



January 1989

Machine Systems for Exploration and Manipulation: A Conceptual Framework and Method of Evaluation

Ruzena Bajcsy
University of Pennsylvania

Susan Lederman
University of Pennsylvania

Roberta L. Klatzky
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

Ruzena Bajcsy, Susan Lederman, and Roberta L. Klatzky, "Machine Systems for Exploration and Manipulation: A Conceptual Framework and Method of Evaluation", . January 1989.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-89-03.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/755
For more information, please contact repository@pobox.upenn.edu.

Machine Systems for Exploration and Manipulation: A Conceptual Framework and Method of Evaluation

Abstract

A conceptual approach to describing and evaluating problem-solving by robotic systems is offered. One particular problem of importance to the field of robotics, disassembly, is considered. A general description is provided of an effector system equipped with sensors that interacts with objects for purposes of disassembly and that learns as a result. The system's approach is "bottom up," in that it has no *a priori* knowledge about object categories. It does, however, have pre-existing methods and strategies for exploration and manipulation. The sensors assumed to be present are vision, proximity, tactile, position, force, and thermal. The system's capabilities are described with respect to two phases: object exploration and manipulation. Exploration takes the form of executing "exploratory procedures," algorithms for determining the substance, structure, and mechanical properties of objects. Manipulation involves "manipulatory operators," defined by the type of motion, nature of the end-effector configuration, and precise parameterization. The relation of the hypothesized system to existing implementations is described, and a means of evaluating it is also proposed.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-89-03.

**Machine Systems For Exploration
And Manipulation
A Conceptual Framework And
Method Of Evaluation**

**MS-CIS-89-03
GRASP LAB 172**

**Ruzena Bajcsy
Susan Lederman
Roberta L. Klatzky**

**Department of Computer and Information Science
School of Engineering and Applied Science
University of Pennsylvania
Philadelphia, PA 19104-6389**

January 1989

Acknowledgements:

**This research was supported in part by Air Force
AFOSR F49620-85-K-0018, US Army grants
DAA29-84-K-0061, DAA29-84-9-0027,
N00014-85-K-0807, NSF grants MCS-8219196-CER,
IRI84-10413-402, INT85-14199, DMC85-17315, NIH
NS-10939-11 as part of the Cerebro Vascular Research
Center, NIH 1-R01-NS-23636-01, NATO grant
0224/85, NASA NAG5-1045, ONR SB-35923-0,
DARPA grant N00014-75-K-0018, and by DEC Corp,
IBM Corp and LORD Corp.**

Machine Systems for Exploration and Manipulation: A Conceptual Framework and Method of Evaluation

Abstract

A conceptual approach to describing and evaluating problem-solving by robotic systems is offered. One particular problem of importance to the field of robotics, disassembly, is considered. A general description is provided of an effector system equipped with sensors that interacts with objects for purposes of disassembly and that learns as a result. The system's approach is "bottom up," in that it has no a priori knowledge about object categories. It does, however, have pre-existing methods and strategies for exploration and manipulation. The sensors assumed to be present are vision, proximity, tactile, position, force, and thermal. The system's capabilities are described with respect to two phases: object exploration and manipulation. Exploration takes the form of executing "exploratory procedures," algorithms for determining the substance, structure, and mechanical properties of objects. Manipulation involves "manipulatory operators," defined by the type of motion, nature of the end-effector configuration, and precise parameterization. The relation of the hypothesized system to existing implementations is described, and a means of evaluating it is also proposed.

Machine Systems for Exploration and Manipulation A Conceptual Framework and Method of Evaluation¹

| | | |
|----------------------------|--------------------|---------------------------|
| Ruzena Bajcsy | Susan Lederman | Roberta L. Klatzky |
| University of Pennsylvania | Queen's University | University of California |
| Philadelphia, PA | Kingston, Ontario | Santa Barbara, California |

1 Introduction

The study of sensate robotic systems is at an early stage, but already there is considerable activity at both the hardware and software levels. What has not yet appeared in this area is a principled theoretical context in which to systematically develop issues and to evaluate proposed solutions. The present paper describes such a context, drawing from the field of behavioral science as well as robotics.

In this paper, we offer a general framework for addressing the question of how a robotic system should solve a problem, by describing one potential system. We assume that the system is equipped with particular sensory and effector capabilities. Drawing in part from experimental work with humans, we suggest approaches to algorithms that might control these capabilities, evaluate results of actions, and learn. We then decompose the general problem into components, and describe how the system might handle each sub-problem and its potential for success and failure, given an analysis of its capabilities and limitations.

In addition, we suggest a means of assessing the contributions of the proposed system components. The approach is to deprive the system of specific capabilities, in either software or hardware, and to determine the effects on the problem domain. A set of preliminary predictions is offered, to be tested in future research.

Our goal is to provide one example of a systematic conceptual approach that might be used within the field of robotics. The potential advantage of this approach is that it embeds any particular problem within a broader theoretical context, one which should facilitate comparisons among various activities of robotic systems and the specific algorithms that are offered to perform them.

The framework we are proposing is sufficiently general that it can accommodate a wide variety of visual systems that deliver three-dimensional information, together with robotic end effectors that are both sensate and capable of movement with several degrees of freedom. The sensors we assume, in addition to vision, are proximity and haptic, the latter including tactile, position, force, and thermal sensing. The force sensor should provide force vectors and torques in three dimensions.

An important initial concern is the relative contribution of data derived from the sensors, on the one hand, and a knowledge base describing the world of objects, on the other. Extensions of artificial intelligence suggest that the knowledge base might be given a high priority, and yet, as sensory data increase in availability and accuracy, the contribution of data-driven processes becomes potentially more potent. We adopt the position that such processes can in fact be given precedence, driving knowledge acquisition by the system.

Our approach is "bottom up," in that we do not assume the system has a priori semantic information about object categories. It does, however, have pre-existing methods and strategies for exploration and manipulation, and it is also capable of tuning its rules and developing new ones by

learning. This knowledge constrains data acquisition, data interpretation, and subsequent actions. Thus, the system is both non- deterministic and intelligent.

An implementation of the proposed framework is currently under development in the GRASP laboratory of the University of Pennsylvania. This system includes a stereo vision component, Puma arm, and a moderately complex hand (Ulrich, Bajcsy, & Paul, 1988), which will be equipped with proximity, position, force, tactile, and thermal sensors. Previous work in this laboratory has implemented precursors to the proposed system. Although these initial efforts lack the full sensing and effector capabilities that we consider here, they are nonetheless informative about the system's potential, and they will be reviewed below.

2 Disassembly as Data-driven Manipulation

In general, the purpose of the system is to interact with objects and hence to learn about them. The domain of objects includes solids that vary in material, structure, and mechanical properties. More specifically, we describe the system's operation in a prototypical task: "disassembly," defined as decomposing objects into their component parts without mechanical damage. This task can be divided into two general phases, namely, initial haptic and visual exploration, followed by task-oriented manipulation. We choose to develop our framework in the context of the disassembly task, because it imposes substantial demands on both intelligent exploration and manipulation, and it fits the data-driven context particularly well. (Assembly, on the other hand, is intrinsically model-driven.) Disassembly further provides a paradigm for learning about objects and constructing representations.

The disassembly problem requires multiple sensing devices (except for trivial cases). Vision is used to guide the end effector to the object and to extract gross structural features; it is also used to guide manipulation and to evaluate its effects. Proximity sensing indicates imminent contact with the object; this is particularly important when the hand obscures vision. Purposive haptic exploration is a necessity with the bottom-up approach, because there are no pre-stored object features. A variety of features are haptically extracted.

The importance of the present bottom-up approach can be seen clearly when we contrast it with how the problem of exploration and manipulation has been treated in systems science. There, these topics are subsumed under "system identification" (Zadeh, 1962; Eykhoff, 1974). The specific task, for example manipulation or disassembly, is considered as a process. The control signals are obtained from sensors and their derivatives that measure position, geometric structure, size, force/torque, and so forth. The process is modeled mathematically using relations based primarily on mechanics and geometry, such as the laws of Newton and Coulomb. The validity of these relations further requires certain assumptions which are expressed in terms of parameters such as compactness and coefficients of friction and stiffness (e.g., Cutkosky, 1988). When modeling a process, the designer typically makes a priori assumptions about the values of these parameters. This is a reasonable procedure, since having initially chosen the material and the design structure, he or she knows the values of these object properties.

However, this is certainly not the case when a robot is placed in an unknown environment. Fel'dbaum (1960) has considered a similar situation, that is, one in which there is no a priori knowledge available about the process and its parameters. Here too, it becomes desirable to investigate (i.e., explore) the characteristics of the process, as well as direct (i.e., manipulate) it. Fel'dbaum uses the term "dual control" to describe the two functions, "investigation" and

”direction”.

This distinction highlights an important problem in dual-control systems, namely, that potential conflicts arise as a result of competition between the investigatory and directional functions of the process. Efficient control can only be realized when actions on the object are coordinated over time; a delay at any point can weaken the control process. Nevertheless, control also requires that the properties of the object be sufficiently known, which necessitates that time be spent on the investigatory function. If the directional function is too precipitous, it may execute an operational movement without appropriately using the results of the object investigation. In contrast, if control proceeds too cautiously, the system will delay longer than necessary. Processing of the information obtained by investigating object properties will be performed at the expense of directing the object to its required state at the correct time. It is necessary to properly coordinate these two functions, so as to maximize a pre-specified quality control criterion.

This analysis indicates that two risk factors must be considered at any given time. The first factor is investigatory risk, which results in unnecessary delays in task performance. The second is directional risk, that is, the risk of action due to insufficient knowledge of the properties of the manipulated object and/or system.

For the remainder of this paper, we will use the term ”exploration” to refer to investigation, and ”manipulation” to refer to direction. Schematically, we can represent the two processes as shown in Figure 1.

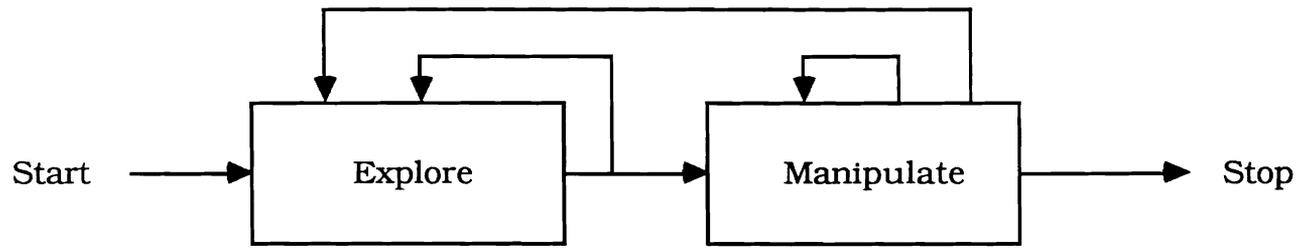


Figure 1: Exploration and Manipulation

Note that each process can repeat itself several times. Further, the coordination between the two processes is controlled by a tradeoff between exploratory (investigatory) risk versus manipulatory (directional) risk.

Our paper next outlines the proposed system's a priori methods of exploration and manipulation and its control strategies. Primitive "exploratory procedures" and "manipulatory operators" are described, as are the primitive attributes extracted during exploration. We then present a stage analysis of the disassembly task and consider previous work that appears relevant to this analysis. Next we evaluate how well existing systems have instantiated the present assumptions. Finally, we describe research avenues for validating the proposed framework.

3 Object Exploration

In its initial exploratory mode, the system's goal is to learn about the structural, material, and mechanical properties of objects. Vision may be used to guide the effector to an object and to extract a structural description. The object's representation is then considerably augmented by haptic sensing. The following discussion concentrates on haptic exploration.

3.1 Exploratory Procedures

Human psychological research (Lederman & Klatzky, 1987) has established that haptic sensing is accomplished through a set of stereotypical patterns of hand movement, each pattern being optimal for delivering a particular object property. These patterns are called "exploratory procedures." In our theory, these are programmed a priori as algorithms for property extraction (see also Klatzky, Bajcsy, & Lederman, 1987, for a preliminary discussion of robotic exploratory procedures). Lederman and Klatzky have documented two types of exploratory procedures in human haptics, which pertain to an object's structure and substance properties. To these we now add a third, procedures for extracting mechanical properties. A summary of the properties extracted in each procedural domain is presented in Table 1. The individual procedures are discussed in detail below.

Table 1: Properties extracted from objects

Structure-related

- shape
- size
- weight (with respect to size)

Substance-related

- hardness
- surface texture
- thermal properties (e.g., thermal conductivity)
- weight (with respect to density)

Mechanical

- elasticity (vs. rigidity)
- brittleness
- viscosity
- coefficient of friction
- part motion (linear vs. rotary)
- degrees of freedom

3.1.1 Structure-related Procedures

The structural properties that have been studied are size, shape, and weight (to the extent it is determined by size). The two haptic exploratory procedures used for this purpose are Enclosure (molding of the end effectors to the object contours) and Contour Following (dynamic edge following).

3.1.2 Substance-related Procedures

These determine information about the following object properties: hardness, surface texture, thermal, and again, weight, to the extent it is determined by density. The thermal property is a complex primitive that may be decomposed into components such as conductivity, specific capacity, and diffusivity.

The exploratory procedures associated with these substance primitives are: hardness – Pressure (application of normal forces or torque to the object); texture – Lateral Motion (lateral back-and-forth rubbing movements, usually on a homogeneous portion of the surface); thermal – Static Contact (static placement of the end effector on the object without contour molding); weight – Unsupported Holding (lifting the object away from a support).

3.1.3 Mechanical Procedures

Mechanical properties are defined as those that reflect the forces applied to an object and the resulting motions. While Klatzky and Lederman did not exhaustively study mechanical properties of objects, they did examine procedures for determining the type of motion of object parts. We do not have substantial human data on the exploratory procedures that elicit such properties at present. However, a set of properties can be defined as shown in Table 1.

Of these mechanical properties, coefficient of friction is of particular importance to the robotics community, which has only now begun to consider the effects of friction – specifically on the stability of control. However, this topic has been studied by other areas in the engineering community, particularly with respect to the friction behavior of a brush-type dc servo-motor driven mechanism (e.g., Armstrong, 1988). Further, the study of non-linear friction in servo mechanisms has a long history. For example, Tustin (1947) examined the effect of backlash and non-linear friction upon feedback control. And Tou and Schulteiss (Tou, 1953; Tou & Schulteiss, 1953) thoroughly studied the consequences of static and kinetic friction on control. Finally, slipping, sliding, and rolling are becoming of considerable concern to roboticists (e.g., Howe, Kao, & Cutkosky, 1988; Brock, 1988; Cole, Hauser, & Sastry, 1988).

We emphasize here, however, that all of the work cited deals either with the impact of friction on control or with the kinematic and dynamic models of finger/object relations; in all cases knowledge of the friction coefficients is assumed. Why should this be so? We believe that there are two reasons. First, robotic sensors that might extract the requisite mechanical information have not been readily available. Second, there is currently a strong emphasis on top-down processing, as exemplified by Iberall, Jackson, Labbe, & Zampano (1988), who allow substance and mechanical properties to be inferred from a priori knowledge rather than learned by active exploration. Nevertheless, a little work on these properties has been done. Stansfield (1987) has implemented a procedure for determining the elasticity of an object, for example.

4 Object Manipulation

We now consider the manipulatory activities involved in our prototypical task, object disassembly.

4.1 Manipulatory Operators

4.1.1 Simple operators

The simplest form of manipulation is to apply a force on an object in some direction with some effector configuration for a period of time. Such simple actions are the result of three types of decisions. One pertains to the nature of the manipulatory action, another to the geometry of the effector configuration, and the final one to selection of a set of parameters related to force, time, and details of the effector configuration. The result of these decisions is a simple manipulatory operator, the components of which are shown in Table 2.

Table 2. Manipulatory Operators

Operator

- Action Type
- translation
- rotary

Effector Configuration

- number of components contacting object
- flexion vs. extension (prehensility)

Detailed Parameterization

- force paramters (magnitude, direction)
- effector parameters (joint angles, range of motion)
- temporal parameters (duration, time between retries)

Common terms for manipulation include "lift", "push", "pull", "turn", and "stabilize". Within our scheme, these are the result of an action type and parameter selection. Lifting, pushing, and pulling are all translations that vary in terms of their parameterization. Lifting is translation counter to gravitational force. Interestingly, pushing and pulling are both translations in a common plane. The distinction in terminology is made with respect to the body, and often pulling is reserved for a prehensile effector configuration. Turning is used to refer to rotary movement. Finally, stabilization is the absence of translation or rotation due to one force being applied counter to some other force.

In the robotics field, Mason (1984) has studied in detail the mechanics involved in pushing and pulling. Lifting is well developed, as noted above. And Salsbury (1981) has described the minimal configuration required for fully stabilizing an object. We know of no work on turning.

Common terms are also used to refer to the effector configuration, including "poke", "pinch", "palm", and "clench". These terms can be differentiated by the variables of number of effector components used to contact the object, and whether they are flexed or extended. Thus, a poke occurs with one extended finger, a pinch with two or three flexed fingers; palm and clench involve flexed and extended fingers of the full hand, respectively. These distinctions have been made for the human hand by Klatzky, McCloskey, Doherty, Pellegrino, & Smith, 1987.

The choice of the manipulatory operator is determined jointly by constraints imposed by goals of the system, the physical properties of the end effector, and the environment. Hager (1988) has explicitly modeled the cost of information for purposes of both exploration and manipulation. This work therefore shares a common theme with the earlier analysis by Fel d'baum. The constraints imposed by the geometry of the object and the physical capacities of the hand have been studied by psychologists. Napier (1956; see also Malek, 1981) suggested that the choice of end-effector geometry is contingent on the amount of force to be applied. How hand shape is dependent on object shape has also been studied (e.g., Klatzky et al., 1987; Kroemer, 1986). Newell (in press) has considered the joint constraints of hand size and object size on the choice of end effector configuration for grasping objects.

4.1.2 Compound operators

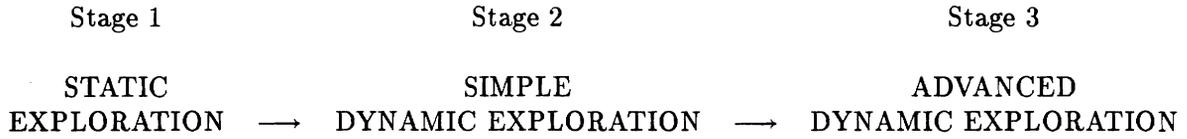
These simultaneously or in sequence apply two or more simple manipulatory operators, such as when bending, squeezing, or twisting. Compound operators are frequently necessary to move one part of an object relative to the rest of it, or to make complex motions such as scissoring (mirror-image pushing and pulling).

5 Stages of Disassembly

5.1 Exploratory Phase

We now consider a stage analysis of the disassembly task, beginning with the exploratory phase. Three stages of exploration, distinguished by the complexity of movement involved, and two stages of manipulation are described. In the present analysis of exploration, we include only static vision, to limit the complexity of the problem. Table 3 provides an overview of the three exploratory stages with the corresponding information about object properties that may be extracted haptically.

Table 3. The exploratory phase of disassembly



| Exploratory Activity | Exploratory Activity | Exploratory Activity |
|--|---|--|
| single contact | multiple contacts (produced by a single moving contact, or by multiple static contacts) | object or part manipulation |
| Outputs | Outputs | Outputs |
| Solid Static coefficient of friction Local crude texture Thermal:conductivity, capacity, etc. Local structure: corners, edges, radius of curvature | Hardness, stiffness coefficient, brittleness, elasticity, visco-elasticity Dynamic coefficient of friction Global and more precise texture; homogeneity of material Global structure: envelope shape; size | Global Structure: precise shape; Weight direction, degrees of freedom |

Stage 1: Static Exploration.

In the first stage of exploration, the system statically attempts to assess the "physics" of the world in which the object exists. When vision is present, it provides structural information about the visible surfaces of the object and its context. Static haptic exploration consists of simple contact. From this, the array sensor can provide local information about texture (using the grain in the image) and local structure. A thermal sensor allows the system to determine thermal properties of the object, for example, conductivity, capacity, and so forth.

As the information that is provided at this stage is limited by the performance of the sensors, one must develop as complete a model as possible of the sensors before implementing Stage 1 (Fuma & Bajcsy, submitted; Krotkov, 1987). The parameters to be modeled are, for example, the sensitivity function, hysteresis, spatial resolution, signal-to-noise ratio, reliability, variability, and the range of admissible values.

Stage 2: Simple Dynamic Exploration.

Next, the relatively simple dynamic exploratory procedures are added. A plausible sequence is as follows.

1. The force and position sensors are used to determine hardness and elasticity. As described by Stansfield (1987), this is done by comparing the initial and final position after the application of force using a Pressure procedure (to extract hardness), and re-contacting the object to determine if its initial position is resumed after the force is withdrawn (to extract elasticity).
2. Lateral Motion is added to determine texture. This procedure enhances the availability and possibly the precision of texture information compared to the static array measure. In addition, it is likely to provide information about the dynamic coefficient of friction.
3. Structural properties are determined by vision and by haptic Enclosure. Vision is used to guide the robotic fingers to enclose the object as symmetrically as possible. The force sensors then are used in an attempt to equalize the forces applied by each finger pad, thus equating the positions of the fingers relative to the object surface. If this can be achieved, the array sensors can be examined unambiguously, and comparison of array outputs can be used to determine if the material is uniform. More complex algorithms must be used if finger positions and forces cannot be equated. Size and gross shape can be determined from the same haptic Enclosure procedure.

Stage 3: Advanced Dynamic Exploration.

At this stage, exploration includes dynamic interactions with the object, to extract details of substance, structure, and mechanical properties. These interactions constitute manipulation, but in the service of exploration.

Dynamic interaction of this sort brings up an important consideration: Grasping must remain stable, while manipulatory forces are applied. Yoshikawa and Nagai (1988) have presented an analysis of fingertip forces that provides a decomposition into two components: those forces that maintain the stability of grasp, and those that effect the desired manipulation. They define the manipulating force as a fingertip force satisfying the following three conditions. First, it will produce

the specified resultant force. Second, it is not in the inverse direction of the grasping force. Third, it does not contain any grasping force component.

The following exploratory procedures are suggested for Stage 3:

1. The object is lifted (Unsupported Holding) to determine weight.
2. Contour Following is used to determine the object's shape in detail and its potential part structure.
3. Specific mechanical procedures (to be defined) determine high-level properties such as the nature of part motion and degrees of freedom at a joint.

5.2 Manipulatory Phase

Thus far, we have defined and described three exploratory stages, which together constitute the object exploration phase of the disassembly task. With the foregoing discussion of manipulatory primitives, we can now consider the second phase of our task, consisting of two manipulatory stages. During these stages, manipulatory operators are applied, driven by the data gathered during the earlier exploratory phase. An overall flow diagram of the manipulatory phase is presented in Figure 2.

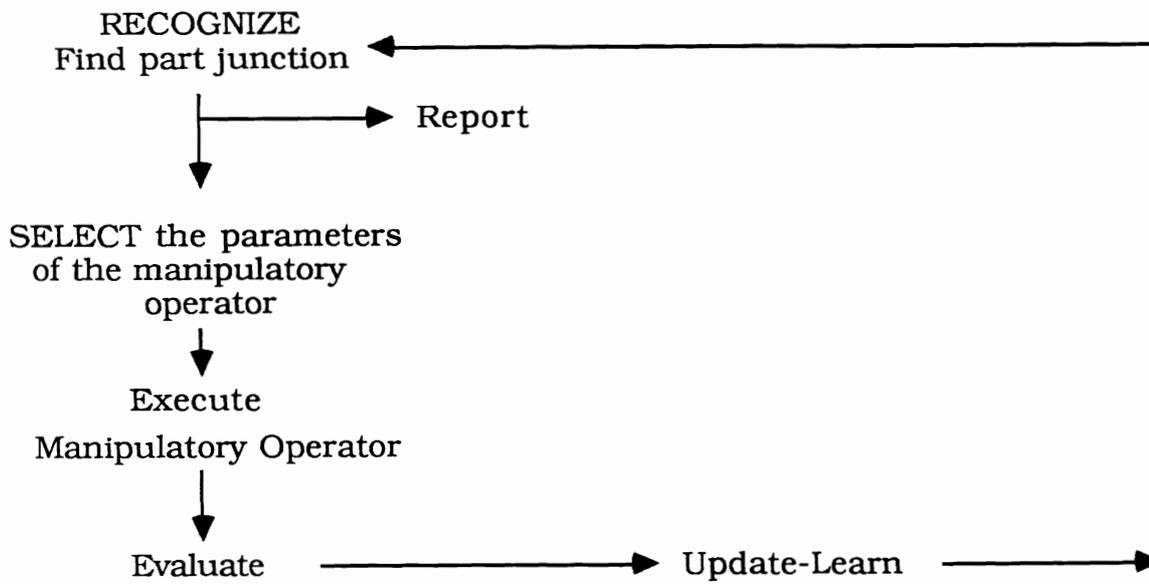


Figure 2: Flow diagram of the manipulatory phase

Stage 4: Task-Oriented Manipulation.

The fourth stage in our task involves the performance of some manipulatory action that is required for disassembly. We conceive of three substages:

1. Find next part juncture to test for disassembly (using vision or haptics). A part is defined structurally as the contour between two concavities and/or by a homogenous textured surface or volume.
2. Select and parameterize a manipulatory operator to apply to the part.
3. Apply manipulatory operator and determine result. Possible results are: the part is not movable, the part is movable but not removable, and the part is both movable and removable. Vision can be used to make this evaluation, but the force and position profile over time should also indicate whether movement has occurred.

Stage 4 requires a set of "executive" rules that are external to the exploratory/manipulatory apparatus per se. These rules reside at a higher level, controlling exploratory and manipulatory activities.

In Substage 1, rules must exist to prioritize junctures on the object for testing. Visual dominance is suggested here; junctures recommended by visual analysis will be tested before those recommended by haptic exploration. We justify this rule on two bases: Vision provides a global perspective, and it also excels at finding structural boundaries. For complex objects having parts that must be disassembled in a specific order, sophisticated rules are needed, such as, start with junctures of external parts that project beyond the outer envelope of the object. Such rules are likely to be available only after learning. Prioritization of parts establishes a starting rule.

Substage 1 also requires a stopping rule, which determines when no further part should be selected. Either a global a priori rule is needed, such as stopping after a fixed time, or the exploratory data must determine the stopping parameter. Basing the stopping point on an analysis of the number of part junctures seems untenable for complex objects, which might have a large number of fine projections to test. Therefore, some rule external to the bottom-up process itself seems necessary.

Guidance in the selection of manipulatory operators in Substage 2 requires a number of heuristic rules. The highest priority should be given to the general principle of economy, that is, do no more than is necessary in consumption of time and energy. Essentially, we assume the executive minimizes some function of time and energy. We offer some potential rules, although caution must be exercised regarding their generality:

- (a) Minimal force: Use as little force as possible, to prevent objects from sliding out of the workspace or breaking. This disfavors lifting, because considerable force would usually be required to overcome gravity.
- (b) Motoric ease: Select operators where the complexity of the end-effector geometry is as simple as possible.
- (c) Minimal system change: Use a configuration that already exists, if possible, or change it as little as possible. Execute an operator that will change the configuration as little as possible (this rule would downgrade scissoring motions, for example).

- (d) Time sharing: Do two things at once, if compatible. Compatibility is a complex concept that needs to be analyzed in further detail. One aspect of compatibility is motoric. For example, lifting and squeezing are two separate acts that may be easily performed in combination.
- (e) Past history of success: Use operators that have previously been successful.
- (f) Don't persist. When the range of usual forces has been exhausted, don't continue. Apply any operator for only a "reasonable" time.

We avoid specific rules that connect structural properties with appropriate manipulations. An example is, parts that project from slots within a plane suggest a translation action (push, pull). Such rules are likely to result from learning (see below).

Conflict resolution must be performed, because rules often disagree. For example, execution of a compound operator violates motoric simplicity but adheres to the time-sharing rule.

We note that the foregoing rules apply to the manipulatory phase of disassembly, but similar rules could be applied to the exploratory phase. As described above, there is a prescribed, exhaustive sequence of exploration that is followed. If the exploratory phase were made strategic, as we envision the manipulatory phase, then exploration would proceed as economically as possible and would terminate early, if sufficient information were obtained to allow a disassembly test. A more sophisticated model would alternate between exploratory and manipulatory phases.

Stage 5 – Learning.

Following the application of a manipulatory operator, learning can occur. We are interested here in learning of a relatively complex sort. The question is, given a sequence of sensory measurements and exploratory and manipulatory procedures, what data reduction and strategic processing will allow an effective set of actions to be repeated or appropriately modified? Modification that generalizes or adapts a learning sequence is particularly important. For example, it might enable an object to be assembled, given its history of disassembly.

Essentially, we view what is stored as a temporally ordered array giving the results of the exploratory procedures, the manipulatory operators attempted with their parameter values, and the results of manipulation. In this way properties of the object are stored in relation to manipulatory activities. Success or failure of the manipulatory operators will adjust the array. An open question is whether a PDP approach lends itself to such a model. A PDP mechanism is well suited to building associations between manipulatory activities and particular objects, but its demands on system memory may make it unfeasible for use with conventional robotics.

The learning mechanism that is sought must be capable of inducing rules at very different levels of complexity. Learning can be general, such that operations that work are increased in strength, while those that do not work are decreased. Learning can also be specific to a particular manipulatory episode, for example, learning that a particular object structure was successfully disassembled with rotary force at a particular part. Note we assume that the system does not learn the entire routine for dealing with a specific object; for example, there is no reason to remember exactly where a stabilizing force was applied. Learning some key rule is all that is needed, as general strategies can then be used to make sure that the rule is executed properly.

Finally, it should be noted that although we have described exploration and manipulation as consisting of two phases, in practice they must occur in alternation or even in tandem, especially

given our bottom-up approach. Manipulation may be interrupted for further exploration; there may be a repeated manipulatory/exploratory cycle.

6 Existing Systems as Theory Instantiations

No existing system fully instantiates the present approach. However, various aspects have been implemented. Below, we describe current implementations including the one being developed in the University of Pennsylvania GRASP laboratory. We also summarize the contributions of its predecessors.

6.1 Hardware

Many different physical principles are currently being explored for robotic sensors. Some substance and structural information can potentially be provided by robotic tactile array sensors (e.g., Winger & Lee, 1988; Tise, 1988; Fearing & Binford, 1988; Clark, 1988; Cameron, Daniel, & Durrant-Whyte, 1988; Begej, 1988; Dario & DeRossi, 1985; Dario & Buttazo, 1987). Unfortunately, very few are commercially available, and these provide very limited spatial detail in all three dimensions. Thus, texture resolution is very crude, while the structural information provided by the array sensor enables detection of contact with an edge or corner, and possibly radius of curvature as well. Size can be discriminated up to the limited extent of the array. Work on sensors that can provide thermal information through static touch has also been initiated (e.g., Russel, (1986); Siegel et al. (1987)).

At the University of Pennsylvania, a complex system involving visual and tactile sensing has been developed. The visual system used has been described in detail by Allen (1985) and Krotkov (1987). The robotic hand is mounted on a PUMA 6-degree of freedom manipulator. Previous work has been done with a Lord sensor mounted on the end of the PUMA arm, and with a gripper equipped with two opposing Lord sensors on its interior surfaces. The Lord sensors provide a 10 X 16 pixel array in a surface approximately 1 in square, as well as force vectors and torques in the planar surface of the array and the axis of the arm.

Ulrich, Bajcsy, and Paul (1988) are currently developing a three-fingered hand (the UPENN hand) with four full degrees of freedom and three coupled joints. Each finger has two joints, proximal and distal to the PUMA arm. The proximal joint flexes and extends, and the distal joint is coupled to this movement. Two of the fingers also move laterally, toward and away from one another. This allows for a variety of effector configurations, including:

- single finger extends or flexes, two curled toward palm
- two fingers flush together, third curled toward palm
- two fingers flush or separated, third opposes
- three fingers oppose symmetrically (i.e., 120 deg apart)
- one or two fingers yoked as hook

The hand will include tactile, force, and thermal sensors. Each finger has a force sensor providing vector and torque information at the distal joint, and a tactile sensor on the tip. The palm will

have both a Lord sensor (described above) and a thermal sensor. One finger has a removable distal phalanx to which appendages can be mounted, such as a "nail" for generating vibration during lateral motions.

7 Existing Algorithms for Exploration and Manipulation

Algorithms for vision are well developed and have been described in books such as Ballard and Brown (1982) and Horn (1986).

Unfortunately, relatively little work has been done to implement haptic exploratory procedures. An excellent argument concerning the need for robotic exploratory procedures was provided recently by Gottschlich and Kak (1988), in a discussion of strategies and errors in high-precision part mating. Part-mating strategies employ both guarded and compliant motion. Guarded motion brings a part into contact with its target environment, while compliant motion slides the part along or through the mating part.

To date, most work in fine-motion planning has incorrectly assumed that motion can occur in a quasi-static manner. However, different considerations arise in a dynamic context. In general, there are different modes of impact that can occur between mating surfaces, leading to unexpected changes in the position of the end-effector and to unanticipated forces. Thus, it is unlikely that either a guarded motion or a compliant motion can satisfy a force constraint precisely. The ensuing sticking can cause unexpected changes in the position of the part, further exacerbating error, especially in a high-force condition.

For compliant moves that are directed toward a defined goal position while maintaining a specific force constraint in the perpendicular direction, force errors are related to position errors through a stiffness "constant". However, it is very questionable whether this coefficient of restitution should be represented as a constant. Typically, the plasticity of a region of deformation depends at any moment on how much it has already been deformed. Thus as deformation continues, the coefficient of restitution should change, making the problem mathematically intractable.

The foregoing arguments suggest that part-mating is complex and non-deterministic. Haptic exploration of the object's substance and mechanical properties would prove very useful in this situation, providing ongoing information about changes in object properties and part relationships.

Some haptic algorithms for Stages 1 and 2 have been developed by Stansfield (1987), who used the Pressure procedure to extract gross texture, hardness, and elasticity. Compliance, which is related to the hardness/elasticity property, has been widely used for control purposes (e.g., Xu & Paul, 1988; Masson, 1981; Whitney, 1977; Paul & Shimano, 1976). Lateral Motion will be implemented in the hand designed by Ulrich, Bajcsy, & Paul (1988) by adding a "nail" to the robotic effector to determine vibratory forces as the nail moves across an object surface. The Stage-3 procedure for Unsupported Holding (lifting) has also been developed (Stansfield, 1987; Tsikos, 1987), although no weight determination has been made in conjunction with Unsupported Holding. Stansfield (1987) and Allen (1987) have implemented one-finger versions of Contour Following, and a version for two fingers with limited degrees of freedom has been developed by Koutsou (1988). However, developing a three-finger algorithm for obtaining contour information remains an open problem.

With respect to Stage 4, rudimentary versions of many of its problems have been handled by Tsikos (1987). Tsikos' task was to remove objects from a table top; thus all "parts" of the display were both movable and removable. The shape and size of the object determined whether it was

pushed (using a gripper in conjunction with a spatula), lifted by prehension (using the bare gripper) or lifted by suction (using a cup held by the gripper). The tools (spatula and suction cup) can be viewed simply as specialized end effectors. The strategies were modeled by a nondeterministic (data-driven) automaton. The automaton was finite state. Its stopping rule was clear – stop when the table is empty.

In work relevant to the executive rules of Stage 4, in particular, Peshkin&Sanderson (1988) have considered the minimization of energy in quasi-static manipulation. They suggest an energy principle for quasi-static systems that is similar to our principle of minimal force. It states that the system should select, from the set of all motions that satisfy the existing constraints, the motion that minimizes the instantaneous power (i.e., the lowest energy or "easiest" motion). Although intuitively appealing, the authors point out that the principle is generally false! For example, in the case of viscous forces, the motion predicted by the minimum power principle will be incorrect. But they further show that the principle is correct in the useful but special case where Coulomb friction is the only dissipative or velocity-dependent force acting in the system.

There have been attempts to deal with learning processes (Stage 5) in the robotics literature, specifically as applied to automation of robot programming (e.g., Bertenstein & Inoue, 1988).

8 Evaluation of the Framework

The biological perceptual sciences have evaluated the relative contribution of various sensory/perceptual systems by using several different experimental paradigms. One method is to use a population of subjects (e.g., the blind) in whom the specific sensory cues are either missing or deficient, due to disease or accident. A second method is to create this same situation in a more controlled manner by temporarily and selectively fatiguing the sensory channel under investigation. A third approach is to create a discrepancy between two sensory systems (e.g., vision vs. touch) or between sensory cues from the same system (e.g. motion parallax vs. binocular disparity cues to visual depth perception), to evaluate the relative contribution of these different sensory inputs to perception.

We view all three paradigms as being of potential value in the assessment of robotic perceptual systems. We have adopted a variant on the first approach above to evaluate the present framework: That is, we consider the impact of eliminating individual components – sensors, exploratory procedures, manipulatory operators, and higher-level algorithms – on system performance. Specific predictions are discussed below.

8.1 Role of the Sensors

8.1.1 Vision

Elimination of vision in our system should have a profound effect, particularly on manipulation. Blind reaching would become necessary to find the object; however, the proximity sensor should minimize risk of damage. The haptic exploratory procedures can still provide structure, substance, and mechanical information, although their burden would be increased by the absence of a gross visual analysis. Structural information should be particularly affected. Manipulation should still be performed effectively. However, it should be more difficult to determine where to apply manipulatory forces, and to evaluate success in disassembly.

8.1.2 Proximity

Proximity and vision provide redundant information about the distance of the effector relative to the object. If both of these non-contact sensors were removed, there would be increased risk of damage to the object through blind reaching. However, the proximity sensor could be eliminated without substantial cost, if vision remained.

8.1.3 Tactile

Loss of the tactile sensor would preclude direct structural information about contact with edges and corners. This could still be derived from the force sensor, but by a more complex algorithm, one which reads torque changes as a function of minute shifts in position. The fingers described by Salisbury (1984) and Bicchi & Dario (1988) both determine edges without a tactile array sensor. With the current quality of array sensing, the loss of texture information from this source cannot be considered devastating. However, gross texture, which might be used, for example, to detect substance boundaries on a flat surface, would no longer be provided. Lateral Motion could still be used to provide texture information through vibration of the force sensors in the joints (e.g., as will be implemented by the nail in the GRASP lab).

8.1.4 Thermal

Elimination of the thermal sensor should primarily hinder the identification of the material out of which the object is composed, through loss of information concerning the associated thermal properties. It is possible that texture and hardness cues provided by the tactile sensors, or colour and reflectance cues provided by vision might provide an alternate, though less precise, source of information.

8.2 Role of the Exploratory Procedures

In our system, an exploratory procedure directs the motor control as to how to move, then reads the sensor information and derives an object property. Although motor movement and sensors might be intact, they would not be functional without this algorithm. Elimination of an individual exploratory procedure means that a particular object property is lost. Hence the importance of the procedures can be ranked by the importance of the object properties with which they are associated. For an example of how the importance of a specific exploratory procedure could be evaluated, consider the effects of eliminating the Contour Following procedure. That is, we do not allow the force and position sensors to systematically follow edges of the object and extract a structural map. Contour Following is used to extract precise information about the contours of the object. With the reduced system, the Enclosure procedure – which we limit to a symmetrical grasp in one position, so that Contour Following is not simulated – can extract grosser structural information. Vision can extract three-dimensional information about the projecting surface. Without motion of the camera, the back surface is hidden.

This situation renders probabilistic any problem in which there is a restricted point or region of disassembly. If the disassembly point is visible, or if it happens to be discovered from the Enclosure, efforts can proceed. Otherwise, the system is stalled. However, a "smart" system might still succeed: It could use the one Enclosure with wrist rotation to turn the back of the object frontwards, hence bringing it under visual control. With a learning history, the system might in

fact operate in this fashion. Clearly, there should be a loss in system efficiency, as measured by time to disassemble.

8.3 Role of the Manipulatory Operators

Without algorithms to guide movement in specific ways, the system would become spastic – uncoordinated and non-purposive in movement. If individual operators were eliminated, the system would become deficient in particular patterns of movement. A single operator might be removed without system failure, but at a cost in accuracy and/or time. For example, lacking a turn operator, but retaining the ability to push, the system could turn an object or part by "tacking."

8.4 Role of Executive Control

Our system has not only sensory and motor capabilities, but a strategic plan for exploring and manipulating objects for purposes of disassembly. The plan is intrinsically flexible, because it is largely data-driven. However, knowledge that prioritizes activities and evaluates outcomes and costs is necessary. Without evaluation rules, infinite persistence on ineffective action is possible.

The plan must incorporate contingencies for acknowledging failure, as well as producing success. As is commonly acknowledged, it is not always possible to find plans for a given task that guarantee success. If errors exceed allowed tolerances, a failure should be signaled.

Donald (1987; 1988) has presented planning strategies that allow for both success and failure outcomes. These strategies are aimed at error detection and recovery (EDR). As others have done, he considers the problem of object assembly, in contrast to disassembly. In Donald's EDR plans, the success and failure of an assembly are very clear; there is no possibility that the plan will fail without the executor realizing it. The EDR framework therefore fills a gap when a "sure-fire" plan cannot be found. It provides a technology for constructing plans that might work, but that fail in a "reasonable" way when failure is inevitable. Donald also outlines considerably more difficult techniques for generating multi-stage EDR strategies in the presence of uncertainty.

In a constrained environment, knowledge about specific objects replaces knowledge about general strategies, and the generation of plans can follow more traditional approaches. However, our goal is to create a system that is, at least initially, bottom-up and extremely flexible. Not only efficiency, but success and failure – and their recognition – are contingent on effective planning and task control. We view disassembly as problem solving; like any such task, it requires specification of start states, end states, and operators for moving through the problem space including the error recovery operators. The more knowledge that is provided to the executive control, the more efficient is the search through the problem space.

9 Another Approach to System Evaluation

The preceding discussion indicates one avenue for a research program based on the present framework – system components are eliminated, and performance is evaluated. A further approach is to re-examine existing problems in disassembly, from the bottom-up perspective advocated here. In effect, this eliminates a major component of other systems: a base of object-specific knowledge.

The bin-picking paradigm is an important example. Here, the objects in the bin constitute decomposable parts. As typically constituted (see, for example, Kelley et al., 1982; Bolles &

Horaud, 1986), this is a problem in pattern recognition. The objects are identical, and a template is available for matching. The critical effort is to determine which object is on top and how it is oriented, so that it can be grasped.

Lacking a pre-existing template, the bottom-up system adopts a different approach to the bin-picking problem. Vision and haptic exploration are used to find junctures between objects, which are treated as potential "parts." Further exploration suggests where forces might be applied to the array. Application of manipulatory operators establishes that the target object can be moved relative to others. Ultimately, the object is grasped and removed.

An interesting difference between the two approaches arises when the objects are aligned in a stack. This makes matters more difficult for the present approach, because the cues to object boundaries are the edges formed by adjacent objects. If the joined surfaces are flat, detection of these edges may be quite difficult. In contrast, the presence of a template allows the top-down processor to anticipate the probable depth of each object, and thus to adjust the grasp.

Now, however, let us suppose that the objects are not uniform, but vary in shape, size, and material. This makes matters more difficult for the top-down approach, because multiple templates are needed. In contrast, the bottom-up approach should find the nonuniformity and uniformity problems to be of roughly the same magnitude. The work of Tsikos (1987), which has been previously described, establishes the success of the bottom-up approach in the non-uniform object case.

10 Concluding Comments

Although we call our approach "bottom up," it is not "unintelligent." The system begins with a repertoire of exploratory and motoric operations, as well as a strategy for applying them. Learning refines and augments these basic capabilities. Rather than "bottom up," the system might be termed "data-driven with constraints." It responds to what it finds in the physical environment; its responses are constrained by its rules and available procedures. The extensive exploration allows our system to adapt to the enormous variability of the world of objects – their varied shapes, sizes, substances, part junctures, and arrangements. While a top-down approach is more efficient in a restricted environment, ultimately the combinatorial explosion in a more natural context makes this approach untenable.

The present approach represents, to our knowledge, the first effort to provide a unified framework in which to view a broad spectrum of problems in robotic exploration and manipulation. This type of framework serves several functions. It specifies the task domain, and it relates tasks to the sensory and motor capabilities of the system. In doing so, this framework also provides a means of experimentally evaluating robotic systems, by considering a specific system in terms of the adequacy and efficiency with which it achieves specified goals.

11 References

1. Allen, P.K. (1987). *Robotic Object Recognition Using Vision and Touch*, Kluwer Academic Publishers.
2. Armstrong, B., (1988). "Friction: Experimental Determination, Modeling and Compensation." *Proceedings 1988 IEEE International Conference on Robotics and Automation*,

Philadelphia, PA, pp.1422-1428.

3. Bajcsy, R. (1984). "What can we learn from one finger experiments?" *Proceedings of the First International Symposium on Robotics Research*, (ed. Brady and Paul), MIT press, 1984.
4. Ballard, D.H. and Brown, C.M. (1984). *Computer Vision*, Prentice Hall, New Jersey, 1982.
5. Bartenstein, O. and Inoue, H., (1988). "Learning-Assisted Robot Programming." In R. Bolles and B. Roth (Ed.) *The Fourth Int. Symposium on Robotics Research*. The MIT Press, Cambridge, Massachusetts, pp.411-419.
6. Begej, S. (1988). "Fingertip-shaped Optical Tactile Sensor for Robotic Applications". *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp. 1752-1757.
7. Bicchi, A. and Dario, P. (1988). "Intrinsic Tactile Sensing for Artificial Hands," *Proceedings of the 4th Int. Symposium on Rob. Research* (ed. R. Bolles and B. Roth), The MIT Press, Cambridge, Mass., pp.83-90.
8. Bolles, R.C. and Horaud, P. (1986). "3DPO: A three dimensional part orientation system." *International Journal of Robotics Research*, 5(3):3-26.
9. Brock, D.L., (1988) "Enhancing the Dexterity of a Robot Hand using Controlled Slip," *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.249-251.
10. Cameron, A., Daniel, R. and Durrant-Whyte, H., (1988). "Touch and Motion." *Proceedings 1988 IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.1062-1067.
11. Clark, J.J. (1988) "A Magnetic Field Based Compliance Matching Sensor for High Resolution High Compliance Tactile Sensing." *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.772-777.
12. Cole, A., Hauser, J. and Sastry, S., (1988). "Kinematics and Control of Multifingered Hands with Rolling Contact." *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.228-233.
13. Cutkosky, M., Akella, P., Howe, R., and Kao, I. (1988). "Grasping as a Contact Sport." *Proceedings of the 4th Int. Symposium on Rob. Research*, (ed. R. Bolles and B. Roth), The MIT Press, Cambridge, Mass., pp.199-206.
14. Dario, P. and De Rossi, D. (1985). "Tactile sensors and the Gripping Challenge." *IEEE Spectrum*, Vol.22, No.8, pp.45-52.
15. Dario, P. and Buttazo, G. (1987). *International Journal Robotics Research*, Vol.6, No.3.
16. Donald, B.R. (1988). "Planning Multi-Step Error Detection and Recovery Strategies." *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.892-897.

17. Eykhoff, P. (1974). *System Identification*, John Wiley and Sons, New York.
18. Fel'dbaum, A.A. (1960). "Dual Control Theory", *Automation Remote Control*, Vol. 21, pp.874-880, 1033-1039.
19. Fearing, R.S. and Binford, T.O. (1988). "Using a Cylindrical Tactile sensor for Determining Curvature," *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.765-771.
20. Gottschlich, S.N. and Kak, A.C. (1988). "A Dynamic Approach to High-Precision Parts Mating", *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp. 1246-1253.
21. Horn, B.K.P (1986). *Robot Vision*, The MIT Press, Cambridge, Mass.
22. Howe, R. D., Kao,I. and Cutkosky,M.R. (1988). "The Sliding of Robot Fingers Under Combined Torsion and Shear Sliding." *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.103-108.
23. Iberall, T., Jacsikon, J., Labbe, L., Zampano, R. (1988). "Knowledge-Based Prehension: Capturing Human Dexterity." *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.82-88.
24. Katz, D. The World of touch.
25. Kazerooni, H. (1986). "Fundamentals of Robust Compliant Motion for Manipulators," *IEEE Journal of Robotics and Automation*.
26. Kelley, R.B., Birk,J.R., Martins, H.A.S., and Tella, R. (1982). "A Robot System which Acquires Cylindrical Workpieces from Bins", *IEEE Trans. on Systems, Man and Cybernetics*, Vol. 12, No.2, pp. 204-213.
Klatzky, R., Lederman, S., and Bajcsy, R. (1987). "Object Exploration in One and Two Fingered Robots." *Proceedings IEEE International Conference on Robotics and Automation*, Raleigh, North Carolina.
27. Klatzky, R.L., McCloskey, B., Doherty, S., Pellegrino, J. & Smith, T. (1987) "Knowledge about hand movements and knowledge about objects." *Journal of Motor Behavior*, 19, 187-213.
28. Kroemer, K. H. E. (1986). "Coupling the hand with the handle: An improved notation of touch, grip, and grasp". *Human Factors*, 28, 337-340.
29. Krotkov, E., (1987). "Exploratory Visual Sensing for determining Spatial Layout with an Agile Camera System," University of Pennsylvania Ph.D. Dissertation, available also as a Technical Report MS-CIS-87-29.
30. Koutsou, K. (1988). "Object Exploration Using a Parallel Jaw Gripper", submitted to *IEEE Journal of Robotics and Automation*.

31. Lederman, S.J. and Klatzky, R.L. (1987). "Hand movements: a window into haptic object recognition", *Cognitive Psychology*, 19, 342-368.
32. Malek, R. (1981). The grip and its modalities. In R. Tubiana (Ed.), *The Hand*, Vol. 1. Philadelphia: W. B. Saunders, pp. 469-480.
33. Mason, M.T., (1981). "Compliance and Force Control." *IEEE Transaction on Systems, Man and Cybernetics*.
34. Mason, M.T., (1986). "Mechanics and Planning of Manipulator Pushing Operations." *Int. Journal of Robotics Research*, Vol. 5, no.3 pp. 53-71.
35. Mason, M.T. and Wang, Y. (1988). *Proceedings 1988 IEEE International Conference on Robotics and Automation* Philadelphia, PA, pp.524-528.
36. Mason, M.T. and Salisbury, J.K. (1985). *Robot Hands and Mechanics of Manipulation*, The MIT Press, Cambridge, Massachusetts.
37. Napier, J. R. (1956). "The prehensile movements of the human hand." *Journal of Bone and Joint Surgery*, 38B, 902-912.
38. Paul, R.P. and Schimano, B. (1976). "Compliance and Control," ACC San Francisco.
39. Peshkin, M.A. and Sanderson, A.C. (1988). "Minimization of Energy in Quasistatic Manipulation. *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.421-426.
40. Raibert, R. and Craig, J.J. (1981). "Hybrid Position/Force Control of Manipulators. " *ASME Journal of Dynamic Systems, Measurement and Control*.
41. Routh, E.J., (1960). "Dynamics of a System of Rigid Bodies." MacMillan, London.
42. Russel, Andrew R. (1986). "A Thermal Sensor Array to Provide Tactile Feedback for Robots", *Int. J. Robotics Res.* 4(3):35-39.
43. Siegel, D. M., Drucker, S.M., Garabieta, I. (1987). "Performance Analysis of a Tactile Sensor", *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 1493-1499.
44. Stansfield, S.A. (1987). "Visually Guided Haptic Object recognition," Ph.D Dissertation, University of Pennsylvania, Philadelphia.
45. Tsikos, C. (1987). "Segmentation of 3-D scenes using multimodal interaction between machine vision and programmable mechanical scene manipulation," Ph.D. dissertation, University of Pennsylvania, Philadelphia.
46. Tise, B. (1988). "A Compact High Resolution Piezoresistive Digital Tactile Sensor," pp.760-764.
47. Tou, J. (1953). "Coulumb and Static Friction in Servo Mechanisms," Ph.D. Thesis, Electrical Engineering Department, Yale University.

48. Tou, J. and Schulteiss, P.M. (1953). "Static and Sliding Friction in Feedback Systems", *Journal of Applied Physics*, 24 (9):1210-17.
49. Townsend, W. and Salisbury, J.K. (1987). "The Effect of Coulumb Friction and Sticktion on Force Control," *Proceedings International Conference on Robotics and Automation*, Raleigh, North Carolina, pp.883-889.
50. Tustin, A., (1947). "of Closed-Cycle Control Systems", *Journal of the Institutions of Electrical Engineers*, 94(2A):143-151.
51. Ulrich, N., Bajcsy, R. and R. Paul (1988). "A Medium-Complexity Compliant End Effector", *IEEE Conference on Robotics and Automation*, Philadelphia, PA.
52. Whitney, D.E. (1977). "Force Feedback Control of Manipulator Fine Motions" *ASME Journal of Dynamic Systems, Measurement and Control*, 91-97.
53. Winger, J.G. and Lee, K-M (1988). "Experimental Investigation of a Tactile Sensor Based on Bending Losses in Fiber Optics." *Proceedings 1988 IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.754-759.
54. Xu, Y., Paul, R.P. (1988). "On Position Compensation and Force Control Stability of a Robot with a Compliant Wrist," *Proceedings IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.1173-1179.
55. Yoshikawa, T. and Nagai, K. (1988). "Analysis of Grasping and Manipulation by Multifingered Robot Hands," Workshop on Dextrous Robot hands, *Proceedings of the IEEE International Conference on Robotics and Automation*, Philadelphia, PA, pp.149-171.
56. Zadeh, L.A. (1962). "From Circuit Theory to Systems Theory", *Proc. IRE*, Vol. 50, pp. 856-865, 1962.

Footnote

During the preparation of this paper, we have come to realize that our diverse disciplines have overlapping goals, but different bodies of knowledge. We acknowledge the support of the Office of Naval Research, who by awarding a contract to us, enabled us to attempt this synthesis. Our efforts are fully collaborative, and our contributions to this paper are equal. Reprints may be obtained by writing to Dr. R. Bajcsy at the University of Pennsylvania, Dept. of Computer and Information Science, Philadelphia, PA 19104, or Dr. S. Lederman, Dept. of Psychology, Queen's University, Kingston, Ontario, Canada K7L 3N6; or Dr. R. Klatzky, University of California, Santa Barbara, CA 93106.