Image Understanding Research at the University of Utah*

William B. Thompson and **Thomas C. Henderson**
Department of Computer Science
University of Utah
Salt Lake City, UT 84112

Abstract

Current work on image understanding at the University of Utah is focused principally on using computer vision and related sensing techniques to aid in problems associated with manufacturing. Three projects are now underway: feature-based reverse engineering, benchmarking of IU for man-made objects, and improved integration of inspection with design and manufacturing. In all three projects, the emphasis is on using information about "real-world" manufacturing equipment and operations, rather than general-purpose methods intended for more generic tasks. In addition, we are completing work on landmark-based navigation for mobile robots.

1 Overview.

The objective of this effort is to develop better methods for using image understanding technologies in a manufacturing environment. Our goal is to exploit information about real-world manufacturing processes, using sensing strategies which instead of extracting generic geometry are tuned to particular application tasks. Results are reported on three problems:

- **Reverse Engineering.** Reverse engineering of mechanical parts requires extraction of information about an instance of a particular part sufficient to replicate the part. Our method creates a CAD model of the original part, allowing easy modification and automatic generation of a process plan for appropriate fabrication techniques.

- **Benchmarking.** We are creating a test set for evaluating image understanding methods applied to man-made parts, particularly those likely to be seen in automated manufacturing operations. Rather than simply collecting data, we are using an innovative approach that provides accurate information about true geometry, control over variability, and the ability to evaluate novel sensors and active vision systems.

- **Integration of sensing with design.** Traditional approaches to visual inspection compare sensed geometry or other surface properties of a manufactured object with nominal descriptions of the intended object. Improvements in speed, accuracy, and reliability can be obtained through the use of sensing strategies which take into consideration the process plan used to control the manufacturing that originally produced the object.

We also report on additional results from the study of landmark-based navigation:

- **Reducing landmark ambiguity.** The visual recognition of landmarks is often subject to ambiguity within a landmark class. For example, it is much easier to recognize a visual feature as a landmark type (e.g., *peak*) than it is to uniquely classify the feature (e.g., *peak-1037*). A number of simple strategies can be used to reduce this ambiguity.

2 Image Understanding for Manufacturing.

Using image understanding technologies to improve manufacturing productivity has been discussed for many years. Except for certain specialized forms of 2–D analysis, however, this potential has seldom been realized. Limitations in processing capacity, the high costs of needed hardware and software, and the lack of a systems approach for developing image understanding methods within the larger context of manufacturing have all contributed to this lack of success. The situation may now be changing. Recent increases in computational power have been enormous. IC technology has brought about drastic decreases in sensor costs. Even more importantly, we now have a better understanding of the unique requirements of manufacturing and are able to generate task-specific approaches tailored to such problems.

There are three general application areas in which image understanding technologies can impact quality and cost in manufacturing. *Visual inspection* provides rapid, non-contact checks on post-production quality and performance monitoring of in-progress manufacturing operations. *Vision-guided manipulation* supports closed-loop control of parts handling

*This work was supported by the Advanced Research Projects agency under Army Research Office grant number DAAH04-93-G-0420 and by National Science Foundation grant IRI-9196146, with partial funding from the Defense Advanced Research Projects Agency.

and automated assembly, increasing the flexibility of automation systems. *Sensor-based reverse engineering* of mechanical parts makes it possible to create CAD models appropriate for computer-aided manufacturing directly from existing physical prototypes or similar objects for which usable CAD models do not exist.

Manufacturing is different than most other problems to which image understanding has been applied. The visual environment is highly structured and the identity of visible objects is often known. Precise object models are frequently available. IU systems must fit within well established methodologies and existing standards for design, geometric modeling, process planning, and fabrication. Alternate methods for solving the same problem usually exist and so speed, accuracy, and cost effectiveness are critical concerns.

These differences impose constraints that must be satisfied by IU systems if they are to be successfully fielded in a manufacturing plant. Cost effectiveness considerations often preclude IU solutions costing more than $50k or requiring more than a few seconds to complete. This places limits on the use of specialized hardware, exotic sensors, and expensive programming. Objects typically have complex geometries not easily modeled using polyhedra or low-order analytic surface patches that are common in existing 3–D IU systems. Parts are often metallic, with highly specular surfaces and additional complications due to dirt, rust, and the like.

3 Feature-Based Reverse Engineering.

The creation of CAD models appropriate for computer-aided manufacturing is an expensive and demanding task. Often these models are related to existing parts for which no usable CAD models exist or to non-computer generated shapes, such as the clay models used in the automotive industry. Automated methods for creating CAD models from existing solid objects (Figure 1) can contribute to the solution of this important problem by drawing on a synergistic combination of image understanding, geometric modeling, and graphics.

For shapes consisting solely of free-form surfaces, replicating an existing part can be done by sensing the geometry of the part with appropriate accuracy and density of samples and then driving an NC mill or other manufacturing device in a way that duplicates the geometry. Commercial three-dimensional position sensors are available which can automatically generate NC code for this task. Most manufactured objects are not free-form, however. They consist largely of manufacturing features such as holes, pockets, and profiles. Effective replication of such parts requires a *reverse engineering* approach in which aspects of the processes that originally created the parts are recovered. An added benefit of this approach is that the resulting CAD model can be easily revised, allowing modifications to be made to an existing part without the need to do a complete redesign.

Many image understanding methods for creating three-dimensional object models from sensed 2–D or 3–D data already exist. For the most part, these methods use generic modeling techniques applicable to a large class of objects and

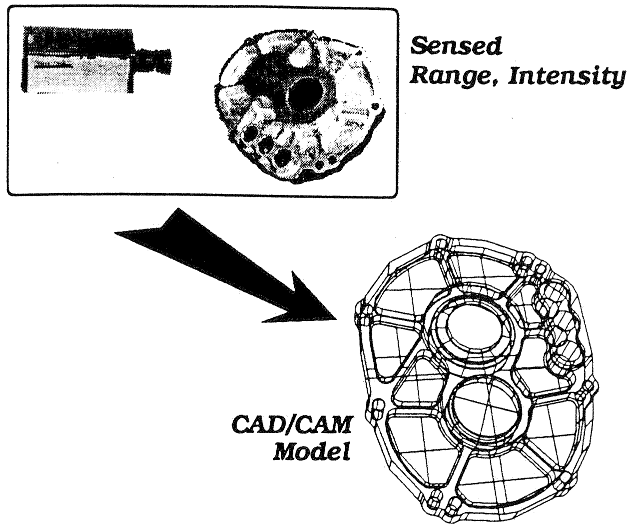

Figure 1: Reverse engineering of mechanical parts by creating CAD models from 3–D position data.

intended to support object recognition and pose estimation. Reverse engineering requires a degree of accuracy that can only be obtained if the geometric modeling technique employed can precisely represent the shapes being analyzed. In order to deal with noisy sensor data, the models should have no more than the number of degrees of freedom necessary to accomplish this representation. Finally, the models must be transferable to full-feature CAD/CAM systems to allow possible modification and the generation of a process plan for manufacturing. The standard modeling techniques in image understanding satisfy none of these requirements. They either cannot precisely represent commonly occurring parts in manufacturing or require a large number of parameters to do so and are thus noise sensitive. The models are difficult or impossible to export to standard CAD system. They fail to capture in any meaningful way information used in modern feature-based design systems.

Most machined parts are made using a relatively small number of manufacturing operations, each of a constrained form. Reverse engineering of such parts can be done using a form of parametric model fitting, where the primitives correspond to these features [Owen *et al.*, 1994]. This avoids inconsistencies between actual object shape and what the models are capable of representing, while leading in a natural and obvious way to representations usable in full-function CAD/CAM systems. (See also [Sobh *et al.*, 1994a, Sobh *et al.*, 1994b, Sobh and Owen, in press, Sobh *et al.*, in press].)

Two separate problems must be solved: a particular set of manufacturing features adequate to produce the part must be chosen and then the parameters of each feature must be estimated from sensed position data. Often, there is no single, unique set of required features. For example, a circular aperture might be the result of a drilling operation with a large bit or a milling operation with a small cutter. Complex manufacturing knowledge may be needed to make an appro-

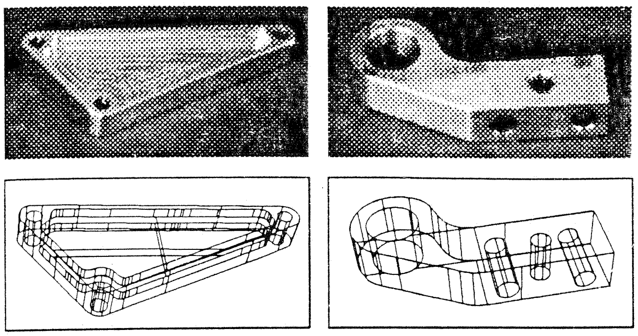

Figure 2: Models created from physical objects using feature-based reverse engineering.

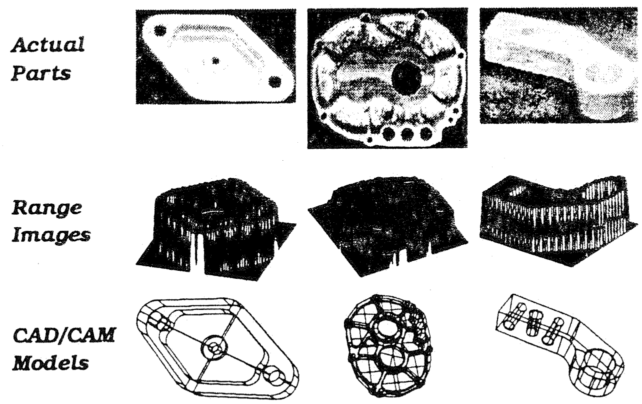

Figure 3: IU test data consisting of actual objects, images, and models.

priate choice. We use an interactive approach to deal with this ambiguity. A machinist or design engineer specifies the features to be used and information about each feature such as tolerances that cannot be determined using data acquired from a single part. The features are then fit to the sensor data automatically and a CAD model is produced for the part. We think of the system as a form of *electronic calipers* — a smart measuring tool, specialized to the job of creating CAD/CAM models using knowledge of the manufacturing process.

Figure 2 shows results from the reverse engineering of two machined parts. Both parts were designed for the Utah mini-Baja/formula SAE racing vehicle. For this experiment, they were simplified by removing chamfers and threads that were too small to be accurately measured with the sensors available to us. The object on the left of the figure is a shock mounting plate used in the rear suspension. The object on the right is part of the steering arm assembly. The images on top are views of the parts taken from a camera system which returns registered range and intensity values. On the bottom are wire frame renderings of the reverse engineered CAD model.

4 "Hard-Copy" Benchmark Suite for IU in Manufacturing.

The development of benchmarks for measuring progress on IU for manufacturing presents both challenges and opportunities compared to other areas of image understanding research. High positional accuracy is usually required. As a result, methods can be effectively evaluated only if precise "ground truth" data about visible objects used for testing is available. Parts and assemblies with complex shapes require views from different vantage points, often with alternate fixturing. The effectiveness of methods for evaluating the integration of multiple sensors, perhaps including both visual and contact sensing modalities, needs to be measured. Image understanding for manufacturing must be able to deal with variations in part appearance due to manufacturing tolerances and outright errors. Testing this requires calibrated data with known and controlled variability. A mechanism is needed to evaluate active positioning of sensor and lighting and the use of novel sensors, and to compare the effectiveness of these techniques with more traditional methods.

The IU community has made substantial efforts to develop test data appropriate for quantifying the performance of image understanding systems. (See [Thompson and Owen, 1994] for a partial listing.) Almost all of this data has been designed to test navigation systems or related technologies such as motion compensation and target recognition. Benchmarking such systems is difficult, since there is little control over the variability of the environment and ground truth about scene properties can be difficult or impossible to obtain.

The traditional approach to creating test data for IU systems involves acquiring a set of images and then, when possible, determining camera models and measuring selected scene parameters. By restricting attention to problems associated with manufactured parts, a different approach can be adopted. Instead of starting with existing objects, objects used for testing can be designed from scratch using modern CAD tools. This provides control over variability and allows for objects designed to test particular aspects of IU algorithms. By using precision NC milling to produce actual objects, the CAD models serve as an accurate representation of the true part geometry.

We are in the process of using this approach to design a set of test data for evaluating the performance of IU systems appropriate for manufacturing tasks. In addition to known object geometry and control over object shape and variability, there exists another advantage unique to this approach. We can easily replicate objects to within the tolerances of the NC milling process, typically about ± 0.001" except on complex sculptured surfaces.[1] As a result, actual objects can be included as part of the test set, allowing evaluation of methods which utilize alternate sensors or active control of sensing (Figure 3). When completed, the data available for distribution will include:

- Machined aluminum parts (the "hard copy").
- CAD models for each part.

[1] We employ the English system of measures in this paper, since that is what is used by most machinists in the United States.

heat source tracker being constructed as part of ARPA's Madefast program. It was machined in our own laboratory using a modern 3-axis mill. The outlined area in the figure is centered around a bump in the supposedly circular outer profile of the mirror. The bump occurred as a result of transients that sometime happen when the motors in the mill's X–Y stage reverse direction. Similarly, other manufacturing features have common failure modes: incorrect tools are selected, tools deflect, high rotational speeds lead to surface roughness, etc. Clearly, inspection systems should focus extra attention on such areas of likely failure.

6 Reducing Ambiguity in Landmark Identification.

Navigation is often based on bearings to distinctive *landmarks* in known locations. A careful choice of landmarks can reduce the uncertainty in position estimates [Sutherland, 1993, Sutherland and Thompson, 1993, Sutherland and Thompson, in press] or the error in following a pre-specified path [Sutherland and Thompson, 1994]. Unless some sort of active beacons emitting distinctive signatures are used as landmarks, landmark identification is a serious problem. For a wide range of navigation problems, the principal problem is ambiguity within a landmark class. It is much easier to detect indoor features such as doors, corridors, etc., and outdoor landmarks such as roads, junctions, peaks, etc., than it is to correctly determine the particular door, junction, or whatever.

The standard approach is to base a unique identification on the more subtle features which distinguish one landmark from another of the same type. This is prone to error and clearly infeasible when several identical looking landmarks are present. Higher order reasoning can be used to reduce the ambiguity, either by utilizing geometric constraints [Thompson *et al.*, 1993] or through an analysis of the mistakes likely to occur in path following [Stuck, 1992]. This involves substantial computational effort, however, particularly when the level of ambiguity is high. Fortunately, a number of simple strategies can be used to reduce within-class landmark ambiguity.

Terrestrial navigation problems involve maneuvering through an essentially two-dimensional space. When a general estimate of current position is available, it is usually possible to determine which landmarks are likely to be in view, even when their location in the view is highly uncertain. In such circumstances, the landmark identification problem simplifies to one of landmark ordering. Figure 5 provides an example involving viewpoint determination in outdoor, unstructured terrain lacking easily recognized cultural features. Views from two different locations are shown. The viewpoints for the two views are relatively close together and the same three peaks are visible in each view. These peaks are marked with a • symbol on the map. Estimating the viewpoint based on a particular view depends on correctly ordering the peaks in the view in terms of their correspondence to peaks **A**, **B**, and C on the map.

A number of simple "common sense" concepts can be used in the process of ordering landmarks such as these in a way

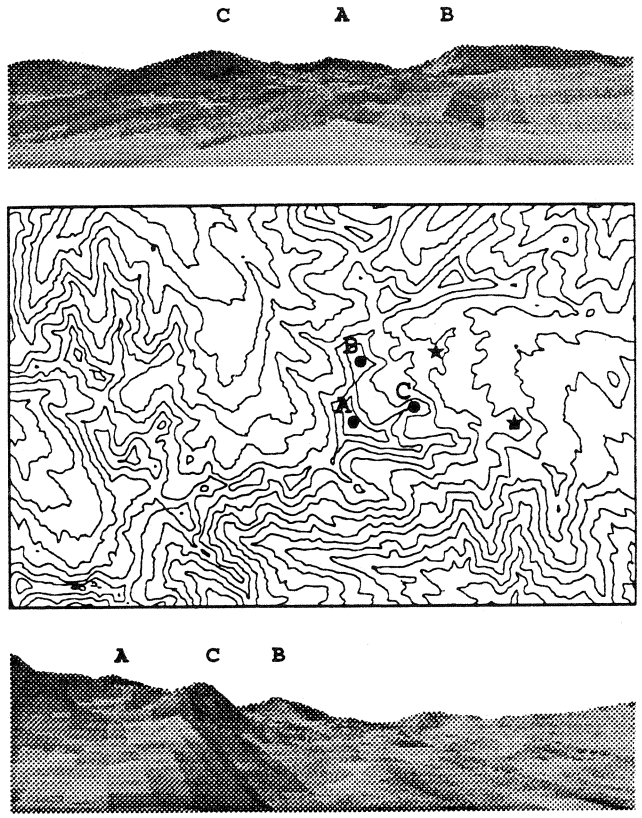

Figure 5: Ambiguity in associating visual landmarks with map locations.

that reduces much of the time consuming task of complete object identification. In the example in Figure 5, if the viewer knows only that landmark **B** is on the right, then the angle from the viewpoint to the outer two features constrains the ordering of the other two features, even without information on the actual viewpoint. If the viewer is looking at the landmarks in the order **ACB**, it is impossible for the angle to the outer two features to be greater than angle **BCA**. Thus, if the angle measures greater than **BCA**, one can immediately deduce that the landmark order is **CAB**. It is, of course, possible for the angle between the outer two features to be less than angle **BCA**, in which case, the angle measure alone will not help in determining landmark order. However, whenever quick, easy measurements such as this can eliminate the need for more complicated procedures, they should be used.

In addition to single angular measurements aiding in the ordering of landmarks, a navigator can take a measurement, move a small amount and take another with the resulting change in the measurements revealing the order of the landmarks. The process involves measuring the angle from the viewpoint to the two leftmost landmarks and the angle from the viewpoint to the two rightmost landmarks then moving toward one of the outer landmark points and remeasuring both angles. This can be easily accomplished by keeping that landmark point in the center of the image plane. Moving toward an outer feature produces the difference in the measures that determines the side of the configuration of landmarks on which the viewpoint lies. If movement was toward the center feature, both visual angles would increase in size. As an example, again consider the map shown in Figure 5 with the only viewer knowledge being that landmark A is in the center. As shown in [Sutherland, 1994a, Sutherland, 1994b], the movement required to observe a difference in angle measure significant enough to determine landmark order is so small that the distinction between views is almost imperceptible to a human observer.

References

[Henderson *et al.*, 1994] T.C. Henderson, T.M. Sobh, F. Zana, B. Brüderlin, and C.-Y. Hsu. Sensing strategies based on manufacturing knowledge. In *Proc. ARPA Image Understanding Workshop*, November 1994.

[Owen *et al.*, 1994] J.C. Owen, P.-P.J. Sloan, W.B. Thompson, and T.C. Henderson. Interactive feature-based reverse engineering of mechanical parts. In *Proc. ARPA Image Understanding Workshop*, November 1994.

[Sobh and Owen, in press] T.M. Sobh and J.C. Owen. Reverse engineering. *Journal of Robotics and Autonomous Systems*, in press.

[Sobh *et al.*, 1994a] T. M. Sobh, M. Dekhil, and J.C. Owen. Discrete event control for inspection and reverse engineering. *Proc. IEEE International Conference on Robotics and Automation*, May 1994.

[Sobh *et al.*, 1994b] T.M. Sobh, J.C. Owen, M. Dekhil, C. Jaynes, and T.C. Henderson. Industrial inspection and reverse engineering. In *Proc. IEEE 2nd CAD-Based Vision Workshop*, February 1994.

[Sobh *et al.*, in press] T.M. Sobh, J.C. Owen, M. Dekhil, and T.C. Henderson. Industrial inspection and reverse engineering. *Computer Vision, Graphics, and Image Processing – Image Understanding*, in press.

[Stuck, 1992] E.R. Stuck. *Detecting and Diagnosing Mistakes in Inexact Vision-based Navigation*. PhD thesis, University of Minnesota, 1992.

[Sutherland and Thompson, 1993] K.T. Sutherland and W.B. Thompson. Inexact navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, May 1993.

[Sutherland and Thompson, 1994] K.T. Sutherland and W.B. Thompson. Pursuing projections: Keeping a robot on path. *Proc. IEEE International Conference on Robotics and Automation*, pages 3355–3361, May 1994.

[Sutherland and Thompson, in press] K.T. Sutherland and W.B. Thompson. Localizing in unstructured environments: Dealing with the errors. *IEEE Trans. on Robotics and Automation*, in press.

[Sutherland, 1993] K.T. Sutherland. Landmark selection for accurate navigation. In *Proc. DARPA Image Understanding Workshop*, April 1993.

[Sutherland, 1994a] K.T. Sutherland. Ordering landmarks in a view. In *Proc. ARPA Image Understanding Workshop*, November 1994.

[Sutherland, 1994b] K.T. Sutherland. The stability of geometric inference in location determination. Technical Report UUCS-94-021, University of Utah, 1994.

[Thompson and Owen, 1994] W.B. Thompson and J.C. Owen. "Hard-copy" benchmark suite for image understanding in manufacturing. In *Proc. ARPA Image Understanding Workshop*, November 1994.

[Thompson *et al.*, 1993] W.B. Thompson, T.C. Henderson, T.L. Colvin, L.B. Dick, and C.M. Valiquette. Vision-based localization. In *Proc. DARPA Image Understanding Workshop*, April 1993.