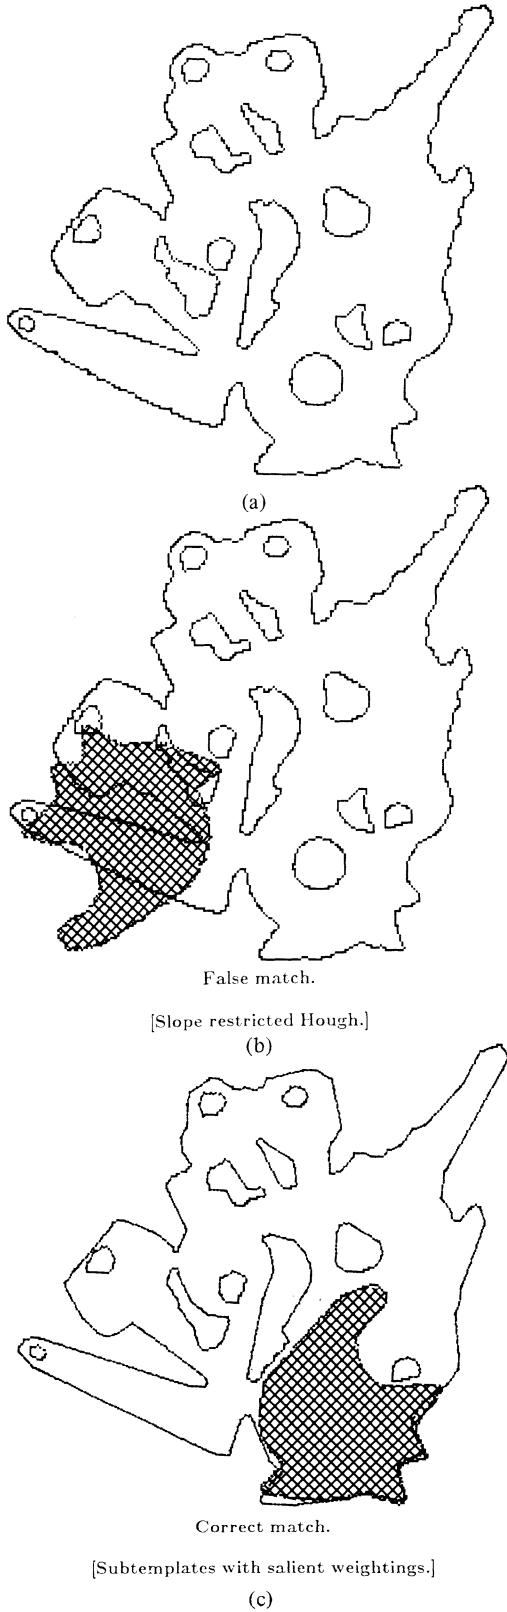


Fig. 14. Experiments with a set of nine parts.

processor were employed in running the algorithm. For example, each processor could be assigned to match the subtemplates for different parts. Strategies for multiprocessors are discussed further in [11].

The experimental results presented suggest that saliency can be applied with advantage in the occluded parts problem. The

concept is not intimately linked with boundary features; other features may be used, for example, internal edges are an obvious addition. Less obviously, preliminary studies suggest that future research consider working with subregions obtained from a three-dimensional sensor. The three-dimensional template is supplied by a solid model database; scaling is possible from the range information.

APPENDIX

Techniques in the numerical solution of least-squares problems include efficient methods for the solution of the following problem:

$$\text{Minimize: } \|F\mathbf{x} - \mathbf{f}\|_2$$

$$\text{Subject to: } G\mathbf{x} \geq \mathbf{g}$$

where F is an $n_1 \times n$ matrix, G is an $n_2 \times n$ matrix, \mathbf{x} is an n vector (to be determined), \mathbf{f} is an n_1 vector, and \mathbf{g} is an n_2 vector. The problem is referred to as a "least-squares with linear inequality constraints" (LSI) in [9], which provides a detailed description of the solution algorithm.

To take advantage of these solution techniques, it is first necessary to show that E is symmetric, positive, and definite.

Theorem: E is symmetric, positive, and definite.

Proof: Since E denotes the product $D^t D$, E is symmetric. If there exists an element of w , w_i that is nonzero, then the product $c_{ij}w_i$ is also nonzero because $c_{ij} > 0$ for all i and j (every τ_i and β_j match to some extent). This product must contribute to some accumulator. Thus, there exists an s such that $\delta_{ij}^s = 1$ and, consequently, such that $a_s > 0$. Therefore, $a^t a > 0$ provided $w \neq \underline{0}$. But from the minimization problem statement, $w^t E w = a^t a$. Thus, E satisfies the definition of a symmetric, positive, definite matrix. •

Since E is symmetric, positive, and definite, Cholesky factorization can be used to factor E into the form $U^t U$ where U is a real upper triangular $|\tau| \times |\tau|$ matrix. Consequently, the minimization problem of Section IV may be rewritten as the following least-squares problem:

$$\text{Minimize: } \|Uw\|_2$$

$$\text{Subject to: } \underline{1}^t w = 1 \quad \text{and } w \geq \underline{0}.$$

This is similar to a special case of the LSI problem, as can be seen by setting

$$F = U, \mathbf{x} = w, \mathbf{f} = \underline{0}, G = \begin{bmatrix} I & \underline{0} \\ \underline{1}^t & 1 \end{bmatrix}, \quad \text{and } \mathbf{g} = \begin{bmatrix} \underline{0} \\ 1 \end{bmatrix}$$

to yield the following:

$$\text{Minimize: } \|Uw\|_2$$

$$\text{Subject to: } \underline{1}^t w = 1 \quad \text{and } w \geq \underline{0}.$$

This is equivalent to the minimization problem, if the inequality $\underline{1}^t w \geq 1$ constraint is replaced with the equality, $\underline{1}^t w = 1$. The following theorem shows that the solution is unaffected by changing the inequality constraint to an equality. The theorem will be stated in terms of minimizing $w^t E w$ rather than in terms of the least-squares form involving U .

Theorem: Given that E is symmetric positive definite, the

problem

$$\text{Minimize: } \mathbf{w}^T \mathbf{E} \mathbf{w}$$

$$\text{Subject to: } \underline{1}^T \mathbf{w} \geq 1 \quad \text{and} \quad \mathbf{w} \geq \underline{0}$$

has the same solution if $\underline{1}^T \mathbf{w} \geq 1$ is replaced by $\underline{1}^T \mathbf{w} = 1$.

Proof: A complete proof is omitted for brevity. Instead, the basis for the validity of the theorem is sketched by referring to the two-dimensional case illustrated in Fig. 10. Clearly, the inequality $\underline{1}^T \mathbf{w} \geq 1$ allows the possibility of solutions above the line given by $w_1 + w_2 = 1$, i.e., when strict inequality holds. Suppose the point X satisfies the constraints $\underline{1}^T \mathbf{w} > 1$ and $\mathbf{w} \geq \underline{0}$ and represents the minimum solution, then there is a hyper-ellipse of constant value passing through X . Draw a line from the origin through X . Since $\underline{1}^T \mathbf{w} > 1$, there exists a point Y on the line between X and the line $w_1 + w_2 = 1$ which also satisfies all constraints. This can be expressed algebraically by

$$\mathbf{w}_Y = k \mathbf{w}_X, \quad \text{where } 0 < k < 1. \quad \text{Thus,}$$

$$\mathbf{w}_Y^T \mathbf{E} \mathbf{w}_Y = k^2 \mathbf{w}_X^T \mathbf{E} \mathbf{w}_X < \mathbf{w}_X^T \mathbf{E} \mathbf{w}_X$$

contradicting the minimality of X . Therefore, the solution point lies on the line $\underline{1}^T \mathbf{w} > 1$. \bullet

Thus, the minimization problem of Section IV can be solved using the LSI solution technique with the above substitutions for $F, \mathbf{x}, \mathbf{f}, G$, and \mathbf{g} .

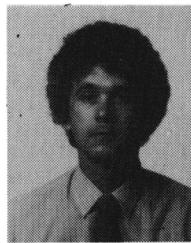
REFERENCES

- [1] A. Arbuschi, V. Cantoni, and G. Musso, "Recognition and location of mechanical parts using the Hough technique," in *Image Analysis*, S. Levialdi, Ed. London, England: Pitman, 1983.
- [2] D. H. Ballard, "Generalizing the Hough transform to detect arbitrary shapes," *Pattern Recogn.*, vol. 13, no. 2, pp. 111-122, 1981.
- [3] H. G. Barrow and R. J. Popplestone, "Relational descriptions in picture processing," *Machine Intell.*, vol. 6, pp. 377-396, 1971.
- [4] E. W. Baumann, "Model-based vision and the MCL language," in *Proc. IEEE Conf. Syst. Man Cybern.*, Oct. 1981, pp. 433-438.
- [5] E. W. Baumann, "CAD model input for robotic sensory systems," in *Proc. Autofact IV*, Nov. 1982.
- [6] H. J. Blinchikoff and A. I. Zverev, *Filtering in the Time and Frequency Domains*. New York: Wiley, 1976, ch. 7.
- [7] P. H. Eichel, "Sequential detection of linear features in two-dimensional random fields," Ph.D. dissertation, Dep. Comput. Inform. Contr. Eng., Univ. Michigan, Ann Arbor, May 1985.
- [8] P. V. C. Hough, "Methods and means for recognizing complex patterns," U.S. Patent 3069654, 1962.
- [9] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems*. Englewood Cliffs, NJ: Prentice-Hall, 1974.
- [10] P. M. Merlin and D. J. Farber, "A parallel mechanism for detecting curves in pictures," *IEEE Trans. Comput.*, vol. C-24, pp. 96-98, Jan. 1975.
- [11] T. N. Mudge and T. S. Abdel-Rahman, "Case study of a program for the recognition of occluded parts," in *Proc. 2nd Annu. IEEE Comput. Soc. Workshop Comput. Architecture Pattern Anal. Image Data Base Management*, Pasadena, CA, Oct. 1983, pp. 56-60.
- [12] W. A. Perkins, "A model-based vision system for industrial parts," *IEEE Trans. Comput.*, vol. C-27, pp. 126-143, Feb. 1978.
- [13] —, "Simplified model-based part locator," in *Proc. 5th Int. Conf. Pattern Recogn.*, Dec. 1980.
- [14] A. Rosenfeld and A. C. Kak, *Digital Picture Processing*. New York: Academic, 1976.
- [15] J. Sklansky, "On the Hough Technique for curve detection," *IEEE Trans. Comput.*, vol. C-27, pp. 923-926, Oct. 1978.
- [16] J. L. Turney, T. N. Mudge, R. A. Volz, and M. D. Diamond, "Experiments in occluded parts recognition using the generalized Hough transform," in *Proc. Conf. Artif. Intell.*, Oakland Univ., Rochester, MI, Apr. 1983.
- [17] J. L. Turney, T. N. Mudge, and R. A. Volz, "Experiments in occluded parts recognition," in *Proc. Soc. Photo-opt. Instrum. Engineers Cambridge Symp. Opt. Electro-opt. Eng.*, Cambridge, MA, Nov. 1983, vol. 449, part II, pp. 719-725.
- [18] R. A. Volz, T. N. Mudge, and D. A. Gal, "Using Ada as a robot system programming language for robot-based manufacturing cells," *IEEE Trans. Syst. Man Cybern.*, vol. SMC-14, no. 6, pp. 863-878, 1984.
- [19] J. Wolter, T. C. Woo, and R. A. Volz, "Gripping position for 3D objects," in *Proc. 1982 Meet. Ind. Appl. Soc.*, Oct. 1982, pp. 1309-1314.



Jerry L. Turney (S'84) received the B.S. degree in physics from Michigan State University, East Lansing, in 1972, the M.A. degree in physics from the University of California, Berkeley, in 1976, and the M.S. degree in computer, information and control engineering from the University of Michigan, Ann Arbor, in 1982.

He is currently working toward the Ph.D. degree in computer, information and control engineering at the University of Michigan. His research interests are computer architecture and computer vision.



Trevor N. Mudge (S'74-M'77-SM'84) received the B.Sc. degree in cybernetics from the University of Reading, England, in 1969, and the M.S. and Ph.D. degrees in computer science from the University of Illinois, Urbana, in 1973 and 1977, respectively.

He has been with the Department of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor, since 1977 as an Associate Professor. His research interests include computer systems architecture (both hardware and software), VLSI design, and computer vision.



Richard A. Volz (M'60) received the Ph.D. degree from Northwestern University, Evanston, IL.

From 1964 to 1976 he was an Assistant Professor at the University of Michigan, Ann Arbor. He is now Director of the Robot Systems Division of the Center for Robotics and Integrated Manufacturing at the University of Michigan. Prior to assuming this position, he held positions as Associate Director of the University Computer Center and Associate Chairman of the Department of Electrical Engineering and Computer Science. His current research includes the software/hardware computer architecture to support robot systems and the use of computer-aided design systems (CAD) for driving robot and sensor programming. Particular projects include CAD model driven systems, automatic determination of grasp points (from CAD information), graphic programming of robots, and distributed systems integration languages for real-time control.