

@Make(article)
@Device(x2700)
@Use[Bibliography="<henderson.proposals>general.bib"]
@Use[Bibliography="<bhanu>bhanu.bib"]
@begin(format)

58

@MajorHeading(Intrinsic Characteristics as the Interface between CAD and
Machine Vision Systems@foot[This work was supported in part by NSF Grants
ECS-83-07483, MCS-82-21750, DCR-83-41796 and MCS-81-21750, DARPA Grant
DAAK11-84-K-0017, and ONR Grant N00014-82-K-0351. All opinions, findings,
conclusions or recommendations expressed in this document are those of the
authors and do not necessarily reflect the views of the sponsoring agencies.])

@begin(center)
Thomas C. Henderson and Bir Bhanu

Department of Computer Science
The University of Utah
Salt Lake City, Utah 84112

@end(center)
@b[Abstract]

@end(format)
Computer Aided Design systems are currently being used to drive machine
vision analysis. We propose that the set of intrinsic 3-D shape
characteristics provided
by the CAD system be used to compare representational power of different
CAD systems so far as the requirements for machine vision are concerned.
@begin(format)

@i(Keywords): CAD models, shape representation, intrinsic characteristics
@end(format)

@Newpage

@Section(Introduction)

Several researchers have recently reported methods for deriving visual
recognition strategies based on CAD representation systems. Such an
approach makes it possible to produce recognition and analysis procedures
without having to scan a physical example of the object. Typically, these
proposed techniques either directly use whatever the CAD model
produces or derive information (e.g., points sampled on the surface of the
object) to drive a particular recognition scheme. In this regard,
there has been some discussion as to an appropriate set of interface data
(e.g., points, surface patches, features, etc.). We propose that a
coherent general solution to this problem is to characterize a CAD system
by the set of intrinsic 3-D shape characteristics (e.g., surface normals,
texture, reflectance properties, curvature, etc.) that the CAD system is
able to provide.

Such a characterization makes it possible to compare CAD systems
irrespective of recognition paradigms, and actually makes it possible to
determine which recognition strategies can be used with a given CAD-based
machine vision system.

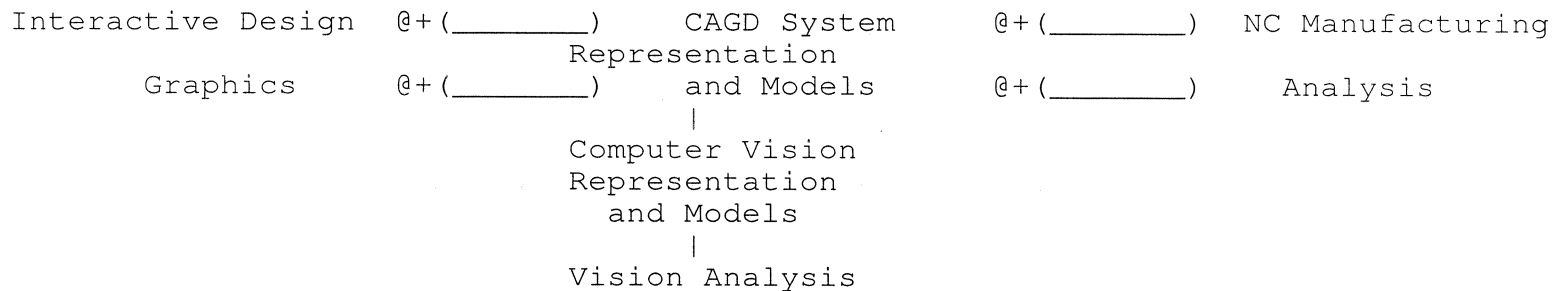
This paper presents the initial results of our work in using Computer
Aided Geometric Design (CAGD) representations and models as a basis for the
visual recognition of 3-D objects for robotic applications.

The representation and analysis of 3-D shape is a strong common concern of
researchers in both Computer Aided Design (CAD) and machine vision.
Unfortunately, the functional requirements of the two user groups led to
the development of representations and models which facilitated the

achievement of different kinds of processing. CAD systems typically emphasize such capabilities as rendering, set operations, etc. while machine vision models must provide information which makes possible the automatic analysis of camera data. Moreover, CAD systems are used to design new shapes, whereas machine vision systems are used to analyze objects already in existence. Thus, many machine vision systems require a "presentation" step in which an example of the object to be modeled is shown to the system, and the corresponding internal representation of the object is generated. Recently, it has become possible to merge these different aspects and requirements to obtain an integrated facility in which an object may be designed and which can support vision analysis. Several proposals for systems along these lines have been made [Bhanu85b, Castore84].

We believe that this approach (i.e., CAD-based machine vision) offers many advantages and makes it possible to have a systematic approach to the analysis of shape across a wide spectrum of requirements. Current CAD systems offer an interactive design environment and provide facilities to create images of the designed parts, perform analysis functions on them (e.g., finite element analysis), and produce numerically-controlled machining information for manufacturing. Figure 1 shows the system we envision.

@begin(verbatim)



@center(@b[Figure 1]. CAGD-Based Vision Analysis System)
@end(verbatim)

In the remainder of the paper, we discuss the nature of CAD and machine vision representations and describe an appropriate interface between them. In particular,

Section 2 reviews CAD representations and models; Section 3 gives an overview of machine vision representations. In Section 4, we argue that intrinsic 3-D shape characteristics provide a convenient interface between CAD and vision analysis. Finally, in Section 5 a summary is given.

@Section(Representation in CAGD)

Constructive solid geometry (CSG) and boundary representations are the best understood and currently most important representation schemes in computer aided design. Present day 3-D wireframe models used in CAD and model-based vision have many deficiencies including ambiguity - it is easy to build a wireframe model that can be surfaced in several ways [RequichaVoelcker82]. In CSG, the basic idea is that complicated solids can be represented as various ordered "additions" and "subtractions" of simpler solids by means of modified versions of Boolean set operators-union, difference and intersection [Requicha80]. For inherent boundary representations a number of different approaches are used. These include Coons patches, bicubic surface patches, Bezier methods and B-splines [BarnhillRiesenfeld74].

Current Geometrical Modeling Systems (GMS) use a limited class of primitives such as rectilinear blocks and conic surfaces (cylinders, cones and spheres). Although these are sufficient to cover a large number of conventional unsculptured parts, a GMS which includes sculptured solids is highly desirable. Also since the sculptured design is surface

oriented, it is easier to incorporate it in a boundary based system. In general, boundary modelers tend to support stepwise construction of the models more easily than CSG modelers but require greater data storage. CSG modelers are inadequate for modeling sculptured parts: they have no capability at all for constructing and using sculptured surfaces as part of the boundary of the solid model. Some advantages of boundary representation are: there are many known surface models available from which to choose@cite[BarnhillRiesenfeld74]; the mathematics of surface representation is well developed and complex shapes can often be represented with a single primitive@cite[Cobb84,Thomas84]; and it results in an intuitive model. A minor disadvantage is that it may be difficult to ensure the validity of a boundary representation of a set. On the other hand, CSG representations are not unique in general, since a solid may be constructed in many ways; the final result may not be easily visualized by looking at the primitives. However, the CSG representation is concise, validity is guaranteed and such a representation can be easily converted to a boundary representation. The comparison of CSG and boundary representation methods can be found in@cite[RequichaVoelcker82,RequichaVoelcker83].

Recently there have been attempts to use a set of manipulative operations for boundary models for solid objects to construct a solid modeling system@cite[MantylaSulonen82,Thomas84]. These are designed for CAD/CAM environments, rather than for computer vision applications. In@cite[MantylaSulonen82], a set of Euler operators is used on the topology of a boundary model, that is on the relative arrangement of its faces, edges and vertices. The operations allow the system to perform arbitrary modifications necessary for boundary representation models, the faces of which are planar polygons.

Until recently it was not possible to carry out Boolean operations on sculptured surfaces. Recent work by Thomas@cite[Thomas84] attempts to combine the best attributes of CSG and surface-based representation systems by using subdivision techniques developed by Cohen et al.@cite[CohenLycheRiesenfeld80]. He uses a uniform boundary representation. The "primitives" are solids bounded by B-spline surfaces. As compared to the other work in solid modeling, his method does not require that the objects being combined have closed boundaries; they must only satisfy a weak completion criterion. Thus this method results in a powerful shape description system which allows the combination of primitives using set operations into arbitrarily-complex objects bounded by curved surfaces and the production of a model which represents such objects. Adjacency information about surface points and the intersection curve between two surfaces as a polyline can be obtained. Although he has used B-spline surfaces, his techniques are applicable to any surface representation scheme@cite[Cobb84]. All this work has been incorporated in the Alpha_1 system@cite[Cohen83]. (More details about Alpha_1 are presented below.) Thus, the advantages of both CSG and sculptured surface representation can be obtained in the shape representation of objects and the combination of objects via set operations. As a result of these significant advances in CAGD, we decided to use the Alpha_1 system for exploring the computer vision application.

Alpha_1 is an experimental CAGD based solid modeler system incorporating sculptured surfaces@cite[Cohen83]. It allows in a single system both high quality computer graphics and freeform surface representation and design. It uses a rational polynomial spline representation of arbitrary degree to represent the basic shapes of the models. The rational spline includes all spline polynomial representations for which the denominator

is trivial. Nontrivial denominators lead to all conic curves. Alpha_1 uses the Oslo algorithm@cite[CohenLycheRiesenfeld80] for computing discrete B-splines. Subdivision, effected by the Oslo algorithm, supports various capabilities including the computation associated with Boolean operations, such as the intersection of two arbitrary surfaces@cite[Thomas84]. B-splines are an ideal design tool, they are simple yet powerful; many common shapes can be represented exactly using rational B-splines. For example, all of the common primitive shapes used in CSG systems fall into this category. Other advantages include good computational and representational properties of the spline approximation: the variation diminishing property, the convex hull property and the local interpolation property. There are techniques for matching a spline-represented boundary curve against raw data. Although the final result may be an approximation, it can be computed to any desired precision (which permits nonuniform sampling). At present, tolerancing information is not included in the object specification in Alpha_1 system. It is planned to be incorporated in the future. Once it is available, we can make our models in terms of classes of objects (rather than a single object) which are functionally equivalent and interchangeable in assembly operations.

Given the CAGD model (perhaps by combining several modeling paradigms), a corresponding set of vision models (with some control structure) is generated. Once these models are available, they provide the basis for standard 2-D and 3-D scene analysis.

An early example of such an interactive system is the ACRONYM system@cite[Brooks81,BrooksGreinerBinford79] designed for applications in computer vision and manipulation. The world is described to ACRONYM as volume elements and their spatial relationships and as classes of objects and their subclass relationships. It uses a hybrid CSG and general sweep scheme for the representation of rigid solids. The representations are CSG-like trees whose leaves are generalized cylinders. Like PADL (a geometric modeling system@cite[Brown82]) it allows variation in size, limited variation in structure and variation in structural relationships of the modeled objects. However, in ACRONYM, it may be difficult to design algorithms for computing properties of objects.

@Section(Representation in Machine Vision)

Geometric modeling is one of the key components of a domain-independent model-based 3-D industrial machine vision system. Here our interest is in the representation, modeling and recognition of rigid, opaque 3-D solid objects.

Three general classes for the representation of 3-D solid objects are (a) surface or boundary, (b) volume, and (c) sweep@cite[BadlerBajcsy78,Requicha80].

In the boundary representation schemes, a 3-D solid object is represented by segmenting its boundary into a finite number of bounded "faces" and describing the structural relationships between the faces.

Another approach to surface representation is to express the surfaces as functions on the

"Gaussian Sphere"@Cite[HornWoodhamSilver78,Smith79].

Volumetric representations include spatial occupancy, cell decomposition and constructive solid geometry@Cite[RequichaVoelcker82,RequichaVoelcker83]. Sweep representations consist of translational sweep, rotational sweep, 3-D sweep and general sweep.

Since a direct model of a 3-D object in terms of its volume (e.g., as a 3-D array) may easily exhaust the memory capacity of a system, representation by oct-trees has been considered@Cite[JackinsTanimoto80]. These may make space array operations more economical in terms of memory space.

A simple approach to analyzing 3-D objects is to model them as polyhedra.

This requires a description of the objects in terms of vertices, edges and faces. Baumgart@cite[Baumgart74] developed a 3-D geometric modeling system ("Geomed") for application to computer vision. He used a face-based representation for planar polyhedral objects, called the "winged-edge" representation. The Euler primitives are used for polyhedron construction and shape operators include union, intersection and difference. Geomed provides many capabilities; for example, arbitrary polyhedra may be constructed, altered or viewed in perspective with hidden lines eliminated. Bolles et al.@cite[BollesHosaharaHannan83] have used a CAD model that contains a standard volume-surface, edge-vertex description as well as pointers linking topologically connected features. Their preliminary model uses a pointer structure similar to Baumgart's "winged-edge" representation. Wesley et al.@cite[WesleyPerez80] have used polyhedral models for automated mechanical assembly in their geometrical modeling system GDP. Their Automated Parts Assembly System (AUTOPASS)@cite[LiebermanWesley77] language has never been implemented. Before automated assembly can be successful, it is essential to have robust representations, models, and general purpose techniques for determining the orientation and position of 3-D objects for a large class of industrial parts.

For curved and more complex objects, other representations and models have been used, such as generalized cylinders (or cones). Generalized cylinders or cones are a quite popular representation in computer vision@cite[Agin72,Brooks81,NevatiaBinford77]. However, there are some problems with this representation. There are infinitely many possible generalized cones representing a single object. More constraints are needed to get a unique description. Although it is possible to represent arbitrary shapes with generalized cones by making them arbitrarily complex, their computation is difficult. They are also not well suited to descriptions of non-elongated objects and objects of arbitrarily deformed surfaces enclosing little volume.

Although sweep representations, such as generalized cylinders, and volume representations, such as constructive solid geometry, imply surface description, they fail to describe the junction or surface peculiarities. In the recognition of 3-D objects from partial views, we detect surfaces first, and only after seeing several different views of the object do we have enough data to obtain volume properties. For objects constructed from thin sheet-like material, surfaces are natural candidates for representation. Further, surfaces are seen first. As such, they are important for computer vision. Hence the need for surface or boundary based representations.

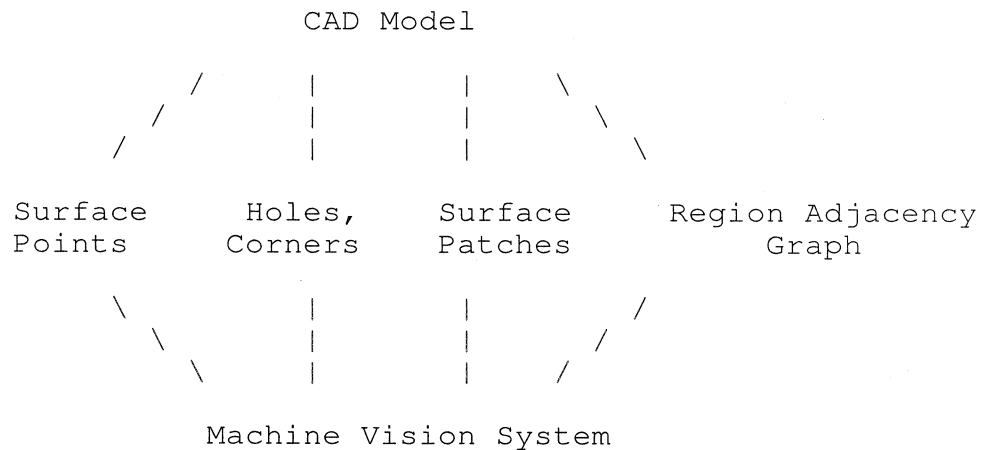
York et al.@cite[YorkHansonRiseman81] have used a structured collection of Coons surface patches for representing 3-D objects whose boundary curves are approximated by cubic B-splines. Their design of Coons patches is cumbersome, since it requires that a simple surface patch be designed on paper before it may be entered into the data base. Brady@cite[Brady84] proposes a symbolic representation of visible surface based on "curvature patches." They are computed locally by determining the tangent vectors that indicate directions in which the surface changes. Example directions include the principal curvature directions and the directions in which the normal curvature is zero. Smooth changes in curvature patch descriptions are obtained to determine the larger scale structure of a surface. It is not clear that curvature patch surfaces are perceptually "fairer" than surfaces developed in CAD. @Section(Intrinsic Characteristics as the Interface)

In order to bridge the gap between shape modeling and shape analysis, it is necessary to give a detailed account of how the machine vision analysis can be performed in terms of the CAD shape model.

Bhanu and Henderson@cite[BHanu85b] have proposed a variety of information that can be produced and used to drive the vision analysis (see Figure 2).

@begin(figure)

@begin(verbatim)



@center(@b[Figure 2.] Typical CAD/Vision Interface)

@end(verbatim)

@end(figure)

In particular, a CAD system can return a set of points sampled from the surface of the shape (i.e., a range finder can be simulated), or a set of surface features or patches can be generated, or even better, a set of surface patches and their connectivity relation can be produced. Which of these is actually produced would depend on the recognition paradigm being used and on the class of shapes to be recognized.

However, in this paper, we propose that intrinsic characteristics of the 3-D shape be used as the interface between the CAD system and the machine vision model. There are several motivations for choosing intrinsic characteristics as the organizing notion:

@begin(enumerate)

The intrinsic characteristic approach has been quite successful as a machine vision paradigm for 3-D object analysis. Although originally proposed by Barrow and Tenenbaum@cite[Barrow76] as a set of registered images, the method applies equally well to 3-D data structures. Moreover, intrinsic characteristics are just those that are viewpoint independent and which are inherent in the particular shape being modeled.

Many shape analysis algorithms are based on the use of one or more intrinsic characteristics; thus, once it is known what intrinsic characteristics can be easily produced by a particular CAD system, then the possible set of recognition techniques is known.

Finally, different CAD systems can be compared on the basis of the set of intrinsic characteristics which they can provide. Thus, the relative tradeoffs in choosing a CAD system can be known more easily with respect to the set of recognition tasks that the system will have to perform.

@end(enumerate)

@Section(Discussion)

We have been studying techniques and algorithms which allow the generation of computer representations and geometric models of complicated realizable 3-D objects in a systematic manner. In order to produce recognition strategies for a machine vision system, it is necessary to specify the interface between the CAD system and the machine vision analysis system. We suggest that this interface be characterized by the set of intrinsic 3-D shape characteristics which can be produced by the particular CAD system

under consideration. Notice that the existence of the CAD model may preclude the necessity of a separate machine vision model, and that by providing intrinsic features, the scene analysis strategies (and executable code) can be directly generated.