

@Make(Article)
@Device(ln01)
@Use[Bibliography="/n/sp/s/u/tch/proposals/general.bib"]
@begin(format)

@MajorHeading(CAD-Based Robotics@foot(This work was
supported in part by NSF Grants MCS-8221750, DCR-8506393, and DMC-8502115.
))

@begin(center)
Thomas C. Henderson, Chuck Hansen, Eliot Weitz and Bir Bhanu
Department of Computer Science
The University of Utah
Salt Lake City, Utah 84112

@b[Abstract]

@end(center)
@end(format)

We describe an approach which facilitates and makes explicit the organization of the knowledge necessary to map robotic system requirements onto an appropriate assembly of algorithms, processors, sensors, and actuators. In order to achieve this mapping, several kinds of knowledge are needed. In this paper, we describe a system under development which exploits the Computer Aided Design (CAD) database in order to synthesize:

@begin(itemize)
recognition code for vision systems (both 2-D and 3-D),

grasping sites for simple parallel grippers, and

manipulation strategies for dextrous manipulation.

@end(itemize)

We use an object-based approach and
give an example application of the system to CAD-based 2-D vision.

@Newpage

@Section(Introduction)

The rapid design of embedded electromechanical systems is crucial to success in manufacturing and defense applications. In order to achieve such a goal, it is necessary to develop design environments for the specification, simulation, construction and validation of multisensor systems. Designing and prototyping such complex systems involves integrating mechanical parts, software, electronic hardware, sensors and actuators. Design of each of these kinds of components requires appropriate insight and knowledge. This in turn has given rise to special computer-based design tools in each of these domains. Such Computer Aided Design (CAD) systems have greatly amplified the power and range of the human designer. To date, however, it is still extremely difficult to address overall system issues concerning how the components fit together, and how the complete system will perform.

It is crucial to develop a design environment in which these multiple facets of system design can take place in a coordinated way such that the description of one component can be easily interfaced to another component, even when they are radically different kinds of things (e.g., a control algorithm, a mechanical linkage and an actuator). The designer should have the freedom to try out ideas at different levels of detail; i.e., from the level of a sketch to a fully detailed design. Figure 1 shows the general set of components required in the system we envision.

```
@begin(FullPageFigure)
@Blankspace(7inches)
@Center(@b[Figure 1.] General Components of a CAD-Based Robotics System)
@end(FullPageFigure)
The use of the CAD database
provides part of the solution to developing such an environment.
```

The system described in Figure 1 consists of four major components:

```
@begin(enumerate)
@b[the Multisensor Knowledge System (MKS)] - this specifies the knowledge of
all available algorithms, processors, actuators, sensors which can be
used in system construction; in addition, other knowledge, such as
environmental constraints (e.g., lighting, obstacles, etc.) can also
be specified,
```

```
@b[the Computer Aided Geometric Design (CAGD) System] - this is an
interactive design system which provides design and analysis capabilities,
```

```
@b[the Requirements Specification Interface] - this is essentially a
primitive task specification language for the user to identify the
task (e.g., 2-D visual inspection for structural defects), and
```

```
@b[the Problem Specific Rules] - these rules take as input the task
requirements specification and use the knowledge in @b[MKS] and the CAGD
system to synthesize a system which satisfies the requirements. Thus,
the rules themselves encode domain-specific knowledge for the
applications of interest.
```

```
@end(enumerate)
```

Finally, the synthesized systems are packaged as "Logical Sensors" and are then available in the @b[MKS] system as future system resources. For example, an edge detector may become a separate Logical Sensor.

The overall objective of this project is to explore the use of various kinds of knowledge, including the use of Computer Aided Geometric Design (CAGD) representations and models as a basis for the manipulation and visual recognition of objects for robotic applications. Specific goals are to:

```
@begin(enumerate)
```

Develop techniques and algorithms which allow the @i(interactive and automatic) generation of computer representations and geometric models of complicated realizable 3-D objects in a systematic manner. These representations and models are obtained using (a) the available CAGD techniques and (b) the data acquired from various range finding techniques.

Develop fast parallel algorithms, reasoning techniques and recognition strategies which are natural to the representations and models developed in 1) for matching 3-D models to the 3-D scene description derived from range data so as to be able to recognize, identify and locate 3-D objects, including partially occluded objects.

```
@end(enumerate)
```

The successful achievement of the goals stated above will be a step forward in the direction of bridging the gap between the fields of CAGD and computer vision in so far as the issues related to representation and modeling of 3-D objects are concerned. This work provides a fundamental understanding of the requirements and the role of boundary/surface models in computer vision. It provides a systematic way of building object models, as opposed to the @i(ad hoc) techniques currently in use. These models can be used in the recognition of objects in 2-D images for tasks such as photo interpretation, navigation and guidance, or as proposed here, for finding the orientation and position of 3-D objects in space for their manipulation by robots. The long-term goal is the automated assembly

of objects with parts designed and manufactured using the knowledge in a CAGD system.

@begin(comment)
@Section(Objects and Methods)

Several distinct programming styles have been developed over the last few years, including applicative-style programming, control-based programming, logic programming, and object-based programming. The object-based style takes the view that the major concern of programming is essentially the definition, creation, manipulation and interaction of objects; that is, a set of independent and well-defined data structures. In particular, a single data structure (or instance) is associated with a fixed set of subprograms (methods), and those subprograms are the only operations defined on that object. This style emphasizes data abstraction combined with message passing@cite[Booch83, Organick83].

An object is a structure with internal state (perhaps called @i(slots) and comprised of name/value relationships) accessed through functions (also called @i[methods]) defined in association with the object. This approach makes management schemes simpler and fewer, easier to implement and use; in addition, individual resources are easier to specify, create (allocate), destroy (deallocate), manipulate and protect from misuse.

It has been effectively argued many times that object-based programming is well-suited to embedded systems processing requirements. In particular, the application of this methodology to the specification of recognition and manipulation systems helps to directly describe most of the important aspects of such systems: parallel processing, real-time control, exception handling, and unique I/O control.
@end(comment)

@Section(CAD-Based Robotics)

In this section, we describe the components of our CAD-Based Robotics approach. The two major components are the Multisensor Knowledge System (@b[MKS]) and the CAGD System. The current CAGD system that we are using is the Alpha_1 system, while work is underway on the construction of @b[MKS]. We are developing the system in the context of several applications projects, including 2-D visual inspection, 3-D computer vision and object manipulation with both simple parallel grippers and more sophisticated dextrous hands.

@SubSection(The Multisensor Knowledge System)

Much of our previous work on multisensor systems has concentrated on the specification of such systems and reasoning about their properties. It is necessary to be able to describe both the parameters and characteristics of individual components of multisensor systems, and to be able to deduce global properties of complete systems. Although it may be possible to deduce such properties (especially static properties like complexity, data type coercion, etc.), we believe that many interesting properties can only be determined by simulating the operation of the complete system.

Thus, we seek a representation that supports:

@begin(enumerate)
@b[multisensor system specification]: this describes the components and interconnection scheme of the particular system being designed,

@b[sensor, algorithm, processor and actuator knowledge representation]: this

structures information about sensor characteristics (e.g., accuracy, hysteresis, dynamic range, etc.), algorithms (e.g., space and time complexity, amenity to parallel computation, stability, etc.) processors (e.g., cycle times, memory limits, address space, power requirements, etc.), and actuators (e.g., actuation principle, power requirements, etc.), and

@b[multisensor system simulation]: this permits one to monitor important parameters and to evaluate system performance.

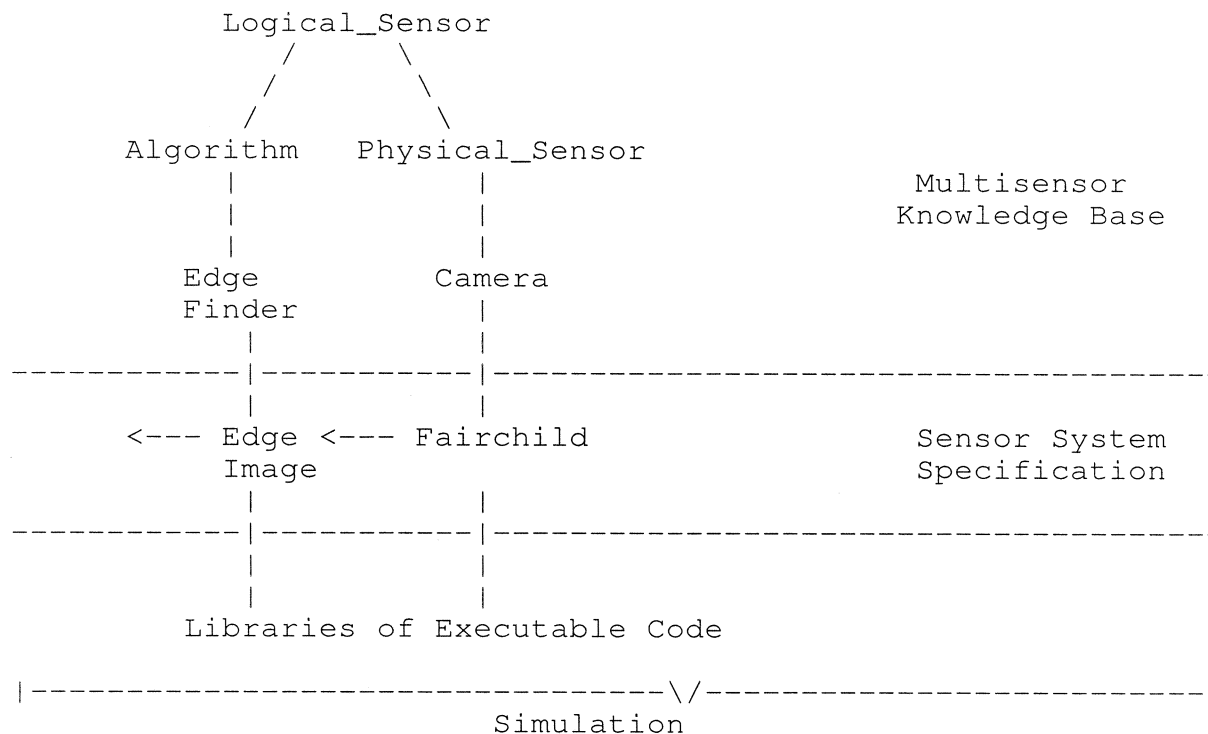
@end(enumerate)

Figure 2 shows the organization of the three capabilities within an object-oriented context.

@begin(FullPageFigure)

@begin(comment)

@begin(verbatim)



@end(verbatim)

@end(comment)

@Blankspace(8inches)

@center(@b[Figure 2]. Multisensor Knowledge System)

@end(FullPageFigure)

In the following subsections, we describe @b[MKS], an object-based approach to providing a unified answer to these three capabilities.

@SubSection(The Multisensor Knowledge Base)

The multisensor knowledge base serves two main purposes:

@begin(enumerate)

to describe the properties of the system components (e.g., sensors, algorithms, actuators and processors), and

to provide class descriptions for the actual devices which are interconnected in any particular logical sensor specification.

@end(enumerate)

That is, the knowledge base must describe not only generic sensors (e.g., cameras), but specific sensors (e.g., Fairchild 9000, Serial No. 28753).

It is then possible to reason about sensor systems at several levels.

Moreover, it is possible that two distinct system specifications require some of

the same physical sensors. In such a case, it is the responsibility of the


execution environment to resolve resource allocation conflicts.

We have chosen a frame-like knowledge representation. Frames relate very naturally to object-based descriptions, and, in fact, can be viewed as a class of restricted objects. It is straightforward to provide hierarchical descriptions of system components. For example, Figure 3 shows the @i(CCD Camera) hierarchy.

@begin(FullPageFigure)

@begin(comment)

@begin(verbatim)



@end(verbatim)

@end(comment)

@Blankspace(8inches)

@center(@b[Figure 3]. Organization of Camera Knowledge)

@end(FullPageFigure)

The @i(CCD Camera) frame has two slots: element spacing and aspect ratio. These slots are specific to CCD cameras and as such do not appear as slots for @i(2-D cameras). These latter have slots for scanning format, scan timing, resolution, output signal, and operating conditions. These slots are inherited by any instance of CCD camera. One level up, we find a frame for @i(Vision) sensors. This frame has specific slots for the spectral band and for the output type (e.g., 2-D byte array, multi-band, etc.). At the highest level of the hierarchy is the @i(Sensor) frame which has a slot for the physics of operation. This slot is used by any particular sensor to allow for an explanation of the physics behind the workings of the sensor. In this way, if reasoning is required about the sensor, it is possible to look in this slot for information.

Note that frames are themselves implemented as objects. Thus, actual devices are instances of some class of objects. This is very concise and conveniently exploits the similarities of frames and objects.

In previous work, we have described a set of generally applicable physical sensor features@cite[Henderson84a]. The manner in which physical sensors convert physical properties to some alternative form, i.e., their transducer performance, can be characterized by:

error,
accuracy,
repeatability,
drift,
resolution,
hysteresis,
threshold,
and range.

These properties can be encoded in the appropriate slots in the frames describing the sensor.

@SubSection(Sensor Specification)

We have previously introduced the Multisensor Kernel System and Logical Sensor Specifications as a means for high-level specification of multisensor systems. The main goals of such a characterization are:

- to develop a coherent treatment of multisensor information,
- to allow system reconfiguration for both fault tolerance and dynamic response to environmental conditions, and
- to permit the explicit description of control.

Logical Sensor Specifications (LSS) permit an implementation independent description of the required sensors and algorithms in a multisensor system. Figure 4 gives a pictorial description of the basic unit: a @i(logical sensor).

```
@begin(FullPageFigure)
@Blankspace(8inches)
@begin(format)
```

```
@center(@b[Figure 4.] Logical Sensor Specification Building Block:
The Logical Sensor)
```

```
@end(format)
@end(FullPageFigure)
```

Sensor data flows up through the currently executing program (one of program@-[1] to program@-[n]) whose output is characterized by the @i(characteristic output vector). Control commands are accepted by the @i(control command interpreter) which then issues the appropriate control commands to the Logical Sensors currently providing input to the selected program. The programs 1 through n provide alternative ways of producing the same characteristic output vector for the logical sensor. The role of the @i(selector) is to monitor the data produced by the currently selected program and the control commands. If failure of the program or a lower level input logical sensor is detected, the selector must undertake the appropriate error recovery mechanism and choose an alternative method (if possible) to produce the characteristic output vector. In addition, the selector must determine if the control commands require the execution of a different program to compute the characteristic output vector (i.e., whether dynamic reconfiguration is necessary).

Logical Sensor Specifications are useful then for any system composed of several sensors, where sensor reconfiguration is required, or where sensors must be actively controlled. The principle motivations for Logical Sensor Specifications are the emergence of significant multisensor and dynamically controlled systems, the benefits of data abstraction, and the availability of smart sensors. (For more on the various aspects of Logical Sensors,

see@cite[Henderson84c,Shilcrat84a,Henderson85a,Henderson85h,Henderson86c,Henderson86d].)

Related work includes that of Albus@cite[Albus81] on hierarchical control, Bajcsy et al.@cite[Bajcsy84a] on the Graphical Image Processing Language, Overton@cite[Overton86] on schemas, and Chiu@cite[Chiu86] on functional language and multiprocessor implementations. For an overview of multisensor integration, see Mitiche and Aggarwal@cite[Mitiche86].

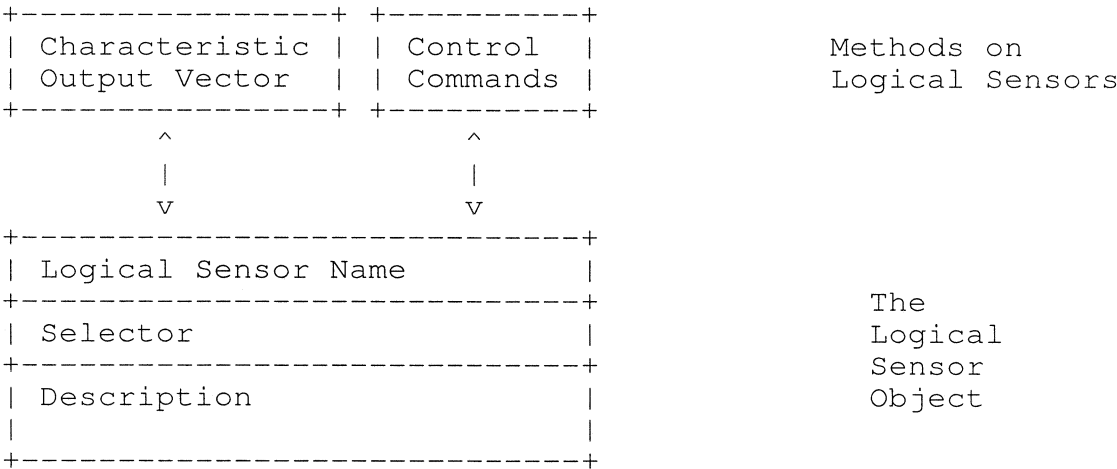
In exploring these issues, we have found that the specification of multisensor systems involves more than just sensor features. It is true that knowledge must be available concerning sensors, but it is essential to also be able to describe algorithms which use the sensor data and the hardware on which they are executed. In addition, the geometric

knowledge provided by the CAD database helps to focus the synthesis of recognition and manipulation strategies. In the rest of the paper, we describe the components of an object-based approach to developing a knowledge system to support these requirements.

An object-based style of programming requires that the logical sensor of Figure 4 be re-described in terms of objects and methods. We shall next give the general flavor of this style, but it must be remembered that any particular sensor is actually an instance of some object class, and, in fact, inherits properties from many levels up.

Each logical sensor is completely specified as shown in Figure 5.

```
@begin(FullPageFigure)
@begin(comment)
@begin(verbatim)
```



```
@end(verbatim)
@end(comment)
@Blankspace(8inches)
@center(@b[Figure 5]. The Logical Sensor Object and Methods)
@end(FullPageFigure)
```

Thus, in order to get data from a logical sensor, the @i(characteristic output vector method) must be invoked. Likewise, to issue control commands to the sensor (e.g., camera pan and tilt parameters), the @i(control commands method) must be used. The role of the @i(selector) is still the same as in previous logical sensor implementations, however, it now, in essence, is invoked to produce the characteristic output vector.

Such a representation makes it very easy to design sensor systems. Moreover, such specifications allow for replacement of sensors and dynamic reconfiguration by simply having the @i(selector) send messages to different objects. Given current object-based programming technology, such systems can be rapidly developed and permit dynamic typechecking (on objects).

Figure 6 shows the Multisensor Knowledge Base, and below the dashed line, a set of particular instances of various algorithms, sensors, etc. (drawn as circles).

```
@begin(FullPageFigure)
@begin(format)
@Blankspace(8inches)

@center(@b[Figure 6]. Logical Sensor Specification Using Object Instances)

@end(format)
@end(FullPageFigure)
A logical sensor specification (indicated as a blocked in subset
```

of the circles) defines a grouping of algorithms, sensors, etc. This newly created logical sensor is an instance of the @i(logical sensor object) and can be sent messages. As mentioned above, there are two methods defined on logical sensors: the @i(characteristic output vector method) and the @i(control commands method). Thus, any logical sensor can be defined recursively in terms of other logical sensors (including itself).

Currently, our main interest is in the automatic synthesis of logical sensor specifications. Given a CAD model of an object, we would like to synthesize a specific, tailor-made system to inspect, recognize, locate or manipulate the object.

Note that the synthesis of a logical sensor specification consists, for the most part, of interconnecting instances of sensors and algorithms to perform the task. This is done by writing the selector to invoke methods on other logical sensors. Given certain constrained problems, most notably the CAD/CAM environment, such a synthesis is possible.

@SubSection(The Alpha_1 CAGD System)

Alpha_1 is an experimental solid modeling system developed at the University of Utah. For the past few years the Computer Aided Geometric Design group has been involved in a concerted effort to build this advanced modeler. Alpha_1 incorporates sculptured surfaces and embodies many theoretical and algorithmic advances. 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 B-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 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. B-splines are an ideal design tool, they are simple, yet powerful. It is also the case that 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).

@SubSection(The Simulation of Multisensor Systems)

Effective simulation plays an important role in successful system development. A key requirement is the support for hierarchical specification of the system and the ability to perform stepwise refinement of the system. In addition, it is necessary to be able to efficiently emulate realtime software that will eventually be embedded in the system. Finally, it would be quite useful to be able to embed physical components in the simulator in order to monitor the system's operation.

An object-oriented simulation methodology is well-suited to satisfy these goals. The multisensor system, that is, the system being modeled, consists of a collection of interacting physical processes. Each such process is

modeled in the simulator by an object, i.e., an independent process. Interactions among physical processes are modeled through messages exchanged among the corresponding objects.

This general paradigm is currently supported in the SIMON simulator developed by Fujimoto@cite[Fujimoto85,Swope86]. A toolkit approach is used in which the simulator is viewed as a collection of mechanisms from which application specific simulation environments are constructed. We are currently exploring the simulation of multisensor systems in the SIMON environment. Simulation can be accomplished by substituting simulation libraries for the run time libraries (see Figure 2).

A crucial aspect of the simulation is the ability to execute specific algorithms on specific hardware. SIMON permits such a direct execution technique in which application programs are executed directly on a host processor rather than through a software interpreter. Performance information is obtained through the automatic insertion of probes and timing software into the program at compile time. These probes perform whatever runtime analysis is required to accurately estimate execution time of basic blocks of code. A prototype implementation using this technique has been developed modeling the MC68010 and 68020 microprocessors. Initial data indicate that application programs may be emulated one to two orders of magnitude more efficiently over traditional register transfer level simulation, while highly accurate performance estimates can still be obtained.

@Section(An Example Application: CAD-Based 2-D Vision)

A simple example which demonstrates some of the power of the Multisensor Knowledge System approach is that of CAD-Based 2-D Vision. The goal is to automate visual inspection, recognition and localization of parts using pattern recognition techniques on features extracted from binary images. Figure 7 shows the scheme pictorially.

```
@begin(FullPageFigure)
@Blankspace(8inches)
@center(@b[Figure 7]. Synthesis of Part Detector)
@end(FullPageFigure)
```

The Multisensor Knowledge System stores knowledge about the algorithms, sensors, processors, etc. This knowledge is used by application specific rules. The systems to be synthesized here require that a model be created for the part to be inspected, and that a robust and (perhaps) independent set of features be chosen along with an appropriate distance metric.

The left side of the figure shows the offline training component. The new part is designed using a Computer Aided Geometric Design system. A set of images are rendered by the CAGD system giving a sample of various views of the part in different positions, orientations, and scales. These serve as a training set to the Multisensor Knowledge System.

A set of rules (or productions) performs an analysis of the views of the part to select a subset of the total set of possible features. Features are used if they are robust, independent and reliable. Once these features have been chosen, a new logical sensor object is created whose only function is to recognize the given part based on an analysis of the selected features. The part detector is then linked into a particular application (e.g., an inspection task at a specific workcell) by sending a message to the appropriate camera.

As a specific example, consider the object shown in Figure 8. It was designed using Alpha_1.

```
@begin(figure)
@begin(format)
```

@Blankspace(4.5 inches)

@center(@b[Figure 8]. The Green Piece)

@end(format)

@end(figure)

The object shown in Figure 8 was rendered at orientations of 0, 22.5 and 45 degrees. These images were analyzed. The set of possible features included are shown in Figure 9.

@Tabclear

@Tabset(1inch)

@begin(figure)

@\Area

@\Diameter

@\Perimeter

@\Perimeter to Area Ratio

@\Holes

@\Thinness

@\Axes of Inertia Ratio

@\Invariant Moment 1

@\Invariant Moment 2

@\Invariant Moment 3

@\Invariant Moment 4

@\Invariant Moment 5

@\Invariant Moment 6

@\Invariant Moment 7

@center(@b[Figure 9.] Available Features for 2-D Inspection)

@end(figure)

The seven invariant moments are those given by Hu@cite[Hu62].

The model produced based on analyzing the images is given in Figure 10.

@Tabclear

@Tabset(1inch, 4inches)

@begin(figure)

@\AREA@67428.335938

@\PERIMETER@2446.000000

@\X_Y_INERT@1.200313

@\HOLES@11

@\MOMENT2@21.311333

@\MOMENT3@20.067244

@\MOMENT4@20.979874

@\MOMENT6@20.173754

@center(@b[Figure 10.] Model Values on Selected Features)

@end(figure)

The result of the analysis was the following weightings are shown in Figure 11.

@Tabclear

@Tabset(1inch, 4inches)

@begin(figure)

@\AREA@0.860404

@\Perimeter@0.001159

@\X_Y_INERT@0.098228

@\Moment 2@0.035425

@\Moment 3@0.000354

@\Moment 4@0.003999

@\Moment 6@0.000432

@center(@b[Figure 11.] Feature Value Weights Used in Distance Function)

@end(figure)

The synthesized logical sensor object merely sends a message to

the segment program for Camera 1 (a Fairchild 3000 CCD camera), then sends a message to each of the features used, then sends a message to the distance function object with the appropriate weights. The system has been implemented in FROBS@cite[Muehle86] (FRames and OBjectS, a system which supports frame objects and is implemented in Common Lisp) using objects and methods. The feature calculations are performed by running C code called from within the instances of the feature objects. The feature giving the number of holes has not been used in the distance calculation, since it is qualitatively different kind of feature. Future work will include incorporating the integration of such information into the analysis.

@Section(Summary and Future Work)

CAD-based robotics offers many advantages for the design, construction, and simulation of robotics systems. We have described many of those. We are currently working on CAD-Based 3-D vision system. That is, we are developing a set of rules which will evaluate the 3-D geometry and function of any part designed with the Alpha_1 CAGD system. In this way, weak recognition methods can be avoided and specially tailored logical sensor objects can be synthesized automatically. Another area of current research interest is the simulation of multisensor systems. We believe that our approach can lead to very natural, straightforward, and useful simulations which can include native code running on the target processors. Finally, we are also investigating the organization of knowledge in the Multisensor Knowledge Base. Certain structuring of the data may lead to improved or simplified analysis.