October 1991

A Multiagent System for Intelligent Material Handling

Ruzena Bajcsy
University of Pennsylvania

Richard P. Paul
University of Pennsylvania

Xiaoping Yun
University of Pennsylvania

R. Vijay Kumar
University of Pennsylvania, kumar@grasp.upenn.edu

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

Recommended Citation

Ruzena Bajcsy, Richard P. Paul, Xiaoping Yun, and R. Vijay Kumar, "A Multiagent System for Intelligent Material Handling", . October 1991.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-91-84.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/437
For more information, please contact repository@pobox.upenn.edu.

A Multiagent System for Intelligent Material Handling

Abstract

The goal of our research is to investigate manipulation, mobility, sensing, control and coordination for a *multiagent robotic system* employed in the task of material handling, in an unstructured, indoor environment. In this research, manipulators, observers, vehicles, sensors, and human operator(s) are considered to be agents. Alternatively, an agent can be a general-purpose agent (for example, a six degree of freedom manipulator on a mobile platform with visual force, touch and position sensors). Possible applications for such a system includes handling of waste and hazardous materials, decontamination of nuclear plants, and interfacing between special purpose material handling devices in warehouses.

The fundamental research problems that will be studied are *organization*, or the decomposition of the task into subtasks and configuring the multiple agents with appropriate human interaction, *exploration*, or the process of exploring geometric, material and other properties about the environment and other agents, and *coordination*, or the dynamic control of multiple agents for manipulation and transportation of objects to a desired destination.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-91-84.

**A Multiagent System For Intelligent Material
Handling**

**MS-CIS-91-84
GRASP LAB 281**

**Ruzena Bajcsy
Richard Paul
Xiaoping Yun
Vijay Kumar**

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

October 1991

A Multiagent System for Intelligent Material Handling

Ruzena Bajcsy, Richard Paul, Xiaoping Yun and Vijay Kumar
General Robotics and Active Sensory Perception
(GRASP) Laboratory
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104

Abstract—The goal of our research is to investigate manipulation, mobility, sensing, control and coordination for a multiagent robotic system employed in the task of material handling, in an unstructured, indoor environment. In this research, manipulators, observers, vehicles, sensors, and human operator(s) are considered to be agents. Alternatively, an agent can be a general-purpose agent (for example, a six degree of freedom manipulator on a mobile platform with visual, force, touch and position sensors). Possible applications for such a system includes handling of waste and hazardous materials, decontamination of nuclear plants, and interfacing between special purpose material handling devices in warehouses.

The fundamental research problems that will be studied are *organization*, or the decomposition of the task into subtasks and configuring the multiple agents with appropriate human interaction, *exploration*, or the process of exploring geometric, material and other properties about the environment and other agents, and *coordination*, or the dynamic control of multiple agents for manipulation and transportation of objects to a desired destination.

I. INTRODUCTION

The goal of this paper is to outline a multiagent robotic system employed in the task of material handling, in an unstructured, indoor environment. In this research, manipulators, observers, vehicles, sensors, and human operator(s) are considered to be agents. Alternatively, an agent can be a general-purpose agent (for example, a six degree of freedom manipulator on a mobile platform with visual, force, touch and position sensors). Mobility is considered to be essential — if an agent is not mobile, it must be possible for it to “piggy-back” on another agent which is mobile. In addition there is a central station which is stocked with a variety of additional sensors, means of illumination, special effectors or tools, that the agents can employ depending on the environment, task and the outcome of the execution of the task.

Some possible applications of such a material handling system and examples of the tasks and environments are outlined below:

- Interfacing between special purpose (but “inflexible”) material handling devices (such as conveyor belts and part feeders) and sophisticated, but stationary, special purpose work cells (which could contain manufacturing machines or robots). Such a system would play the role of the human operator by off loading

components from the special purpose material handling equipment, loading workpieces on fixtures for machines/work cells, and fetching appropriate tools. In addition, the agents can perform tasks such as palletizing, retrieval from inventory, storage and clean-up operations.

- Handling of waste and hazardous materials in nuclear sites. This involves transportation of chemicals, waste material, old radioactive equipment and other toxic substances. Such a system could also perform routine inspections with a variety of sensors (including visual, temperature, touch, radiation and chemical sensors) and monitor, for example, air quality or radiation levels. It could be used for buried waste retrieval, in which case it would be able to excavate, explore, identify and sort objects, in addition to being able to transport objects. Similarly, decontamination, which requires disassembly followed by removal of old equipment, would be facilitated with such a material handling system.

Central to the multiagent concept is *organization* or the decomposition of the task into subtasks and configuring the multiple agents optimally with appropriate human interaction. For example, this includes the determination of the number of agents required, their spatial and temporal distribution, the required effectors/sensors/tools and the organization of the task. With each task, there are two key aspects: *exploration* and *task execution*. These two phases are not independent and must be interwoven (and indeed, in some situations they can be carried out concurrently). The exploration allows the system to gather information about its environment including, for example, the type of material and the geometry of the object that must be handled. Of course, it is possible, that based on this information, a reorganization of the task may be required. In the second phase of task execution, the goal is the final destination of the transported object. It is assumed that a human operator outlines the general path from the initial to the final position, but local modifications of this path and redirecting during the transport is permitted if the need should arise.

In Figure 1 we show a typical scenario involving multiple agents. For example, consider the lifting and transportation of an object, say a long pipe (approximately 5 meters long, 15 cm OD, 25 kg mass) inside a warehouse. Assume that the approximate location of the pipe (for example, the location of the room where the pipe is stored) is known but the exact position is not available. The multiple agents are first *organized* into an *exploration* task

Acknowledgements: Navy Grant N0014-88-K-0630, AFOSR Grants 88-0244, AFOSR 88-0296; Army/DAAL 03-89-C-0031PRI; NSF Grants CISE/CDA 88-22719, IRI 89-06770; and Du Pont Corporation

Fig. 1: A Typical Task Scenario of Multiple Agents

in which they attempt to locate the desired object. In the process, they encounter several objects and using, for example, touch and vision, they are able to discriminate against "wrong" objects. An observer agent, periodically informs the human operator, who is a "super agent" in this scenario, of the status of the system. For example, this can be done *via* a graphical display. If necessary, the operator can intervene by interrupting the search and providing more information in order to speed up the search. If the display consists of visual images obtained by the observers(s), he/she can send agents directly to the object. Once the object is located, relevant properties, including the size, shape, position and orientation of the pipe are obtained. Agents with manipulatory and tactile capability extract necessary material and mechanical properties. The mass (weight) and inertia of the pipe are first inferred from size, shape and material information. With this information, the agents are *organized* for the *task execution* phase. Based on the payload characteristics and the capabilities of each agent, the system determines the number of agents, the stance (pose) for each agent, and allocates the load between the agents for the lifting task. This selection is based on a sufficing principle, although optimality and redundancy considerations are important. It is possible that the lifting task cannot be accomplished due to a poor estimate of the payload. Or, if the pipe is flexible, a failure may be reported due to a sag in the pipe. In either case, the system assigns additional agents and reconfigures them automatically. For example, if the pipe sags, additional agent(s) are commanded to support the pipe at appropriate points in order to remedy the problem. If the execution of the lifting task fails, the observer, which perceives that the system is not functioning as desired, could alert the operator. Alternatively, if the op-

erator continuously monitors the system, although at a much lower bandwidth, he/she is able to intervene to assist the multiagent system, if necessary. The approximate path for the transportation is specified, possibly by the human operator. One or more observers position themselves appropriately so that their view is not occluded by obstacles. The proprioceptive sensing in each agent along with force/torque/tactile information allows dynamic, coordinated, control in the manipulation task. Obstacles in the path can be cleared by agents that are not involved in the transportation. If these obstacles are large, a modified path is sought by the system. In this approach, the multiagent system is intelligent in the sense that it is capable of learning by exploring and the agents can coordinate with each other. At the same time, the framework allows the robotic agents to interact with a human agent(s), who possesses superior intelligence. As time progresses, this interaction is reduced — the multiagent system becomes more sophisticated while the role of the human operator is reduced. This synergism increases the reliability, facilitates programming, and makes it possible to have a working system with less effort. What follows is an elaboration of some of our ideas and preliminary results in exploration, human interaction with robotic agents as well as robotic agents functioning in autonomy, and representation of the task and agents.

II. INTELLIGENT EXPLORATION OF ENVIRONMENTS

The ultimate goal is to build an intelligent material handling system that can function in partially or completely unstructured environments. It is essential to incorporate in the multiple agent system the ability to explore in an unknown environment for two reasons. First, without adequate knowledge of the system, the work organization, and the spatial distribution of agents for a desired task (for example, transporting a large object such as a pipe) may not be possible. Second, the controller design, irrespective of the actual control algorithms used, requires an accurate dynamic model of the agent(s) and the environment (including the objects to be handled). In most cases, the performance of the algorithms is sensitive to uncertainties and unmodeled dynamics, and in unstructured environments, a model of the environment is typically not available. Further, although a known model of the agent usually resides in the controller for that agent, this model is usually not available to other agents.

A. Representation

The exploration task involves the determination of environmental properties which may be categorized into three classes:

- *geometric properties*
such as shape, size and volume of an object.
- *material properties*
for example, stiffness, viscosity, inertia, static and kinetic coefficients of friction and compressibility.
- *kinematic properties*
for example, the mobility (the number of independent parameters that describe the configuration of the system) of an unknown linkage in the environment, or the geometry of an uncalibrated agent.

B. The Exploration Task

Exploration is not just a simple problem in system identification [1]. First, it must be driven by the knowledge of the specific task for which the information is required. And, to avoid exception handling and the resulting combinatorial explosion, the investigation must be sensor driven. Second, when no *a priori* knowledge is available about the external environment, the control of the robot for the manipulation task (which now involves interaction with an unknown external system) that is inherent in the exploration poses a problem. This is because we are now requiring that the two functions of investigation as well as execution be carried out simultaneously and both these functions lead to conflicting demands on the controller. Finally, the problem of exploration is intrinsically different from measuring the required properties in a physics laboratory using standard measurement instruments since the identification process by a robot must not only be performed in real-time, but it also must not rely on sophisticated cumbersome equipment that is practical only in a laboratory environment.

1). *Exploratory Procedures* Psychological studies [13] have indicated that haptic sensing is accomplished through a set of patterns of hand movement, called exploratory procedures. We have approached the problem of exploration in robotics by trying to establish such a set of exploratory procedures (EP's) [1]. The basic objective here is to develop a "bottom-up" approach to exploration using such a set of EP's. The exploratory procedures are constructed and implemented through a series of *motion primitives* or control algorithms. Our EP's differ from those described in [13] since the motivation in this project (incorporating intelligence into material handling systems) is quite different from that leading to the work in [13] (study of human behavior). Further our EP's are not necessarily based only on haptic recognition - we allow for position, velocity, force, touch and visual sensing. Indeed it is quite possible and desirable to allow for flow of information between different sensing modalities. For example, the effects of manipulation can be felt through force and touch sensors as well as seen by cameras that observe deformations and other movements.

In order to describe the geometric properties of the environment, we have developed a new surface and volume segmentation algorithm that uses three dimensional data points (obtained from a sensor such as a laser scanner) to develop a best possible description (with minimum residuum within the desired error tolerance) in terms of the parametric shape primitives. The surfaces are fitted to a constant, planar or biquadratic function [14], while the volumes are fitted to a superellipsoid [3, 10, 9]. This segmentation algorithm can deal with an arbitrary scene of multiple objects and parts, each of which is decomposed into individual superellipsoids and surfaces as described above. All this is done without any *a priori* assumption on the objects and or scene.

We have made considerable progress in the extraction of material properties. Stansfield [4, 19] has demonstrated EP's for extracting the compliance of an unknown object and characterizing the behavior as plastic or elastic. Tsikos [21] treated manipulation as a physical segmentation exploratory procedure. He showed that the connection between perception and action in a simple manipulatory world can be adequately modeled by a non-deterministic finite state automaton, very similar to the

work of Brooks [6]. Campos [7] has developed and integrated a set of exploratory procedures that are more tailored to robotic systems rather than humans. In particular, a new thermosensitive sensor gives this system the ability to discriminate between different materials, such as metallic, plastic and others. This, in conjunction with the EP for estimating shape/volume will allow the determination of the weight of the object. EP's for establishing the hardness are driven by measurements of strain and stress in orthogonal directions (as opposed to only measuring pressure [19]).

C. Multiagent Exploration

In any exploration task, the deployment of multiple agents makes it possible to introduce redundancy and therefore, improved reliability, and efficiency at the cost of increased complexity. Consider the exploration of a contaminated site. The whole area is visually scanned for all "interesting" objects. If an object is not recognized, other sensing modalities are employed to extract more information. For example, by touching, the material properties can be learned. When the combination of manipulatory and visual exploratory procedures fails, the operator is alerted and he/she determines the identity of the object. This search procedure is most efficiently performed by multiple agents. And in general, different exploratory procedures can be pursued in parallel by multiple agents. The parallelism and the data driven nature of the exploration process make its organization a challenging problem. The optimal (or near-optimal) organization of this process will be pursued in our study.

Sometimes, the exploratory procedure for a single property is intrinsically a multiagent task. For example, consider again the exploration of a mechanical assembly such as a pair of shears or vice grips. Here it is necessary to identify the nature of the mobility in the assembly. In order to induce relative motion between the components [12], one manipulator must hold one end while another robot grasps the other end and manipulates it. The concept of one agent holding and securing an object while the other explores is quite general and can be seen in humans too. Similarly, disassembly of a mechanical assembly with the objective of exploring and learning requires more than one agent. The cooperation between the agents and the coordinated control will be studied in the course of this project.

Exploration often requires tight coupling between different sensing modalities [17]. For example, vision can provide starting points for exploration, since visual sensors encode (rather quickly and simply) positional and orientational information of the object as well as shape parameters, such as surface descriptors [2]. The effect of manipulation (rubbing, pressing, or inducing relative motion between components) can be detected through visual sensing. In fact the sequence of visual examination, manipulation followed by another visual inspection of the same scene and the process for detecting the changes caused by manipulation is a very basic sequence. The visual sensor drives the manipulator and the manipulator's actions drive the visual examination process. Similarly, coupling between touch sensors and manipulators with position sensors is beneficial. We will develop a general control framework which accommodates coupling between the different sensing modalities.

In summary, the multiagent approach is a natural and versatile approach to exploration. Our preliminary work

has given us a good understanding of the problems underlying the exploration task. Future work will be directed at coordinating multiagent exploration, integration of different sensing modalities and task driven exploration.

D. Human Agent

We shall use a generalization of the *teleprogramming* technique [8] to allow the human agent(s) to interact with the other agents. Teleprogramming is used to control remote robotic workcells by providing the human operator with a graphical simulation of the remote environment, offering immediate visual and kinesthetic feedback, regardless of the communication bandwidth (and thus possible delays) with the actual remote manipulator. Operator's actions in the geometric model are interpreted within the context of a given task and automatically translated into a stream of instructions for execution by a remote agent, possibly delayed in time.

Previous approaches to human-assisted remote control of robotic systems, especially in situations involving a time delay, rely on predictive displays, which offer only visual feedback [5], extensive dynamic simulation of the remote environment [11], and relatively low-level (servo or trajectory level information) communication with the remote workcell. The design of the teleprogramming control paradigm represents a departure from these techniques in that it offers immediate kinesthetic, as well as visual, feedback to the operator. Further, detailed knowledge of the environment dynamic properties is not necessary, and that it communicates with the target system by sending a stream of symbolic instructions.

In the case of the multiple agents, described in this paper, the target system can be any agent (manipulator, vehicle, sensor, etc.). The operator can move the agents around in a geometric model at rates that may be much higher than the corresponding execution rates. If the operator can do this then he or she can organize the activity of many agents simultaneously by sequentially attending to different agents, which in turn follow along at a slower execution rate.

As in teleprogramming, the operator is provided with force feedback as well as providing for positional input. Thus if an agent became "stuck" then the operator could be alerted by the agent itself or the observer agent and the master input device attached to the agent so that the operator could "feel" the constraints on the motion of the agent and could provide guidance, at a lower level, to extricate the agent.

In this manner an operator could organize many agents and could supervise the overall activity. If needed, the operator could move to an even lower level of control, directly causing the agent to exert forces and to initiate motions. Object pick-up could also be handled directly by the operator in the same manner in which we handle the control of a single manipulator.

As agent autonomy develops the operator may be able to interact at a higher level with the agents organizing themselves to perform tasks. If agent autonomy is less than expected the operator may interact at a lower level, such as in planning individual trajectories for each agent.

E. Representation of Tasks and Agents

We characterize a task by looking at its decomposition into subtasks. We have the following four cases.

- **Single subtask:** In this case, all the available resources (agents, sensors and effectors in the central station)

can be allocated for the task, and there is no penalty for selecting one agent versus other agents. The goal is to choose an appropriate number of agents which will suffice to achieve the task. Optimality is not a major concern. In contrast, redundant agents will be employed to increase the robustness and reliability of the system.

- **Spatially distributed multiple subtasks:** In this case, multiple subtasks need to be carried out simultaneously (at least starting at the same time). The available resources have to be shared among all the subtasks.
- **Temporally distributed multiple subtasks:** The multiple subtasks are initiated sequentially, but the performance durations of subtasks may overlap. In this case, while organizing a subtask, it may be necessary to "save" an agent(s) for subtasks that are to be initiated later. Also, it is necessary to plan for possible failures of subtasks and delays, which could result in unforeseen overlaps between processes.
- **Spatially and temporally distributed multiple subtasks:** This is the combination of the last two cases.

In order to perform multiple subtasks simultaneously, agents will need to be distributed among the different subtasks. For example, while lifting a large object, it is first required to determine the number and type of agents which must be employed. While in simple cases, this problem can be easily solved by the operator, in general, the problem is complicated, especially since the strength of an agent varies as a function of its configuration or the position in its workspace. It is evident that, for example, algorithms derived from workspace and dynamics considerations [22, 15, 20] can serve as aids for the human operator. The key questions are:

- What information about each subtask must be included in its representation?
- What are the important considerations (kinematic, dynamic, workspace, strength, mobility, type of end-effector, sensing capabilities, etc.) for determining the allocation of the agents between each subtask?

From the second question, it is clear that it is necessary to establish a model, i.e., a database which has relevant information about each agent. The representation of a manipulator agent could, for example, include its strength capacity, payload, number of degrees of freedom, type of end-effector and sensors, and the size of workspace. The important point is that while specificities, such as the exact kinematic and dynamic properties need not be included in an agent's representation, it should be possible to infer the manipulator's general capabilities from its representation. The manufacturer's specifications, or their equivalent, provide a starting point. The exact pieces of information that should be included in the database and how to represent an agent's capabilities are research problems that will be studied.

F. Autonomous Operations of Agents

The level of autonomy that different agents possess may vary. An agent without navigation capability will be either used in a workspace that is relatively free of obstacles, or it will be used in conjunction with another agent which

has some navigation capability. Nevertheless, our minimum requirements of the agents in the multiagent system are that any mobile agent is able to follow a path specified by the operator, any manipulator agent can track a position or a force trajectory, while any sensory agent can position itself in a desired position in order to obtain relevant information.

When different agents must cooperate, deciding on an appropriate number of agents for a particular task is a key problem. We will pursue a sufficing requirement, but we will prefer near-optimal solutions to this problem. Even when this is resolved, the organization of the task involves the spatial and temporal distribution of the different agents. While it is clear that considerations of task dynamics and the capabilities of each agent are required, there is no obvious method of pursuing this problem. This is a research problem that will be studied.

Another important research problem in the organization is the deliberate, but judicious, use of redundancy for robustness and reliability. Deploying more agents than the minimum to a task increases the reliability of the system in case of the failure of an agent. For example, if two agents are marginally capable of lifting an object, the use of three agents will increase the robustness to errors that might have been generated during the exploration process.

The degree of coordination between human and robotic agents is an important research issue. While each agent has some ability to operate autonomously, there can be situations in which human intervention can be requested. Once possibility is when a set of agents are "stuck" and must be extricated from a situation that they are not equipped to deal with. This can happen when a tool or end-effector gets jammed, an agent fails, or when the observer is incapable of observing the task (for example, its view is occluded by an obstacle that it cannot circumvent). Another likely event is a change in the (dynamic) environment which makes the path that was previously designated by the operator impossible to follow. When the operator is summoned, he/she can resolve the problem by repositioning the agents or reorganizing the task. In each situation, the level of interaction can be different. It can be in the form of a minor high-level change in the organization or specification of a subtask(s), or could be some type of low-level interaction with a specific agent.

G. Monitoring — the Role of an Observer Agent

The function of an observer agent is to *monitor*, or *observe* the correct execution of the task or subtask. For example, the task or subtask can be: holding an object cooperatively or following an agent at a fixed distance. These tasks and subtasks are modeled as discrete event dynamic systems (DEDS) [16, 18]. In this case, the states are relations between the agents and the manipulated object and those between an agent and the environment. The events are movements of the agents, which cause the state to change. The desirable states will be those that are required for the successful execution of the task, for example, holding the object in an upright position. The undesirable states will be those situations that should be avoided. For example, these include situations in which constraints on a manipulatory task are violated or ones in which the observer is ill positioned with respect to the task and the desired view is obscured. In the first case the observer reports the problem and possibly alerts the operator about the undesirability of the executed action. And in the second case the observer attempts to correct its position and orientation. The DEDS theory gives us a

powerful framework that allows us to infer the observability of the system which is then used to organize then the multiple agents.

III. CONCLUSIONS

We are presenting here an outline for a multi-agent robotic system employed in the task of material handling in an unstructured, though indoor environment. An agent can be: human(s), robotic manipulator(s), vehicle(s) or and observer such as a camera system. Question can be raised: Why multi-agent when there are still many unsolved problems with a single agent? Our answer is: in order to reduce the weight, thereby the dynamics, flexibility, speed of performance and the cost and yet keep the payload, one must consider distributed manipulatory agent systems. These agents must work in cooperation. Until recently the problem of control amongst many active agents was unsolved. We are now beginning to attack successfully these control problems which in turn enables us to deal with distributed agent systems. However, in an unstructured environments there are too many unpredictable situations that at this time it is not practical to have a completely autonomous system. Hence we propose a hybrid system where one or more agents is a human. This by itself poses some interesting problems in terms of communication, representation and coordination among humans and robot-agents, which we have tried to identify. We have some preliminary results, but are only at the beginning.

REFERENCES

- [1] R. Bajcsy, S. Lederman, and R. L. Klatzky. *Machine Systems for Exploration and Manipulation: A Conceptual Framework and Method of Evaluation*. Technical Report MS-CIS-89-03, Department of Computer and Information Science, University of Pennsylvania, January 1989.
- [2] R. Bajcsy and F. Solina. Three dimensional object representation revisited. In *Proceedings of the First International Conference on Computer Vision*, pages 231–240, London, England, June 1987.
- [3] R. Bajcsy, F. Solina, and A. Gupta. *Analysis and Interpretation of Range Images*, chapter Segmentation versus Object Representation – Are they Separable? Springer-Verlag.
- [4] R. Bajcsy and S. Stansfield. Object apprehension using vision and touch. In *Proceedings of the JPL/NASA Space Telerobotics Workshop*, Pasadena, CA, January 1987.
- [5] A. K. Bajczyk, W. S. Kim, and S. C. Venema. The Phantom robot: Predictive displays for teleoperation with time delays. In *IEEE International Conference on Robotics and Automation*, pages 546–551, Cincinnati, Ohio, 1990.
- [6] R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, 1986.
- [7] M. Campos and R.K. Bajcsy. *A robotic haptic system architecture*. Technical Report MS-CIS-90-51, University of Pennsylvania, Dept. of Computer and Information Science, Philadelphia, Pa., June 1990.
- [8] J. Funda and R. P. Paul. Teleprogramming: overcoming communication delays in remote manipulation. In *Proceedings of the First Workshop on Mobile Robots for Subsea Environments*, October 1990. Monterey, CA.
- [9] Alok Gupta and Ruzena Bajcsy. Part description and segmentation using contour, surface and volumetric primitives. In *Proceedings of the Conference on Sensing and Reconstruction of 3D Objects and Scenes*, pages 203–214, SPIE, Santa Clara, CA, Feb 1990.

- [10] Alok Gupta, Luca Bogoni, and Ruzena Bajcsy. Quantitative and qualitative measures for the evaluation of the superquadric models. In *Proceedings of the IEEE Workshop on Interpretation of 3D Scenes*, pages 162–169, Austin, TX, 1989.
- [11] G. Hirzinger, J. Heindl, and K. Landzettel. Predictive and knowledge-based telerobotic control concepts. In *IEEE International Conference on Robotics and Automation*, pages 1768–1777, 1989.
- [12] Vijay R. Kumar, Xiaoping Yun, and Ruzena Bajcsy. Exploration of unknown mechanical assemblies through manipulation. In *Advances in Intelligent Robotics Systems Symposium, SPIE*, Philadelphia, PA, November 1989.
- [13] S. J. Lederman and R. L. Klatzky. Hand movements: a window into haptic object recognition. *Cognitive Psychology*, (19):342–368, 1987.
- [14] Ales Leonardis, Alok Gupta, and Ruzena Bajcsy. Segmentation as the search for the best description of the image in terms of primitives. In *Proceedings of the Third International Conference on Computer Vision*, pages 121–125, IEEE, Osaka, Japan, December 1990.
- [15] Y. Nakamura. Minimizing object strain energy for coordination of multiple robotic mechanisms. In *Proceedings of 1988 American Control Conference*, pages 499–504, Atlanta, GA, June 1988.
- [16] C.M. Ozveren. *Analysis and Control of Discrete Event Dynamic Systems: A State Space Approach*. PhD thesis, Massachusetts Institute of Technology, August 1989.
- [17] R. Paul, H.F. Durrant-Whyte, and Max Mintz. The integration of sensing with actuation to form a robust, intelligent robot control system. In *Third International Symposium of Robotics Research*, Gouvieux, France, 1985.
- [18] T. M. Sobh and R. Bajcsy. Visual observation of a moving agent. In *Proceedings of the European Robotics and Intelligent Systems Conference (EURISCON '91)*, Corfu, June 1991.
- [19] S. A. Stansfield. A robotic perceptual system utilizing passive vision and active touch. *International Journal of Robotics Research*, 7(6), April 1988.
- [20] T. J. Tarn, A. K. Bejczy, and Z. Li. Dynamic workspace analysis of two cooperating robot arms. In *Proceedings of 1988 American Control Conference*, pages 489–498, Atlanta, GA, April 1988.
- [21] C.J. Tsikos and R.K. Bajcsy. *Redundant Multi-Modal Integration of Machine Vision and Programmable Mechanical Manipulation for Scene Segmentation*. Technical Report MS-CIS-88-41, University of Pennsylvania, Dept. of Computer and Information Science, Philadelphia, Pa., 1988.
- [22] I. D. Walker, R. A. Freeman, and S. I. Marcus. Dynamic task distribution for multiple cooperating robot manipulator. In *Proceedings of 1988 International Conference on Robotics and Automation*, Philadelphia, PA, April 1988.