SAND89-1087 UC-32
Unlimited Release
Printed June 1989

SAND--89-1087

DE89 013875

ROBOTIC GRASPING OF UNKNOWN OBJECTS: A KNOWLEDGE-BASED APPROACH

S. A. Stansfield
Intelligent Machine Principles Division
Sandia National Laboratories
Albuquerque, NM 87185

Abstract

In this paper, we demonstrate a general-purpose robotic grasping system for use in unstructured environments. Using computer vision and a compact set of heuristics, the system automatically generates the robot arm and hand motions required for grasping an unmodeled object. The utility of such a system is most evident in environments where the robot will have to grasp and manipulate a variety of unknown objects, but where many of the manipulation tasks may be relatively simple. Examples of such domains are planetary exploration and astronaut assistance, undersea salvage and rescue, and nuclear waste site clean-up. This work implements a two-stage model of grasping: stage one is an orientation of the hand and wrist and a ballistic reach toward the object; stage two is hand preshaping and adjustment. Visual features are first extracted from the unmodeled object. These features and their relations are used by an expert system to generate a set of valid reach/grasps for the object. These grasps are then used in driving the robot hand and arm to bring the fingers into contact with the object in the desired configuration. Experimental results are presented to illustrate the functioning of the system.

1 Introduction

Research into general-purpose grasping and robot haptics has applicability to several problem domains. Short term success will enhance our under-

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

standing of what is required, in terms of hardware, software, and models, to make articulated, sensate robot hands more useful and will guide us in designing the next generation of these devices. Basic research in robot haptics may also help us to determine how much of the control of teleoperated manipulators can be accomplished automatically and how much must be assumed by the human operator. Long term success in these areas will have broad application to the area of automated assembly, especially in small batch production. But it is in the area of unstructured, hazardous environments where this type of research will find its most useful application. There is a common thread in such diverse applications as nuclear waste site clean-up, planetary exploration and astronaut assistance, and undersea salvage and rescue – the environment is highly variable. Take, for example, the problem of nuclear waste site clean-up. The numbers and types of materials to be handled are not well known and cannot be modeled for the robot. The robot must be capable of handling objects with which it is not familiar, which may be broken or covered with clots of soil, or which may be temporarily fixed in some way to objects nearby. On the other hand, unlike an application such as automated assembly, the robot does not need to handle these materials for fine manipulation. Grasping and transporting objects for sorting and packing, and carrying out elementary tasks such as reorienting, separating, and brushing, may frequently be all that is required.

The research presented in this paper addresses the problem of general-purpose robotic grasping for unstructured environments. We have designed and implemented a system which generates reach/grasps for unmodeled objects. The system integrates visual perception of the object to be grasped with high-level knowledge about the relationships between extracted object features and the set of valid grasps for the object. Results from studies of human grasping and perception have been integrated into the work when we believe that they provide a good model or help to simplify the solution. The remainder of this paper discusses the components of our system in detail and presents a set of experimental results which illustrate the functioning system.

2 Related Work

Much of the previous and current work in robotic grasping is analytical and studies how a desirable, stable grasp may be chosen. Several methods for determining grasps have been proposed. Nguyen [Ngu88] determines force-closure grasps for polyhedral objects by finding independent regions of contact which totally constrain the motion of the object. Li and Sastry [Li 87] propose three quality measures for choosing an optimal grasp, including a task-oriented measure which takes the expected forces and moments of the task into account. Salisbury [Sal85] uses screw theory to determine how internal grasp forces may be used to constrain the set of stable grasps, while Hanafusa and Asada [Han77] utilize the potential energy in compliant fingers to determine grasp stability.

The above researchers address the issue of choosing a stable grasp. They do not, however, address the question of how this stable grasp is to be acquired. There are two distinct approaches to the problem of grasp acquisition – analytical and senor-based. Mason [Mas85] uses quasi-static, planar pushing operations to ensure that the object will move into the proper position to allow stable grasp contacts to be made. Trinkle and Paul [Tri89] use the same quasi-static analysis, however they consider the case of multiple pushers and negligible friction. Neither work requires sensory input. Fearing [Fea86] utilizes local tactile sensing of the surface normals of polygonal objects to determine the finger motions required to reach feasible grasping locations. Initial contact with the object is assumed. Lozano-Perez et al. [Loz87] use a range sensor to locate a known object from a pile of objects. The system generates a set of grasps and regrasps for a parallel-jaw gripper to move the object safely from its start position to some desired end position. Metrically accurate polyhedral models of all objects are assumed and the set of stable grasps is limited by the gripper. Ikeuchi et al. [Ike86] use photometric and binocular stereo vision systems to extract such information as surface orientation and range from an object. This information is then used to bring a parallel-jaw gripper into contact with the object in a legal grasp configuration. Legal grasps are again limited by the forces which the gripper can apply to the object. Rao et al. [Rao88] also use a vision system to extract information about the object to be grasped. A range acquisition system is used to extract 3D data which is then fit to a Generalized Cone. Four grasp modes are defined for the Belgrade-USC hand and the vision system is used to choose from among them for the given object. Mode selection is done either via a table look-up or via a sorted grasp mode list for the object.

In the final work described above, the authors also propose the use of heuristics to choose an appropriate grasp from the sorted grasp mode list. Other researchers have also proposed the addition of heuristics to a grasping system. Tomovic et al. [Tom87] and Iberall et al. [Ibe88] both propose knowledge-based systems for use in grasp selection. Tomovic incorporates the idea of reflex control, while Iberall focuses on planning the posture. Cutkosky [Cut88] presents an expert system which performs grasp selection in the manufacturing domain. Information about grasp parameters is input by the user and then utilized by the system to choose an appropriate grasp. In none of the latter three works, however, is the information provided by the expert system used to drive actual devices.

The approach presented in this work integrates knowledge and perception to drive robotic grasp. A structured-lighting vision system is utilized to extract aspects, or views, of the object to be grasped. This set of aspects is then used by a rule-based system to generate a set of grasps for the object. The parameters generated by this knowledge-based system are used by a set of lower-level motor modules to drive preshaping and grasp adjustment for the robot hand and wrist orientation and reach for the robot arm in order to bring the hand into contact with the object in the desired configuration. At the heart of our system is the symbolic representation constructed by the visual system and utilized by the rule-based system in generating grasps. The system requires neither an a priori knowledge of the object to be grasped, nor a dense set of sensory data such as is often needed to build geometric models.

3 A Two-Stage Model of Grasping

The utility of general-purpose, flexible manipulators versus specialized end-effectors is highly task dependent. One can easily argue for specialized end-effector tools when the task, the set of manipulations, and the objects to be manipulated are well known. In an environment which is not so highly structured, a single, flexible manipulator capable of carrying out a number of different tasks would seem to be more useful. Likewise, one might find arguments in support of both the analytical approach to robotic grasping and a more anthropomorphic approach which attempts to incorporate human techniques into the robotic system. Analytical methods have a firmer mathematical foundation. On the other hand, because the analysis can quickly become unwieldy, simplifying assumptions are often made. For example, grasps are often assumed to be planar. They are often modeled as point contacts or as idealized soft-finger contacts. Coulomb friction, or no friction at all, is assumed. The difficulties in controlling a multiple-degree-of-freedom manipulator capable of providing such grasps is not taken into account, nor are the inaccuracies of these devices or the inherent uncertainties in the environment itself. Finally, simplified object models are often assumed to be available. Obviously, the less structured the environment becomes, the less likely it is that these assumptions will fit the reality of the situation.

An alternate approach is to look at studies done by psychologists and cognitive scientists in human grasping. Based on these studies, we may develop a set of heuristics to help both in simplifying the synthesis and control of robotic grasps and in decreasing the number of limiting assumptions we must make about the world. Studies of human grasping provide a number of useful insights for the researcher in robotic grasping. For example, it has been noted that humans tend to use a predetermined set of hand configurations in the initial stages of a grasp [Nap56, Jea78, Arb83]. Often, several fingers are coupled, reducing the degrees of freedom of the system [Ibe87]. High-level knowledge about the task, the object to be grasped, and the perceived state of the world affect grasp choice and execution [Nap56, Kla86, Cut87]. And finally, sensory information is utilized at all stages of a grasp.

This research takes the second approach to implementing robotic grasp. We have built a knowledge-based system which incorporates many of the techniques observed in human manipulation. The system is based both on models of human grasping and of human perception. Perceptual capabilities allow the robot to adapt more readily to changes in its environment, whether these changes are caused by encountering new conditions in an unstructured environment or by a system or sensor malfunctioning within the robot itself. Reasoning provides the mechanism by which this flexibility may be realized, as well as providing a "hook" into the larger, intelligent system of which grasping and manipulation is but a subsystem.

The model of human grasping which we have used in designing our system was originally proposed by Jeannerod [Jea78] and later reexplored by Arbib [Arb83]. This model divides the grasp into two stages: the transport stage and the manipulation stage. In the transport stage, viewer-relative, extrinsic properties of the object such as spatial location and orientation are used to guide hand/wrist orientation and a ballistic reach toward the object. In the *manipulation* stage, *object-relative, intrinsic properties* such as shape and size are used to preshape the hand in anticipation of the grasp. Vision is used to extract the appropriate information and is utilized in an open-loop, feedforward manner during the preshaping and reach stages. We refer to preshape and reach as the precontact stage of grasping and it is this stage which we have implemented in the current research. The final, or postcontact, stage of a grasp occurs after the fingers are in contact with the object. In this stage, haptic information is used in a closed-loop, feedback manner to adjust and maintain the grasp during manipulation. We do not address the postcontact stage in this work.

4 Robotic Hand Preshaping and Reach

Studies of human grasping have been presented in both the robotics and medical literature. Common to all of these works is the idea that there is a fixed set of grasp configurations, or a grasp taxonomy, into which all grasps may be fit. Researchers differ, however, in their ideas about which parameters are used to choose particular instances of grasps and on the relative importance of these parameters. Taylor and Schwarz [Tay55] present a set of six grasps originally proposed by Schlesinger [Sch19]. In this taxonomy, object shape is the most important parameter used in selecting a grasp. A spherical object, for example, is grasped using a spherical grip. Jeannerod [Jea78] also suggests that shape and size are important parameters in choosing an initial grasp configuration. Napier [Nap56] places grasps into one of two categories: power grasps and precision grasps. Power grasps are used when stability is important, while precision grasps are used when dexterity is required. Napier believes that the task, rather than the shape of the object, is the important factor is choosing a grasp. He notes that in removing a lid from a jar, two different grasps are utilized: A power-type grasp is used to loosen the lid, a precision type grasp is then used to unscrew it. Cutkosky and Wright [Cut87] also point out the importance of task in grasp selection. Based on a study of grasps used by machinists in a small batch machine shop, they provide a taxonomy which further divides Napier's two categories and note that once a choice between a power or precision grasp has been made, object shape and task requirements become more equally weighted. Lyons [Lyo85] defines two functional indices for categorizing grasps, which are very similar to Napier's grasp taxonomy. The first index is firmness or power required. For example, the grasp required to use an object as a tool must be firmer than the one used simply to transport it. The second index is precision. The grasp used to insert a bolt requires more dexterity than the one that is used simply to lift it. Based on these indices, Lyons suggests three grasps with the appropriate functionality: the encompass grasp, the lateral grasp, and the precision grasp. Finally, Iberall [Ibe87] suggests that grasp posture is constrained by the way in which the hand can apply opposing forces around an object. These forces are defined by both the capabilities of the hand to create them and by the forces required to carry out the given task. She defines a set of three methods for attaining these forces: pad opposition, palm opposition, and side opposition. The chosen grasp reflects the use of one or more of these oppositions. Again, it would seem that object shape and size will affect the hand's ability to apply such oppositions and so will be equally as important as task in choosing a grasp.

Given the above studies, we propose the following ordering on these two factors: That intrinsic properties of the object such as shape and size are used to generate the set of all possible grasps for an object, and that task requirements are then used to choose one grasp from this set. Thus an object property is a necessary, but not sufficient condition for a grasp. For example, if object shape does not allow a precision grip of the object, then task requirements calling for this grasp are meaningless. If, on the other hand, shape allows several different possible grasps, then task requirements will dictate which is used. This, then, is our approach: we use object properties to generate a complete set of valid grasps for the object. The task and the state of the environment may then be used to prune this set.

We next address the problems of specifying the set of hand preshapes to be known to the system and of determining the parameters which will be used to define them. In choosing our preshapes, we have once again utilized the analyses of human grasping presented above. Hence, the set of three hand preshapes which we have implemented – wrap, pinch and grip – are similar to those proposed by Lyons. The defining parameters of a preshape are the number of virtual fingers which it uses and the type of contact which the fingers make with the object. The concept of virtual fingers is due to Iberall [Ibe87] and involves the coupling of one or more real fingers to apply the desired forces to an object. Since the coupled fingers move as a single unit, the control of the grasp is made simpler.

In addition to the defining parameters, we also specify adjustment parameters for a grasp. Adjustment parameters are used once a preshape has been invoked to adjust the grasp to the specific object. Currently, we implement two adjustment parameters within our system. Aperture adjustment increases the distance between the top and bottom fingers, while maintaining the grasp preshape, in order to insure that an object will fit into the grasp. Span adjustment increases or decreases the distance between the top two fingers to insure that all fingers make contact with the object and to maximize stability. We summarize our three grasps below.

Figure 1: Wrap hand preshape.

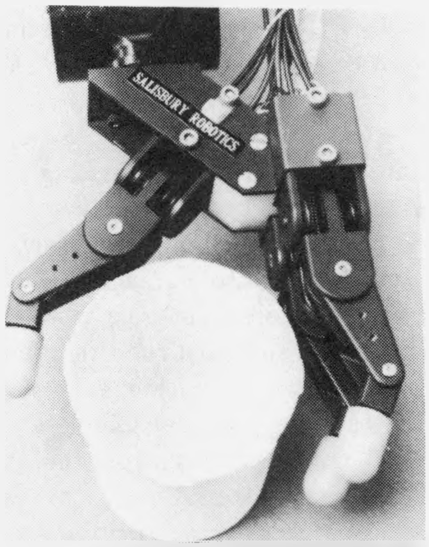

4.1 Wrap

The wrap grasp is a power grasp in Napier's taxonomy. The hand attempts to enclose the object. This grasp provides maximum stability, but minimum dexterity for further manipulation. The wrap utilizes three virtual fingers and contact with the object is extended (i.e. there are multiple areas of contact along the fingers.) Both aperture and span adjustments are executed. Figure 1 shows the wrap preshape as implemented on the Salisbury robot hand.

4.2 Grip

The grip is a precision grasp. Three virtual fingers are used and there is only one area of contact per finger. The object is held at the fingertips, so there is greater dexterity. Manipulation of the object with the fingers is also possible. Both aperture and span adjustments are performed for this grasp. Figure 2 shows the grip preshape as implemented on the Salisbury robot hand.

Figure 2: Grip hand preshape.

4.3 Pinch

The pinch grasp is somewhere between a power and a precision grasp. Two virtual fingers are used: The top two fingers are coupled to form a single unit which opposes the bottom finger. The object is held at the fingertips and there is only one area of contact per finger. Because the top two fingers are coupled, only aperture adjustment is performed. Figure 3 shows the pinch preshape as implemented on the Salisbury robot hand.

4.4 Wrist Orientation and Reach

The other stage of our grasp involves orientation of the wrist and a ballistic reach toward the object in order to place the fingers in the proper position to perform the grasp. There are three parameters for this stage. The target point is the point on the object above which the palm of the hand will be centered. The approach plane is the plane normal to the axis along which the hand will approach the object. Finally, the oppositions determine the orientation of the hand for finger placement. We define our oppositions as a set of planes. This is different from the way in which Iberall uses the term. The principle, however, is essentially the same. In Iberall's terminology, oppositions define the forces which the hand can apply around an object.

Figure 3: Pinch hand preshape.

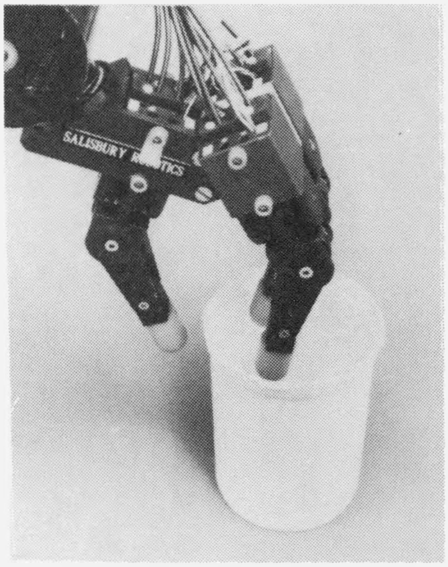

In our terminology, oppositions define finger placement, which in turn determines the forces which will be applied. The main difference is that our usage of the term is object-based, rather than manipulator-based.

5 Visual Perception

Visual information is obtained using a structured-lighting vision system. Hardware consists of a laser scanner and translation and rotation tables. This allows us to obtain scans of an object from several different views. The result of a scan is a set of three-dimensional points calibrated to the world space of the Puma robot. This data is processed by the visual perception system to create a representation of the sensed object which is utilized in generating and executing grasps of the object. The vision system, as we use it in this work, is passive. Visual processing of an object occurs at the beginning of an experiment. This information is then used in a feedforward manner by succeeding modules.

In Stansfield [Sta88c], we present a model for robotic perception based upon Fodor's [Fod83] proposed model of the human perceptual system. In this model, the perceptual system is organized into a hierarchy of problemsolving modules each of which is domain-specific and informationally-encapsulated. Processing within the system proceeds via the assignment of a set of intermediate levels of representation of the sensed world, beginning with low-level primitives and ending with an abstract, symbolic representation to be used by the cognitive system. It is the function of each module to either extract some feature or primitive or to process features from lower-level modules into more abstract representations. The final output of this system is an *apprehension* of the object being sensed. By apprehension we mean the identification of the features of an object and the relations among them.

We have based the structure of our visual perception system upon this model. The primitives of the system are the 3D points obtained by scanning an object. The final output is a symbolic representation consisting of the features of the object and their spatial relations. Hence, the system apprehends, rather than recognizes, the object to be grasped. We discuss both the representation and the processing below.

5.1 Object Representation

An important facet of our system is that it does not require geometric models of the objects to be grasped. Indeed, in unstructured environments, it is not reasonable to expect that we will be able to model all of the objects which the robot might encounter; nor can we assume, in such inherently noisy environments, that we will be able to gather sufficiently accurate sensor data to match these models even if they were available. What we have done instead is to provide the system with perceptual information in the form of a set of defining features loosely coupled spatially into a set of 2-1/2 dimensional views, or aspects.

The representation is currently created solely by the visual system, although it is easily extended to contain information from other perceptual systems such as touch. It is a feature-based, hierarchical representation. At the highest level is information about the perceived object as a whole. The next level contains descriptors for the components which define this object. The lowest level contains the features which parameterize these components. These features are organized into a set of view-dependent aspects which we refer to as the *aspect polyhedron.*

The concept of aspects is due to Koenderink [Koe79]. The idea is that all of the infinite 2-dimensional views of a 3-dimensional object can be grouped into a finite set of equivalence classes. These equivalence classes represent the aspects of the object. The aspect polyhedron is a collection of aspects for the sensed object created by taking multiple scans of the object at fixed positions and then extracting the object features for each scan. Informally, one might think of it as a set of projections of an object onto the faces of a polyhedron inside of which the object is centered. In Stansfield [Sta88b] we present this representational paradigm and show how it can be used for the tactile exploration and recognition of generic objects. In this work, we do not attempt to match the object to a model nor to attach a label to the object as a whole. Rather, we build the aspect polyhedron and use it directly to reason about and to drive manipulation. An example of an aspect polyhedron created by our system is shown later in the paper. First, we discuss the visual processing in more detail.

5.2 Visual Processing

We currently assume that we are dealing with a single, isolated object which may contain multiple components and features. A maximum of five aspects for the object may be obtained: top, front, back, right, and left. (The system requires only one aspect to function.) An aspect is obtained by scanning an object from a particular viewpoint. Once an aspect is obtained, the features and relations which it contains are extracted. Figure 4 shows the flow of processing for the vision data within an aspect. Aspects are related to each other via the aspect polyhedron, even when frames for all five faces are not constructed. After the initial sensing has been done and the aspect frames created, the system constructs the frames for the object and components based on all of the information available from the individual aspects.

Initial Segmentation The set of 3D points obtained from the structured-lighting system must be segmented and classified. We have chosen to do the segmentation using well-known reflectance-image processing techniques. The first step, after performing a scan of the object, is to create a 2D binary image from the 3D points. Standard region growing and segmenting techniques are then used to segment this 2D image into a set of 2D regions. Figures 5 - 7 show the results of this processing for three different aspects of an object having a hollow, cylindrical body and one part. The set of 3D range points is then segmented by associating each point with the region of which its mapped pixel is a member.

Classification of Components and Features The perceptual system extracts and synthesizes a fixed alphabet of primitives and features. This

Figure 4: Flow of control for visual processing.

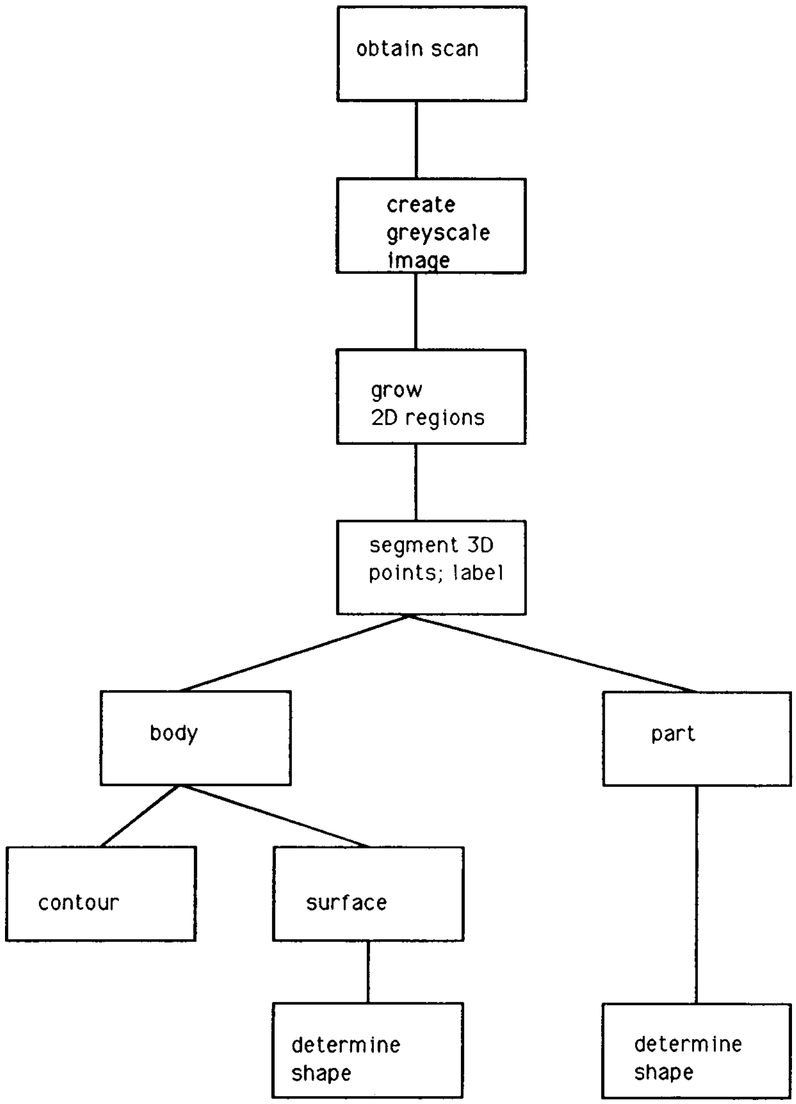

Figure 5: Top view of sensed object.

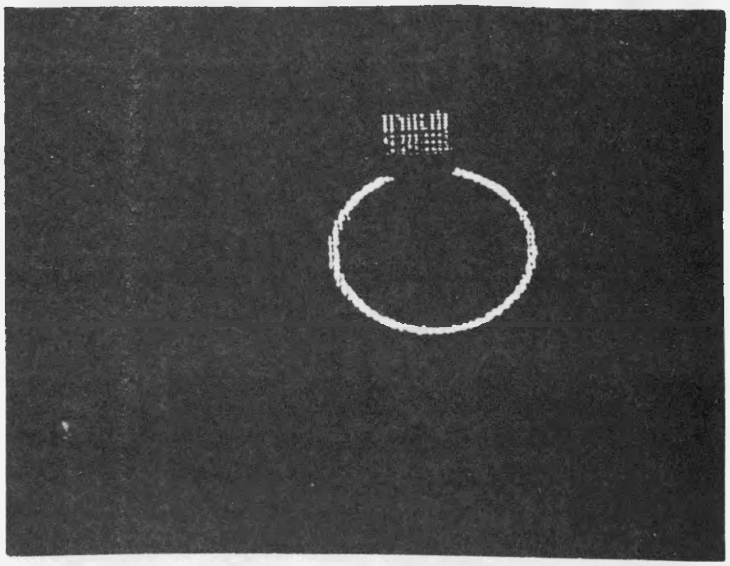

Figure 6: Front view of sensed object.

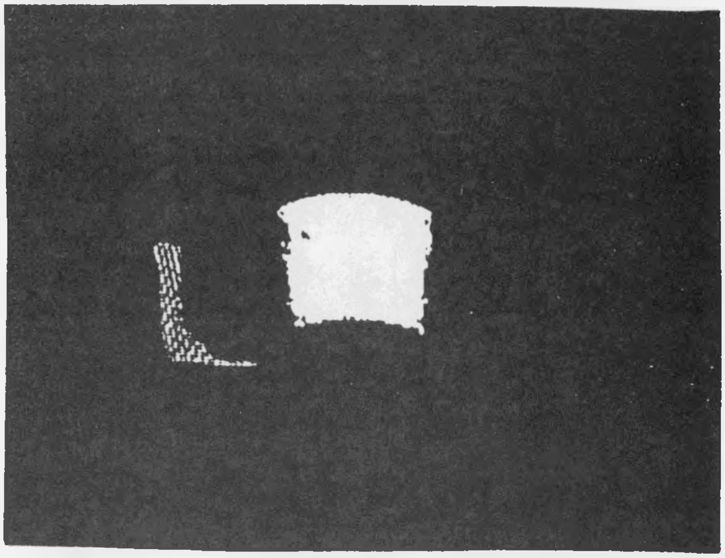

Figure 7: Right view of sensed object.

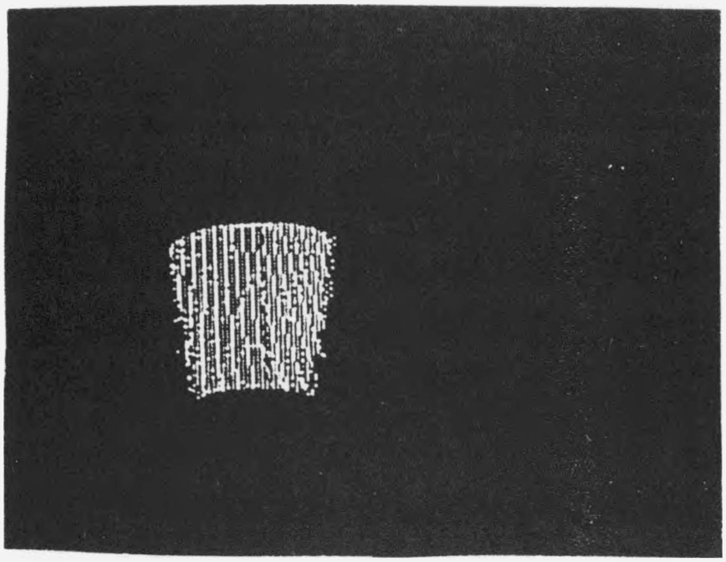

set is restricted primarily by the limitations of the devices which comprise the system. (Humans, for example, do not usually speak of the infrared properties of an object, since the human system is not capable of extracting this property.) Figure 8 summarizes the hierarchical definition of an object based upon the set of features currently extracted by our visual system. As we stated earlier, an object is composed of a set of components. These components are obtained from the set of segmented 3D range points. A component is labeled as either a body or a part. Each object may have only one component labeled body. A component is labeled based upon its size in 3D space and its area in 2D image space. The set of features which define the components are then identified and parameterized. A body component is currently composed of rim contours and planar or curved surfaces. The determination of whether a feature is a rim contour or a surface is based upon the ratio of its area to its extent in 2D space. If the feature is labeled as a surface, then its shape is obtained using the 3D points. If the component is labeled as a part, then it is parameterized by its shape. A one-extended part has dimension in only one direction, while a two-extended part has dimension in two directions.

Features are extracted from each aspect obtained for the object. The result is a hierarchy of frames containing the object, its components, and its

Figure 8: Components and features defining an object.

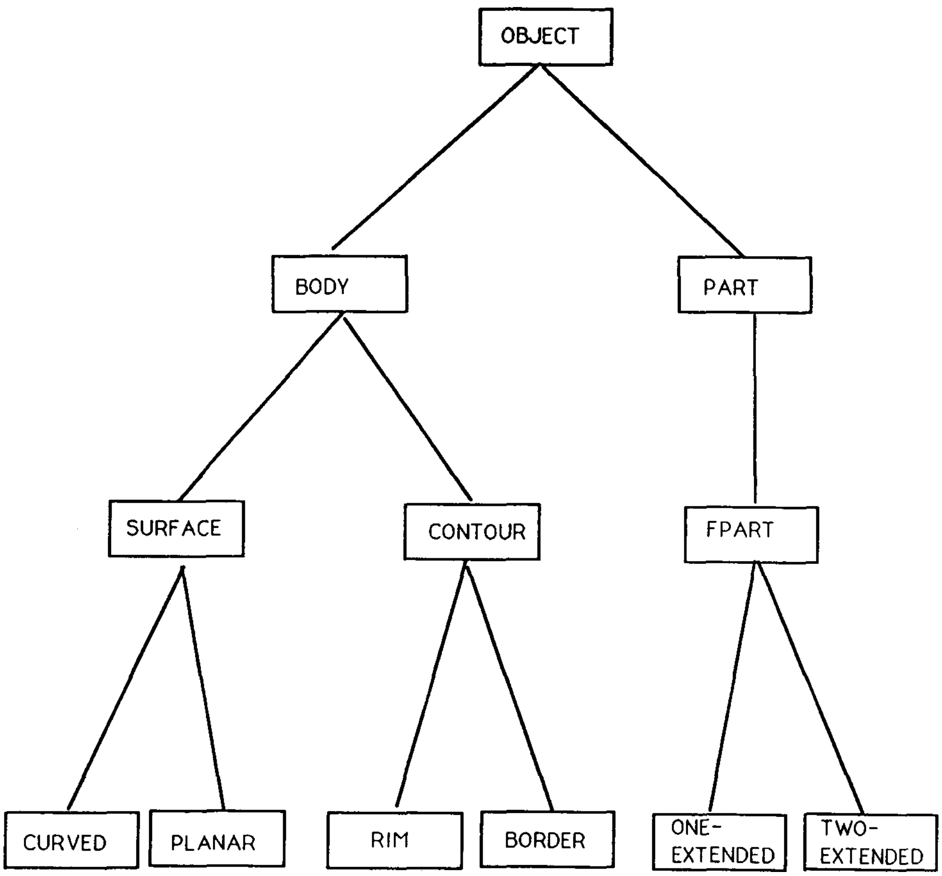

view-dependent features grouped according to the aspect of which each is a member. An object is parameterized by its spatial extent, its dimensionality (a very thin object has dimension 2, for example,) and its set of components. A component is parameterized by its spatial extent and type. A feature is parameterized by its type and shape. A set of spatial relations (above, below, left-of, and right-of) is also extracted for the components. Figure 9 shows the final output of the vision system for the object in Figures 5 - 7.

We would like to conclude our discussion of the visual processing stage with an observation. First, while researchers in machine vision are making daily strides forward, the state of the art is still very primitive. Our system is capable of functioning even with a very limited visual capability. Second, the structure for perception which we have set forth here will allow us to incorporate not only improved visual processing, but also haptic and other forms of sensory input as they become available.

6 Reasoning for Grasp Generation

Once the object has been visually perceived and the aspect representation has been built, a rule-based expert system is invoked to generated a set of valid grasps for the object. The system is implemented in Prolog. The rules operate on the aspect polyhedron created by the perceptual system and embody a knowledge of which features must be present, and what the relations among these features must be, in order for a particular hand preshape and reach to be valid. This simple set of heuristics for grasping objects forms the heart of our system. It also provides the system with both power and flexibility. From information as sparse as a single view of an object, and without the necessity of matching to a stored model, we are able to generate grasps for the object.

Because we are dealing with unstructured environments, we do not impose the condition that an object must be recognized before it can be grasped. Indeed, such a condition would, in general, be quite limiting – imagine having to know what an object was before you could handle it! In unstructured environments in particular, we cannot expect to model all of the objects which the robot might encounter. In addition, we must expect that sensor errors and objects which deviate from the model will be regularly encountered. And finally, to match an object to a single hypothesis may require more initial sensing than is currently necessary or may be possible.

Figure 9: Final visual perception of the object.

```
object frame
enclosing volume: [81,122,98]
dimension: 3
components: [body part]

component frame
component: body (255)
enclosing volume: [81,80,81]

contour frame
component: body (255)
view is top
type is rim
shape is undetermined

surface frame
component: body (255)
view is front
shape is curved

surface frame
component: body (255)
view is back
shape is curved

surface frame
component: body (255)
view is right
shape is curved

component frame
component: part (128)
enclosing volume: [37,43,98]

part frame
component: part (128)
view is top
shape is one_extended

part frame
component: part (128)
view is front
shape is two_extended

part frame
component: part (128)
view is back
shape is two_extended
```

In this work, we show that it is not necessary for the robot to recognize an object in order to grasp it. It need only apprehend that object. In Stansfield [Sta88a], we present a system which attempts to recognize the object before generating grasps. A benefit of recognition is that we may instantiate a full model of the object to be used in further reasoning. For example, from the model we would have knowledge of unsensed portions of the object. We could then generate grasps which place fingers on these portions. Our system currently does not place fingers on unsensed parts of an object.

What we foresee in the future is a hybrid system. The robot will carry around models for a small set of objects which are part of its personal domain – a set of tools for example. This will allow the robot to identify a tool if it is misplaced during use and needs to be retrieved. In addition, our frame-based representation will allow us to store other, non-perceptual, information about these objects, such as their use, proper storage position, etc. The majority of the grasping work, however, will be carried out on objects which are part of the environment, rather than the robot's personal domain. The robot must be able to grasp these objects without recognizing them.

The rules which the system utilizes are quite simple. Intrinsic properties of the object, such as its size and the set of features which comprise it, are used to generate the hand preshape. For example, a curved object – represented as a set of curved surfaces – is grasped using a wrap grasp. An object with a planar surface or a rim may be grasped using a grip grasp. And a flat, thin object may be grasped between the fingers in a pinch grasp. Extrinsic properties of the object, such as its location and orientation, are used to generate the reach. The set of rules is completely device independent. The only requirements are that the hand be capable of executing the desired preshape and that the span of the hand be known. In addition, the system can generate grasps for objects given any number of aspects between the minimum of one and the maximum of five. The lack of aspects, and hence of information about the object, only causes fewer grasps to be generated.

6.1 Example Rules

How these rules operate is perhaps best illustrated by a set of examples. Figure 10 shows the psuedo-code rule for invoking a pinch grasp of a rim feature with reach from above. This is a simple grasp and there is only one condition: The object must have a rim contour in the aspect labeled top.

Figure 10: Rule for a pinch grasp of a rim contour.

```
Pinch_Top_Body(Object):-
  If aspect labeled Top has Rim_Contour
     THEN valid preshape is Pinch for component Body
          approach Body from Above
          target is Point on Rim_Contour
          oppositions are Point Inside Rim_Contour
                      and Point Outside Rim_Contour.
```

Figure 11: Rule for a grip grasp of a rim or border contour.

```
Grip_Right_Body(Object):-
  If aspect labeled Right has Rim_Contour OR
     aspect labeled Right has Planar_Surface AND
     object fits into span of hand
     THEN valid preshape is Grip for component Body
          approach Body from the Right
          target is center of Contour
          oppositions are Contour from Front
                      and Contour from Back.
```

A pinch grasp is the most easily executed: Fingers are placed on either side of the rim and the grasp is closed. The approach, target, and oppositions define the wrist orientation and reach.

Figure 11 shows the rule for invoking a grip grasp of a contour. Contours are created either by rim features or by the borders of planar surfaces. In this case, the hand spans the object at the position of the contour and so a check is made that the object fits into the hand in the proper dimension.

The wrap grasp is the most complex grasp. This is because the hand attempts to enclose the object, creating multiple points of contact along each finger. It is in this configuration that collision with other components not being grasped is most likely to occur. Figure 12 shows the rule for invoking a wrap grasp of a curved object reaching from the left. The conditions are that there be a set of three adjacent curved surfaces and that the left aspect have no other components. Recall that it is the left aspect which defines

Figure 12: Rule for a wrap grasp of a curved surface.

```
Wrap_Left_Body(Object):-
  If aspect labeled Left has Curved_surface AND
     aspect labeled Left contains no other features AND
     aspect labeled Front has Curved_Surface AND
     aspect labeled Back has Curved_Surface AND
     object fits into span of hand
     THEN valid preshape is Wrap for component Body
          approach Body from the Left
          target is center of Curved_Surface
          oppositions are Curved_Surface from Front
                      and Curved_surface from Back.
```

Figure 13: Rule for a wrap grasp of a two-extended part.

```
Wrap_Front_Part_2E_Left(Object,Part):-
  If aspect labeled Front has Two_Extended Part AND
     Part is Left_Of Body
     THEN valid preshape is Wrap for component Part
          approach Part from the Left
          target is center of Part
          oppositions are Part from Front
                      and Part from Back.
```

the approach and placement of the palm. Other components of the object which show up in this aspect will be in the way of the desired grasp of the curved surface.

The three rules discussed above have been invoked by different features of a body component. Figure 13 shows the rule for a wrap grasp of a part component. In this case, the part has a two-extended shape in the aspect labeled front and is to the left of the main body component of the object. The reach is generated so that the palm approaches the part from the left and the fingers wrap around it from the front and back.

6.2 Example of a Generated Grasp Set

Figure 9 in Section 5 shows the final results of the visual perception of an object having multiple components and a variety of features. The object is a hollow, cylindrical tube with a two-extended part. The object is oriented, relative to the robot, so that the rim of the tube is facing up and the part is to the left of the body. Figure 14 shows the set of grasps generated for this object. A wrap grasp has been generated for the curved body of the object with approach from the right. No other wrap grasp for the body is generated because the part would be in the way of the fingers. Three grasps have been generated for the rim contour. The pinch grasp will place at least one finger inside of the tube. Two different grip grasps are generated, placing the hand/wrist in two different orientations. All fingers are on the outside of the rim. And finally, a wrap grasp of the part is generated with approach from the left. A wrap grasp of the part from above is not generated, because the part is not large enough in that aspect.

7 Driving the Robot Hand and Arm

Figure 15 summarizes the flow of control for the entire system. The object is visually perceived and its defining features, along with their spatial relations, are extracted. The result is an apprehension of the object which is represented as a hierarchy of frames and a set of aspects. This symbolic representation is input to a rule-based reasoning system which generates a set of valid grasps (hand preshape and reach parameters) for the object. Each one of these grasps may then be input to the lower-level motor control modules and used to drive the robot hand and arm in executing the grasp.

The hardware used to implement the haptic system consists of a six-degree-of-freedom PUMA 560 robot arm, currently under position control, and a Salisbury [Sal85] robot hand which is mounted on the wrist of the PUMA. The hand has three fingers, each with three joints, for a total of 9 degrees-of-freedom. The hand is also position controlled, with a guarded move implemented using the tendon tensions.

There are three distinct stages to a grasp: hand preshaping, preshape adjust, and wrist orientation and reach. For ease of execution, we have further divided the preshape adjust stage into two substages, aperture adjust and span adjust. Figures 18 - 20 illustrate the entire grasping sequence. The object to be grasped is a a rectangular polyhedron with an irregular hole cut out of it. Figure 16 shows the results of the visual perception of this

Figure 14: Set of grasps generated for cylindrical tube with two-extended part.

```
Use hand preshape Wrap for component Body
Approach target is the center of the curved
  surface from the right
Opposition 1: curved surface from the front
Opposition 2: curved surface from the back

Use hand preshape Pinch for component Body
Approach target is point on rim from top
Opposition 1: inside of chosen rim point
Opposition 2: outside of chosen rim point

Use hand preshape Grip for component Body
Approach target is center of contour from top
Opposition 1: contour from left
Opposition 2: contour from right

Use hand preshape Grip for component Body
Approach target is center of contour from top
Opposition 1: contour from front
Opposition 2: contour from back

Use hand preshape Wrap for component Part
Approach target is the center of the part
   from the left
Opposition 1: part from the front
Opposition 2: part from the back
```

Figure 15: Overall flow of control for the grasping system.

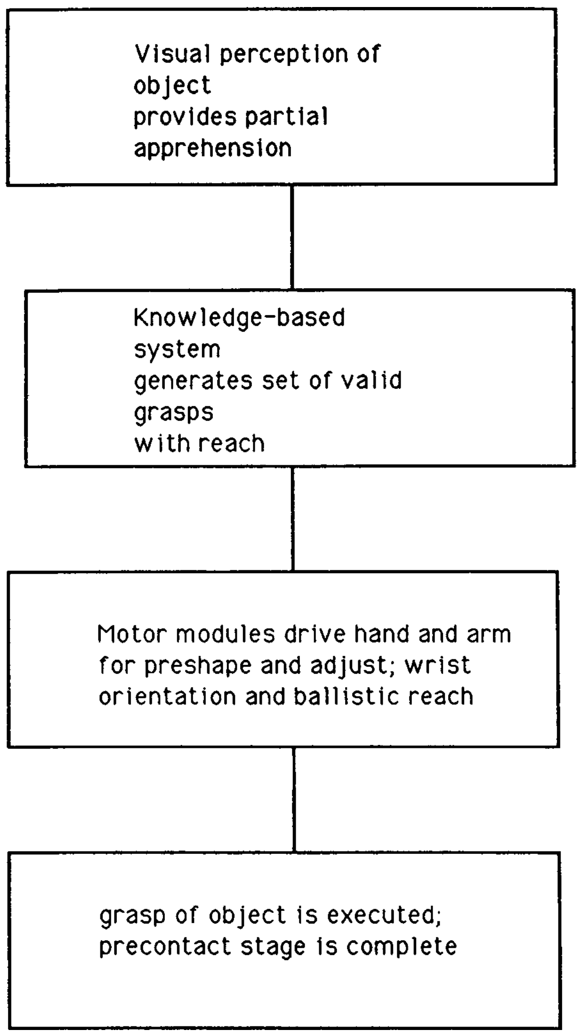

Figure 16: Results of visual processing for polyhedron with irregular cut-out.

```
object frame
enclosing volume: [139,232,79]
dimension: 3
components: [body]

component frame
component: body (255)
enclosing volume: [139,232,79]

contour frame
component: body (255)
view is top
type is rim
shape is undetermined
```

object (only one aspect was created,) while Figure 17 shows the set of grasps generated by the expert system. Note that only one grip grasp for the object from above has been generated, since the object is too large to fit into the hand in any other orientation. A pinch grasp of the rim is also generated. We use the execution of the grip grasp to illustrate all of the stages of a grasp.

Recall that the information provided by the knowledge-based system consists of the following: A preshape, a target point and approach plane, and a set of oppositions. The motor modules also have access to the symbolic perceptual representation which contains the bounding volumes for the object and its components, as well as the aspect polyhedron. At the motor level, the system also retains information concerning the positions of these features in the world space of the robot (at the symbolic reasoning level it is not necessary for the system to have access to this information.) In the first step of the grasping sequence, the desired preshape is executed on the robot hand. In this case, it is a grip preshape shown in Figure 18. In the next step, the aperture of the preshape is widened to insure that the object will fit into the span of the hand. A transform is then built to drive the wrist orientation and reach of the arm. The approach target provides us with the (x,y,z) components of this transform, while the approach plane

Figure 17: Grasp set generated for polyhedron with cut-out.

```
Use hand preshape Pinch for component Body
Approach target is point on rim from top
Opposition 1: inside of chosen rim point
Opposition 2: outside of chosen rim point

Use hand preshape Grip for component Body
Approach target is center of contour from top
Opposition 1: contour from front
Opposition 2: contour from back
```

and oppositions are used to calculate the Euler angles. The reach is then performed. These steps are illustrated by Figure 19, which shows the hand and arm after the preshape aperture has been adjusted, as the arm begins the wrist orientation and reach. Finally, the span adjustment of the top two fingers is then carried out and the grasp is executed. The result is shown in Figure 20.

8 Experimental Results

In this section we present a series of experimental results which illustrate the functioning system. We have exercised all of the features of the system by presenting a diverse set of objects to be perceived and grasped. Thus we have used objects with planar and curved surfaces in various combinations, objects having various rims, objects having multiple components of different shapes, and objects of both three and two dimensions. In addition, by varying the numbers of aspects extracted for different objects, we show that the system is capable of working with varying degrees of partial information.

8.1 Experiment I: Rim Contours

Figures 21 and Figure 22 show the results of executing pinch grasps of two different rim contours. The tube in Figure 21 contains a smoothly curved rim surrounded by curved surfaces. The polyhedral object in Figure 22 is the one used to illustrate the examples in Section 7. In the case of this

Figure 18: Initial grip hand shape.

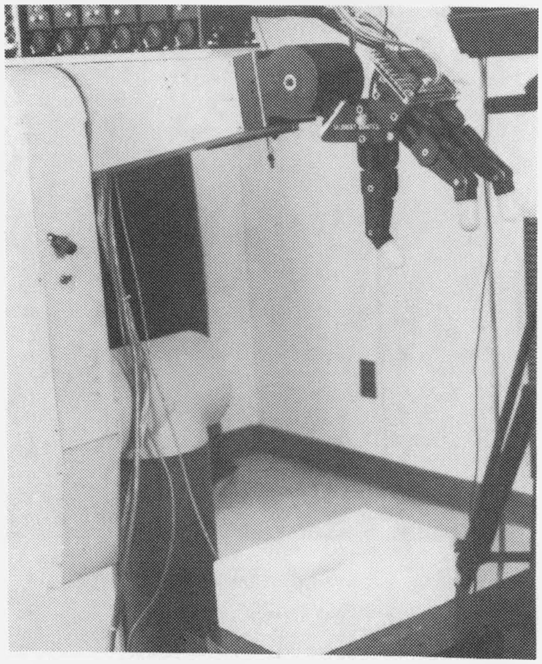

Figure 19: Aperture adjust, wrist orientation and reach.

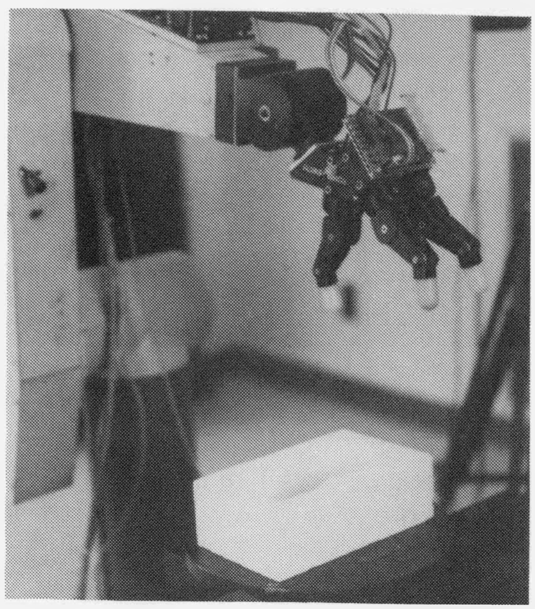

Figure 20: Span adjust and final grasp.

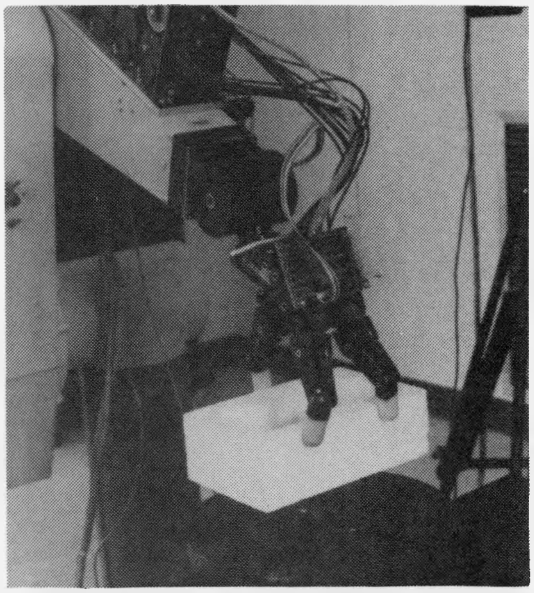

object, the rim is irregularly curved and is surrounded by planar surfaces. The system is capable of handling both cases.

8.2 Experiment II: Freeform Object

This set of experiments shows multiple grasps of a freeform object containing both curved and planar surfaces. Figure 23 shows the output of the visual sensing for this object (note that four aspects were obtained), while figure 24 shows the set of grasps generated for the object by the knowledge-based system. Figure 25 shows the execution of the wrap grasp of the object from above. Figure 26 shows the execution of the grip grasp of the object's planar surface from the left.

8.3 Experiment III: Ellipsoidal Object

We have thus far shown that the system is capable of grasping 3D objects containing planar surfaces and a combination of planar and curved surfaces. This experiment shows the grasp of an ellipsoidal object containing only curved surfaces. In this case, only wrap grasps of the object are generated. Figure 27 shows the execution of a wrap of the object from above. This is

Figure 21: Pinch grasp of curved tube with curved rim contour.

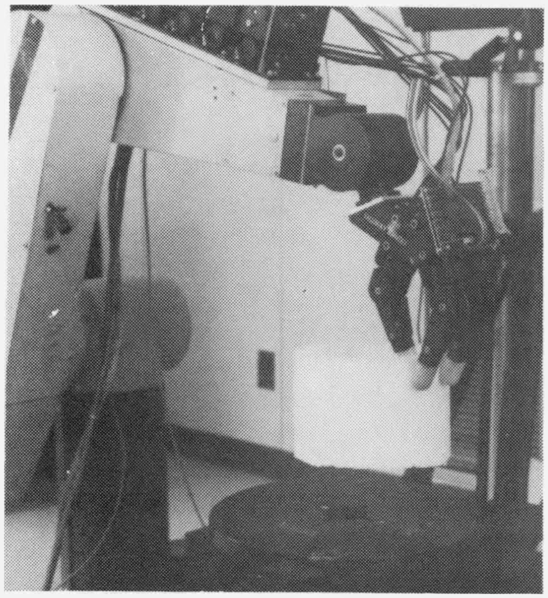

Figure 22: Pinch grasp of polyhedron with irregular rim contour.

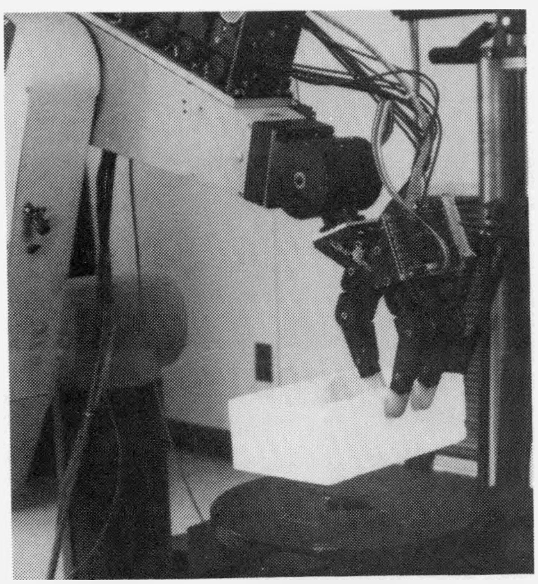

Figure 23: Results of visual processing for freeform.

```
object frame
enclosing volume: [106,95,150]
dimension: 3
components: [body]

component frame
component: body (255)
enclosing volume: [106,95,150]

surface frame
component: body (255)
view is top
shape is curved

surface frame
component: body (255)
view is front
shape is curved

surface frame
component: body (255)
view is back
shape is curved

surface frame
component: body (255)
view is left
shape is planar
```

Figure 24: Grasp set generated for freeform.

```
Use hand preshape Wrap for component Body
Approach target is the center of the curved
   surface from the top
Opposition 1: curved surface from the front
Opposition 2: curved surface from the back

Use hand preshape Grip for component Body
Approach target is center of contour from left
Opposition 1: contour from front
Opposition 2: contour from back
```

Figure 25: Wrap grasp of freeform from above.

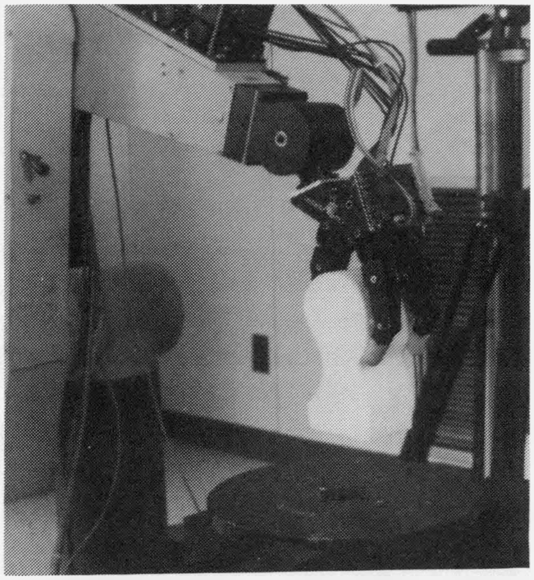

Figure 26: Grip grasp of freeform from the left.

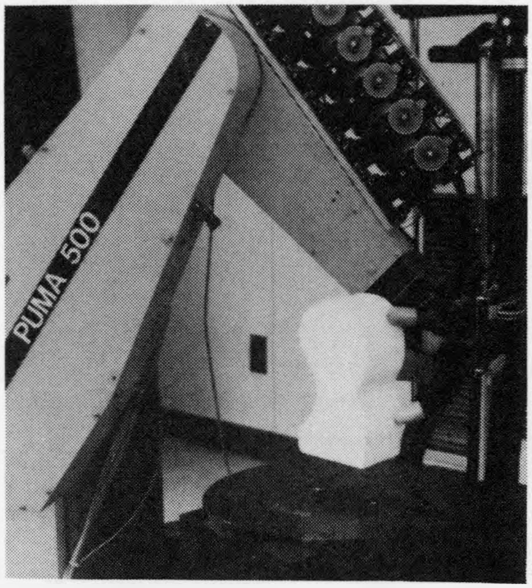

the most difficult grasp, because the shape of the object does not allow it to be enclosed from this direction.

8.4 Experiments IV and V: Objects with Multiple Components

The next set of experiments show multiple grasps of objects having multiple components.

Experiment IV: Object with One-Extended Part Figure 28 shows the results of the visual processing for a solid, cylindrical object having a single part whose shape is one-extended. Figure 29 shows the set of grasps generated for this object. The system generates a wrap grasp of the part from above. This is shown in Figure 30. A grip grasp of the cylindrical body from above is also generated. This is shown in Figure 31. The system does not generate a wrap of the body, because it has not created aspects for the front and back of the object – the system does not attempt to place fingers on unsensed portions of an object.

Figure 27: Wrap grasp of ellipsoidal object from above.

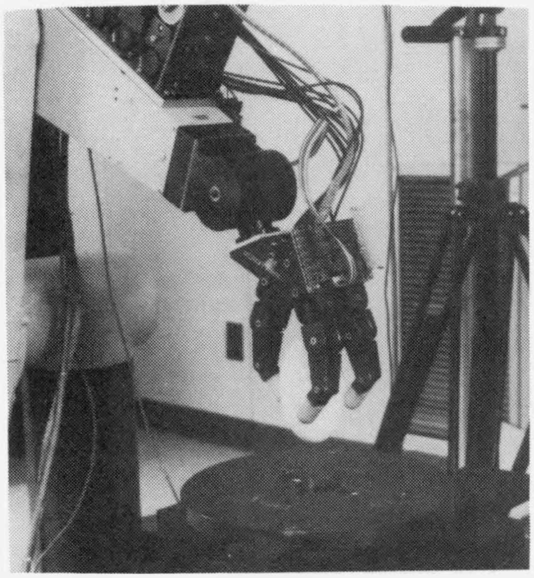

Experiment V: Object with Two-Extended Part Figure 32 shows the results of the visual processing for a solid, polyhedral object having a single part whose shape is two-extended. Figure 33 shows the set of grasps generated for this object. Figure 34 shows the wrap grasp of the part from the left.

8.5 Experiment VI: Small Object

This set of experiments shows that the system can also deal with small (relative to the hand,) irregularly shaped objects. Figure 35 shows the results of the visual processing for a gear shaped object standing on edge. Figure 36 shows the set of grasps generated for the object. The pinch grasp of the gear from above is shown in Figure 37. The object was also sensed laying flat. (The object is placed on a base to allow the hand to be positioned without causing the fingers to collide with the table.) Figure 38 shows a grip grasp of the gear, in this configuration, from above.

Figure 28: Results of visual processing of cylinder with one-extended part.

```
object frame
enclosing volume: [125,247,45]
dimension: 3
components: [body part]

component frame                  component frame
component: body (128)            component: part (255)
enclosing volume: [125,98,11]    enclosing volume: [33,77,4]

surface frame                     part frame
component: body (128)             component: part (255)
view is top                       view is top
shape is planar                   shape is one_extended

surface frame
component: body (128)
view is left
shape is curved
```

Figure 29: Grasp set generated for cylindrical object with one-extended part.

```
Use hand preshape Grip for component Body
Approach target is center of contour from top
Opposition 1: contour from left
Opposition 2: contour from right

Use hand preshape Grip for component Body
Approach target is center of contour from top
Opposition 1: contour from front
Opposition 2: contour from back

Use hand preshape Wrap for component Part
Approach target is the center of the part
   from top
Opposition 1: part from the front
Opposition 2: part from the back
```

Figure 30: Wrap grasp of one-extended part from above.

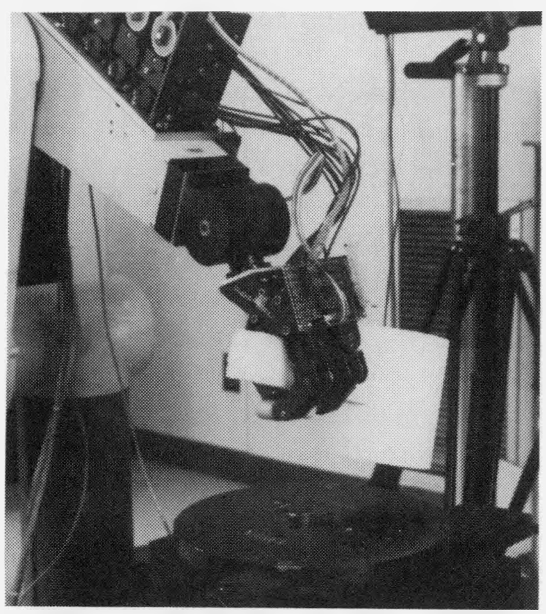

Figure 31: Grip grasp of cylindrical body from above.

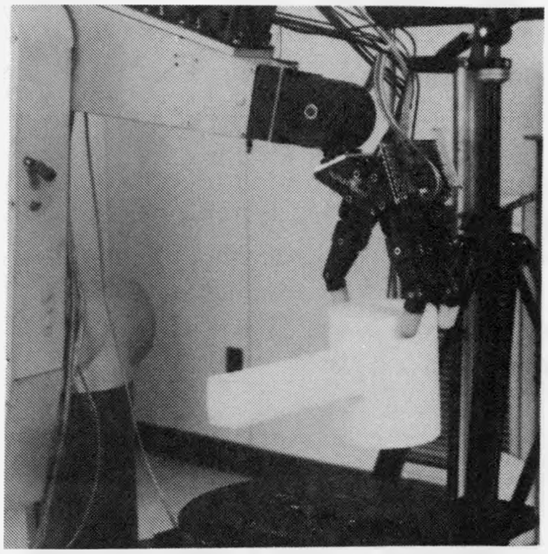

9 Summary and Discussion

In this paper, we have presented a general-purpose robotic grasping system. The system implements a two-stage model of human grasping in which the grasp is divided into a hand shaping phase and a reach phase. The system does not require models of the objects to be grasped. Rather, we have integrated perception and knowledge into the grasping task to allow the robot to deal with unknown objects and sparse sensor data. Thus, the object to be grasped is first perceived visually and a representation is built which consists of the features of the object and their spatial relations in the form of a set of aspects. This symbolic representation is then utilized by a rule-based expert system to generate a set of valid grasps for the object. The knowledge embodied in this system concerns which features must be present, and what the relations among these features must be, in order for a particular grasp to be valid. The information generated by the expert system is used to drive a robot arm and hand in executing a grasp of the object. We consider this work to be only a first step toward the development of a truly versatile, intelligent robotic grasping system. We have presented a set of experiments which show that the system, as it is, is capable of handling a variety of objects. Extensions are possible at all stages to make

Figure 32: Results of visual processing of polyhedron with two-extended part.

```
object frame
enclosing volume: [87,230,154]
dimension: 3
components: [body part]

component frame
component: body (255)
enclosing volume: [87,99,136]

surface frame
component: body (255)
view is top
shape is planar

surface frame
component: body (255)
view is back
shape is planar

surface frame
object: obj
component: body (255)
view is right
shape is planar
```

```
component frame
component: part (128)
enclosing volume: [32,97,144]

part frame
component: part (128)
view is top
shape is one_extended

part frame
component: part (128)
view is back
shape is two_extended
```

Figure 33: Grasp set generated for polyhedron with two-extended part.

```
Use hand preshape Grip for component Body
Approach target is center of contour from back
Opposition 1: contour from left
Opposition 2: contour from right

Use hand preshape Grip for component Body
Approach target is center of contour from right
Opposition 1: contour from front
Opposition 2: contour from back

Use hand preshape Grip for component Body
Approach target is center of contour from top
Opposition 1: contour from front
Opposition 2: contour from back

Use hand preshape Grip for component Body
Approach target is center of contour from top
Opposition 1: contour from left
Opposition 2: contour from right

Use hand preshape Wrap for component Part
Approach target is the center of the part
   from the left
Opposition 1: part from the front
Opposition 2: part from the back

Use hand preshape Wrap for component Part
Approach target is the center of the part
   from the top
Opposition 1: part from the front
Opposition 2: part from the back
```

Figure 34: Wrap grasp of two-extended part from the left.

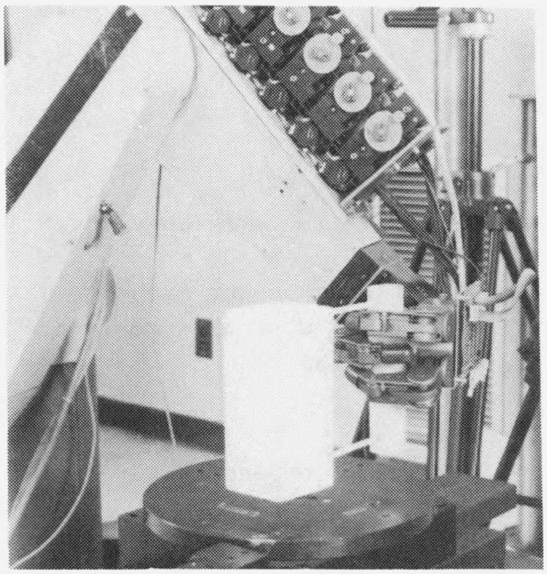

Figure 35: Results of visual processing for gear.

```
object frame
enclosing volume: [27, 83,32]
dimension: 2
components: [body]

component frame
component: body (255)
enclosing volume: [27,83,32]

surface frame
component: body (255)
view is front
shape is planar
```

Figure 36: Grasp set generated for gear.

```
Use hand preshape Grip for component Body
Approach target is center of contour from front
Opposition 1: contour from left
Opposition 2: contour from right

Use hand preshape Pinch for component Body
Approach target is body from top
Opposition 1: body from front
Opposition 2: body from back

Use hand preshape Pinch for component Body
Approach target is body from left
Opposition 1: body from front
Opposition 2: body from back

Use hand preshape Pinch for component Body
Approach target is body from right
Opposition 1: body from front
Opposition 2: body from back
```

Figure 37: Pinch grasp of standing gear from above.

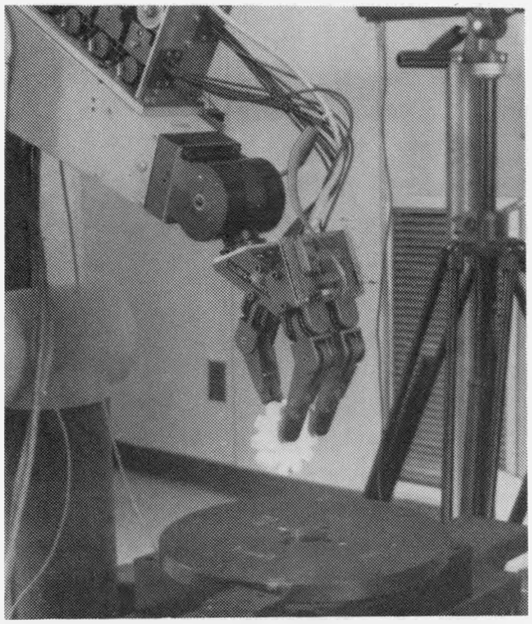

Figure 38: Grip grasp of repositioned gear from above.

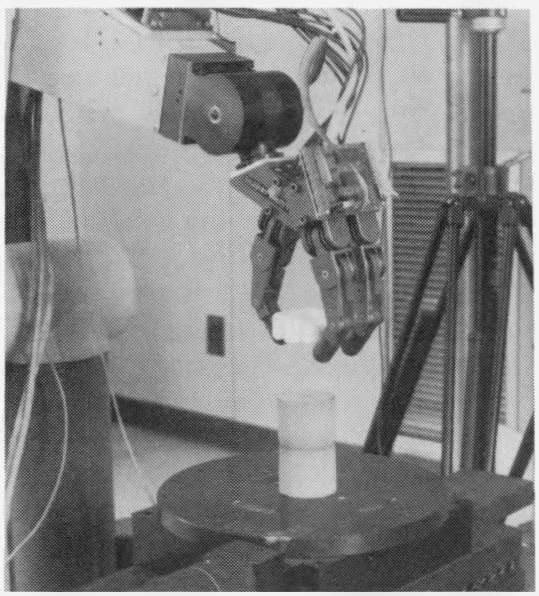

the system more powerful.

At the perceptual level, a better visual system will allow us to extract more and better initial information about the object. We do not consider our work to be computer vision research, and so we have developed the minimum vision system necessary. The model is easily extended to include better visual processing as this becomes available. The area of robotic perception which we do hope to incorporate as future work, however, is the development of a haptic perception system. This will be carried out in two phases. The first step will be the integration of tactile exploration into the initial perceptual stage. The haptic system is capable of extracting information, such as weight and texture, which the visual system is not. Yet such information is extremely useful in executing and maintaining a grasp. An initial haptic exploration of the object will provide more information to the system to be used in choosing grasps. The second phase of the haptic research will involve implementing the postcontact stage of the grasp. In this stage, haptic information about manipulation forces and slip are used to dynamically adjust and maintain the grasp during task execution. The same perceptual system will be used as in the haptic exploration phase, but the processing will no doubt be very different. We currently have a tactile sensor mounted on the robot hand and are in the process of designing the haptic perception system.

We will also be extending the knowledge-based system in several ways. First, both the representation and the reasoning must be further developed to allow more complex objects to be apprehended and grasped. The set of hand preshapes may also be enlarged. In addition, the system is easily extended to incorporate rules for pruning the set of grasps down to the one which will be executed. As we stated earlier in this paper, the task and the state of the world must be taken into account. Rules which select a grasp based on the power or dexterity required can be added, as well as rules which embody a knowledge of how to grasp objects which are to be used as tools. For example, there might be a rule which states that if the task is to brush the soil from an object, then the whisk must be grasped by the handle. Also, in a cluttered environment, the placement of other objects must be taken into account. Certain grasps will not be possible if there is another object occluding the approach of the hand. In the rule-based system as we have structured it, such extensions will not be difficult to incorporate.

Finally, limitations are placed on the system by the current state of development of the robotic devices. It is our hope that an additional benefit of this research will be to allow us to develop a cohesive set of specifications for the next generation of robot hands. What we have shown, we believe, is the utility of integrating perception and knowledge into the grasping task.

10 Acknowledgements

This work was performed at Sandia National Laboratories and was supported by the U.S. Department of Energy under Contract DE - AC04 - 76DP00789. We gratefully acknowledge the contributions of Greg Starr, who designed the controller for the Salisbury robot hand and of William Davidson and Peter Boissiere, who developed the integrated robotic system used in this work.

References

[Arb83] Arbib, M.,Iberall, T. and Lyons, D. *Coordinated Control Programs for Movements of the Hand.* Technical Report COINS 83-25, CIS Dept., UMass Amherst, Amherst, Ma., August 1983.

[Cut87] Cutkosky, M. and Wright, P. Modeling manufacturing grips and correlations with the design of robotic hands. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1533–1539, Raleigh, NC, April 1987.

[Cut88] Cutkosky, M. *On Grasp Choice, Grasp Models and the Design of Hands for Manufacturing Tasks.* Technical Report CDR 4.8.88, Stanford University, 1988.

[Fea86] Fearing, R. Simplified grasping and manipulation with dextrous robot hands. *IEEE Journal of Robotics and Automation*, RA-2(4):188–195, 1986.

[Fod83] Fodor, J. *The Modularity of Mind.* MIT Press, Cambridge, Ma., 1983.

[Han77] Hanafusa,H. and Asada, H. Stable prehension by a robot hand with elastic fingers. In *Proceedings of the 7th ISIR*, pages 361–368, Tokyo, Japan, October 1977.

[Ibe87] Iberall, T. The nature of human prehension: three dextrous hands in one. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 396–401, Raleigh, NC, April 1987.

[Ibe88] Iberall, T., Jackson, J., Labbe, L., and Zampano, R. Knowledge-based prehension: capturing human dexterity. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 82–87, Philadelphia, PA, April 1988.

[Ike86] Ikeuchi, K., Nishihara, H., Horn, B., Sobalvarro, P., and Nagata, S. Determining grasp configurations using photometric stereo and the prism binocular stereo system. *International Journal of Robotics Research*, 5(1):46–65, 1986.

[Jea78] Jeannerod, M. Intersegmental coordination during reaching at natural visual objects. In Long, J. and Baddley, A., editor, *Attention and Performance, IX*, Lawrence Erlbaum, Hillsdale, N.J., 1978.

[Kla86] Klatzky, R., McCloskey, M., Doherty, S., Pellegrino, J., and Smith, T. *Hand Shaping and Object Processing.* Technical Report Cognitive Science 8605, University of California, Santa Barbara, 1986.

[Koe79] Koenderink, J. and van Doorn, A. The internal representation of solid shape with respect to vision. *Biological Cybernetics*, 32:211–216, 1979.

[Li 87] Li, Z. and Sastry, S. Task oriented optimal grasping by multi-fingered robot hands. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 389–394, Raleigh, NC, April 1987.

[Loz87] Lozano-Perez, T., Jones, J., Mazer, E., O'Donnell, P., Grimson, W., Tournassoud, P., and Lanusse, A. Handey: a robot system that recognizes, plans, and manipulates. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 843–849, Raleigh, NC, April 1987.

[Lyo85] Lyons, D. A simple set of grasps for a dextrous hand. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 588–593, St. Louis, MO., April 1985.

[Mas85] Mason, M. Manipulator grasping and pushing operations. In *Robot Hands and the Mechanics of Manipulation*, MIT Press, Cambridge, MA, 1985.

[Nap56] Napier, J. The prehensile movements of the human hand. *Journal of Bone and Joint Surgery*, 38B(4):903–913, 1956.

[Ngu88] Nguyen, V-D. Constructing force-closure grasps. *International Journal of Robotics Research*, 7(3):3–16, 1988.

[Rao88] Rao, K., Medioni, G., Liu, H., and Bekey, G. Robot hand-eye coordination: shape description and grasping. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 407–411, Philadelphia, PA, April 1988.

[Sal85] Salisbury, J. Kinematic and force analysis of articulated hands. In *Robot Hands and the Mechanics of Manipulation*, MIT Press, Cambridge, MA, 1985.

[Sch19] Schlesinger, G. Der mechanische aufbau der kunstlichen glieber. In Borchardt et al., editor, *Ersatzglieder und Arbeitshilfen*, Springer, Berlin, 1919.

[Sta88a] Stansfield, S. Reasoning about grasping. In *Proceedings of the AAAI-88 Conference*, pages 768–773, St. Paul, MN, August 1988.

[Sta88b] Stansfield, S. Representing generic objects for exploration and identification. In *Proceedings of the IEEE Robotics and Automation Conference*, pages 1090–1095, Philadelphia, PA, April 1988.

[Sta88c] Stansfield, S. A robotic perceptual system utilizing passive vision and active touch. *International Journal of Robotics Research*, 7(6):138–161, 1988.

[Tay55] Taylor, C. and Schwarz, R. The anatomy of mechanics of the human hand. *Artificial Limbs*, 2:22–35, May 1955.

[Tom87] Tomovic, R., Bekey, G., and Karplus, W. A strategy for grasp synthesis with multifingered robot hands. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 83–89, Raleigh, NC, April 1987.

[Tri89] Trinkle, J. and Paul, R. The initial grasp liftability chart. *IEEE Journal of Robotics and Automation*, 5(1):47–52, 1989.

Distribution: (UC-32)

1. Dr. Peter Allen
CS Dept.
Columbia University
New York, NY 10027

2. Dr. Ruzena Bajcsy
CIS Dept.
University of Pennsylvania
Philadelphia, PA 19104

3. Dr. George Bekey
CS Dept.
University of Southern California
Los Angeles, CA. 90089

4. Dr. Mark Cutkosky
Mechanical Engineering Dept.
Stanford University
Stanford, CA 94305

5. Dr. Paulo Dario
Centro E Piaggio
Universita Di Pisa
Via Diotisalvi, 2-56100 PISA
Italy

6. Dr. Rod Grupen
CIS Dept.
University of Massachusetts
Amherst, MA 01003

7. Mr. Rob Howe
Mechanical Engineering Dept.
Stanford University
Stanford, CA 94305

8. Dr. Thea Iberall
CS Dept.
University of Southern California
Los Angeles, CA. 90089

9. Dr. Roberta Klatzky
Psychology Dept.
University of California
Santa Barbara, CA 93106

10. Dr. Susan Lederman
Psychology Dept.
Queen's University
Kingston, Ontario Canada K7L 3N6

11. Dr. Mike McCarthy
Mechanical Engineering Dept.
University of California
Irvine, CA 92715

12. Dr. Richard Paul
CIS Dept.
University of Pennsylvania
Philadelphia, PA 19104

13. Dr. Ken Salisbury
AI Lab
545 Technology Square
MIT
Cambridge, MA 02139

14. Dr. Tom Speeter
Room 4E-638
AT&T Bell Labs
Crawfords Corner Rd
Holmdel, NJ 07733

15. Dr. Harry Stephanou
School of Information Technology and Engineering
George Mason University
Fairfax, VA 22030

16. Dr. Jeff Trinkle
SIE Dept.
University of Arizona
Tucson, AZ 85721

17. Dr. S. T. Venkataraman
Tele/Autonomous Systems Group
JPL
4800 Oak Grove Drive
Pasadena, CA 91109

1400	E. H. Barsis
1410	P. J. Eicker
1411	S. A. Stansfield
1411	J. J. Wiczer (20)
1412	P. J. Eicker (Actg)
1414	R. W. Harrigan
1415	K. T. Stalker
1420	W. J. Camp
3141	S. A. Landenberger (5)
3141-1	C. L. Ward (8) for DOE/OSTI
3151	W. I. Klein (3)
8524	J. A. Wackerly