LOCATING PARTIALLY VISIBLE OBJECTS: THE LOCAL FEATURE FOCUS METHOD

Robert C. Bolles
SRI International, Menlo Park, California 94025

ABSTRACT

A new method of locating partially visible two-dimensional objects has been designed. The method is applicable to complex industrial parts that may contain several occurrences of local features such as holes and corners. The matching process is robust, because it bases its decisions on groups of mutually consistent features, and it is relatively fast, because it concentrates on key features that are automatically selected on the basis of a detailed analysis of CAD type of models of the objects.

I. INTRODUCTION

There are several tasks that involve locating partially visible objects. They range from relatively easy tasks, such as locating a single two-dimensional object, to the extremely difficult task of locating and identifying three-dimensional objects jumbled together in a bin. In this paper, we describe a technique to locate and identify overlapping two-dimensional objects on the basis of two-dimensional models.

Sequential [1,2,3] and parallel [4,5] approaches have been taken to solve this problem. In the sequential approach, one feature after another is located and as much information as possible is derived from the position and orientation of each feature. This approach is fast because it locates the minimum number of features; however, if the objects are complicated, determining the order of the features to be located may be difficult. Development of the location strategy becomes even more difficult when mistakes are taken into account.

In the parallel approach, all the features in an image are located, and then large groups of features are used to recognize objects. Graph-matching [4,6], relaxation [5,7], or histogram [8,9] techniques can be used to determine the feature groups. This approach is robust because it bases its decisions on all the available information, and the location strategy is straightforward because all the features are used. For even moderately complex objects, however, the quantity of data to be processed makes use of this approach impractical on current computers.

Described here is a method called the Local Feature Focus (LFF), that combines the advantages of the sequential and parallel approaches, while avoiding some of their disadvantages. This is achieved by careful analysis of the object models and selection of the best features.

II. LOCAL FEATURE FOCUS METHOD

The basic principle of the LFF method is to locate one relatively reliable feature and use it to partially define a coordinate system within which a group of other key features is located. Enough of the secondary features are located to uniquely identify the focus feature and determine the position and orientation of the object of which it is a part. Robustness is achieved by using a parallel matching scheme to make the final decisions, and speed is achieved by carefully selecting information-rich features.

The idea of concentrating on one feature is not new; it has been used in several special-purpose vision programs [10,11]. What is new is the ability to generate the focus features and their secondary features automatically from object models. This automatic feature selection, when perfected, will significantly reduce the need for people to program recognition procedures and thus will make possible quick and inexpensive application of the LFF method to new objects.

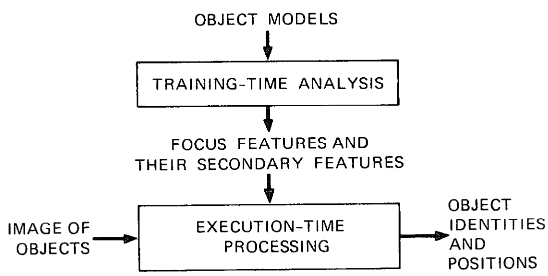

Figure 1 THE TOP-LEVEL BLOCK DIAGRAM

As Figure 1 shows, the analysis of object models is performed once during "training time" and the results of the analysis are used repeatedly during "execution time," making this approach particularly attractive when large numbers of objects are to be processed. In the rest of this paper, we concentrate on the training-time analysis.

III. ANALYSIS

The goal of the analysis is to examine a model of an object (or objects), such as the one in Figure 2, and generate a list of focus features and

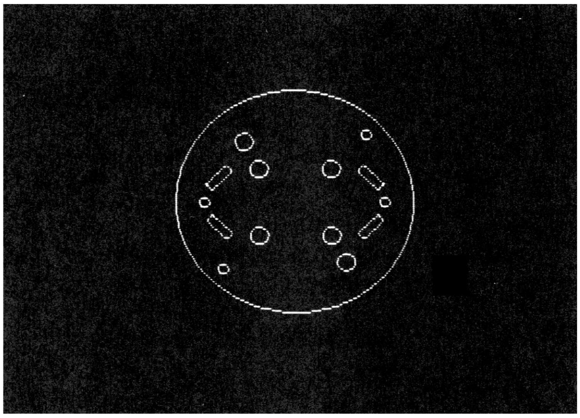

Figure 2 AN OBJECT MODEL

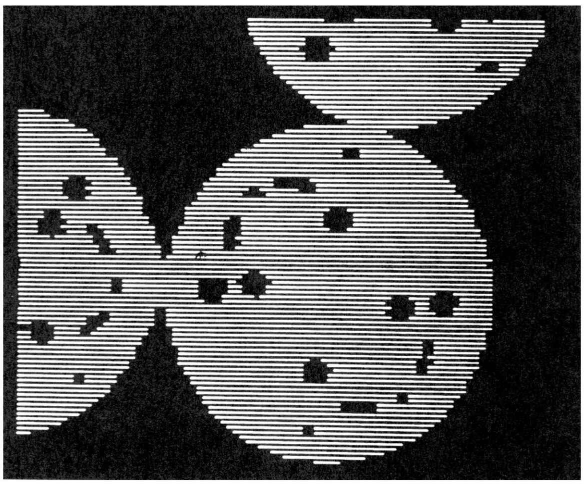

Figure 3 AN EXAMPLE TO BE PROCESSED

their associated secondary features. Given this information and a picture such as the one in Figure 3, the execution-time system tries to locate occurrences of the objects. In the current implementation of the system, objects are modeled as structures of regions, each of which is bounded by a sequence of line segments and arcs of circles. The execution-time system uses a maximal-clique graph-matching method to locate the groups of features that correspond to occurrences of the objects. Therefore, the analysis is tailored to produce the information required by the maximal-clique matching system. In particular, the description of each secondary feature includes the feature type, its distance from the focus feature, and a list of the possible identities for the feature. The analysis to produce this information is performed in five steps:

(1) Location of interesting features

(2) Grouping of similar features

(3) Rotational symmetry analysis of each object

(4) Selection of secondary features

(5) Ranking of focus features.

The purpose of the first step is to generate the set of all features of the objects that could be located at execution time. Typical features include holes, corners, protrusions, and intrusions. For the model in Figure 2, the set of features contains all 14 internal holes.

In the second step, the set of features is partitioned into subsets of "similar" features. Features are defined to be similar if they are likely to be indistinguishable at execution time. For the model in Figure 2, feature detectors can distinguish at most three types of holes: "slots," "small holes," and "large holes." Therefore, the set of interesting features is partitioned into three subsets, each defining a possible focus feature.

In the third step, a complete rotational symmetry analysis of each object is performed [12]. The rotational symmetry is used to determine the number of structurally different occurrences of each feature. Because the model in Figure 2 is twofold rotationally symmetric, the features occur in pairs, the members of which are indistinguishable on the basis of the relative positions of other features of the object. Instead of four types of small holes, there are only two, one on the axis between between the slots and one off that axis.

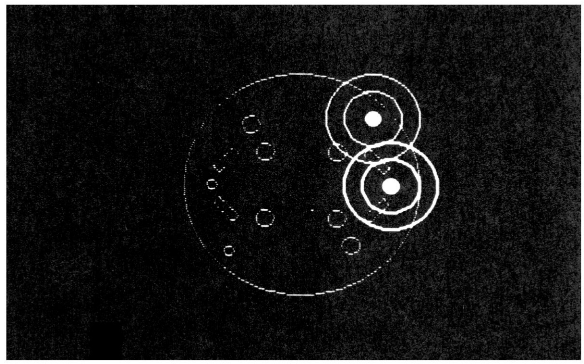

Figure 4 SECONDARY FEATURES FOR SMALL HOLES

The fourth step in the analysis is the most complicated. The goal is to select secondary features for each focus feature. The secondary features must distinguish between the structurally different occurrences of the focus feature and determine the position and orientation of the object. In Figure 2, for example, given an occurrence of a small hole, what nearby features could be used to determine whether it is one of the holes on the axis or off of it? There are two slots close to the small hole on the axis and only one near the off-axis occurrence. In addition, the slots are at different distances from the holes. Let D1 be the distance between the on-axis small hole and its slots and let D2 be the distance from the off-axis small hole to the nearest slot. Figure 4 shows circles of radii D1 and D2 centered on the two different types of small holes. Tabulated below are the feature occurrences that are sufficient to determine the type of the small hole and compute the position and orientation of the object.

ON-AXIS SMALL HOLE	OFF-AXIS SMALL HOLE
Two slots at D1	No slots at D1
No slots at D2	One slot at D2

The analysis in step 4 locates secondary features in two substeps. First, it performs a rotational symmetry analysis centered on each structurally different occurrence of a focus feature. This analysis builds a description of the object in terms of groups of features that are similar and equidistant from the focus feature. Figure 5 shows the groups of features produced by the current system when focusing on one of the small holes. In the second substep, the analysis iteratively selects groups of features from these descriptions to be included in the set of secondary features associated with the focus feature. Groups are selected for their contribution to identifying an occurrence of the focus feature or determining the position and orientation of the object.

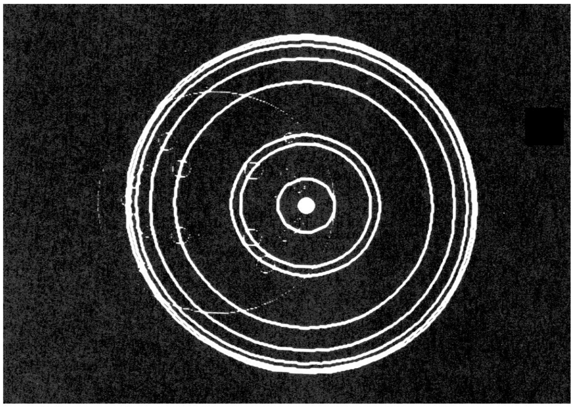

Figure 5 FEATURE GROUPS ABOUT A SMALL HOLE

The fifth and final step in the training-time analysis is the ranking of the focus features. The goal is to determine the order in which the focus features should be checked at execution time. The current system simply ranks them according to the number of secondary features required at execution time.

IV. DISCUSSION

The LFF method is a simple combination of the sequential and parallel approaches. It offers the reliability of a parallel approach and most of the speed of a sequential approach. The speed is achieved by using the location of the focus feature to define a coordinate system within which the other features are located. Quickly establishing a coordinate system significantly reduces the time required to find secondary features.

The utility of the LFF method depends on the reliability of locating focus features and the number of structurally different occurrences of these features in the objects. Fortunately, most industrial parts have good candidates for focus features. The problem is to find them at training time so they can be used at execution time. In fact, the more information gathered at training time, the more efficient the system at execution time. Also, as the training-time analysis is made more automatic, correspondingly less time is required of a human programmer.

The current implementation of the training-time analysis forms the basis for a completely automatic feature selection system. Several extensions are possible. For example, the system could select extra features to guarantee that the execution-time system would function properly even if a prespecified number of mistakes were made by the feature detectors. The system could use the orientation of a focus feature, if it exists, to determine the orientation of the feature-centered coordinate system. The system could also select two or more groups of features at one time, which is necessary for some more difficult tasks such as distinguishing an object from its mirror image. Finally, the system could incorporate the cost and reliability of locating a feature in the evaluation of the feature.

In conclusion, the LFF method is a combination of the sequential and parallel approaches that provides speed and reliability for many two-dimensional location tasks. The automatic selection of features makes it particularly attractive for industries such as the aircraft industry that have hundreds of thousands of different parts and cannot afford a special-purpose program for each one.

REFERENCES

[1] S. Tsuji and A. Nakamura, "Recognition of an Object in a Stack of Industrial Parts," Proc. IJCAI-75, Tbilisi, Georgia, USSR, pp. 811-818 (August 1975).

[2] S. W. Holland, "A Programmable Computer Vision System Based on Spatial Relationships," General Motors Research Pub. GMR-2078 (February 1976).

[3] W. A. Perkins, "A Model-Based Vision System for Industrial Parts," IEEE Transactions on Computers, Vol. C-27, pp. 126-143 (February 1978).

[4] A. P. Ambler et al., "A Versatile Computer-Controlled Assembly System," Proc. IJCAI-73, Stanford, California, pp. 298-307 (August 1973).

[5] S. W. Zucker and R. A. Hummel, "Toward a Low-Level Description of Dot Clusters: Labeling Edge, Interior, and Noise Points," Computer Graphics and Image Processing, Vol. 9, No. 3, pp. 213-233 (March 1979).

[6] R. C. Bolles, "Robust Feature Matching Through Maximal Cliques," Proc. SPIE's Technical Symposium on Imaging Applications for Automated Industrial Inspection and Assembly, Washington, D.C. (April 1979).

[7] S. T. Barnard and W. B. Thompson, "Disparaty Analysis of Images," To appear in IEEE Transactions on Pattern Analysis and Machine Intelligence (July 1980).

[8] R. O. Duda and P. E. Hart, "Use of the Hough Transform To Detect Lines and Curves in Pictures," CACM, Vol. 15, No. 1, pp. 11-15 (January 1972).

[9] S. Tsuji and F. Matsumoto, "Detection of Ellipses by a Modified Hough Transformation," IEEE Transactions on Computers, Vol. C-27, No. 8, pp. 777-781 (August 1978).

[10] J. T. Olsztyn and L. Rossol, "An Application of Computer Vision to a Simulated Assembly Task," Proc. IJCPR-73, Washington, D.C. (October 1973).

[11] D. F. McGhie, "Programmable Part Presentation," in the SRI Ninth Report on Machine Intelligence Research Applied to Industrial Automation, pp. 39-44 (August 1979).

[12] R. C. Bolles, "Symmetry Analysis of Two-Dimensional Patterns for Computer Vision," Proc. IJCAI-79, Tokyo, Japan, pp. 70-72 (August 1979).