

12


Morphing Hands and Virtual Tools (or What Good is an Extra Degree of Freedom?)

Olac Fuentes and Randal C. Nelson

Technical Report 551
December 1994



UNIVERSITY OF
ROCHESTER
COMPUTER SCIENCE



19950118 079

Morphing Hands and Virtual Tools (or what good is an extra degree of freedom)

Olac Fuentes

Randal C. Nelson

The University of Rochester
Computer Science Department
Rochester, New York 14627

Technical Report 551

December 1994

Abstract

Manipulators with large numbers of degrees of freedom, from the human hand to the trunk of an elephant, are common in the biological world. These freedoms allow highly flexible and robust performance of complex tasks. However, progress in developing and controlling artificial high-degree-of freedom manipulators has been slow. The main problem is that traditional robotics has focussed on the solution of systems of kinematic equations where there is a unique solution. Such approaches tend not to generalize well to situations with a high-dimensional solution space, and controlling redundant systems has acquired a reputation as a hard problem. However, this need not be the case. In this paper, we describe a behavioral method for using extra degrees of freedom to simplify rather than complicate manipulation problems, while at the same time obtaining more flexibility than would be available with a simpler system. The method is developed in the context of a high DOF robot hand, but it has the potential to generalize to other sorts of manipulators.

The basic idea is based on the observation that, for a particular task, using a custom-designed fitting can greatly simplify the control problem. Using a wrench sized for a particular nut is an extreme example. We use the extra degrees of freedom to dynamically configure or "tailor" the manipulator to match the particular object and task at hand. This creates a *virtual tool*. The tailoring is accomplished by imposing low-level, task-specific constraints on the degrees of freedom. These constraints are selected dynamically from a large set of potential constraints in response to the demands of the current task. The process of smoothly transitioning from one virtual tool to another in the course of task execution is referred to as *morphing*. We apply the technique to the control of a 16-DOF Utah/MIT hand, and perform fine manipulations on a range of objects using virtual tools that are dynamically instantiated on the basis of sensory information.

This material is based upon work supported by DARPA under research grant no. MDA972-92-J-1012, and ONR under research grant no. N00014-93-1-0221. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or ONR.

Accession For		
NTIS	CRA&I	☒
DTIC	TAB	☐
Unannounced		☐
Justification _____		
By _____		
Distribution /		
Availability Codes		
Dist	Avail and/or Special	
A-1		

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE December 1994	3. REPORT TYPE AND DATES COVERED technical report		
4. TITLE AND SUBTITLE Morphing Hands and Virtual Tools (or What Good is an Extra Degree of Freedom?)		5. FUNDING NUMBERS N00014-93-1-0221, MDA972-92-J-1012		
6. AUTHOR(S) O. Fuentes and R.C. Nelson				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESSES Computer Science Dept. 734 Computer Studies Bldg. University of Rochester Rochester NY 14627-0226		8. PERFORMING ORGANIZATION		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESSES(ES) Office of Naval Research ARPA Information Systems 3701 N. Fairfax Drive Arlington VA 22217 Arlington VA 22203		10. SPONSORING / MONITORING AGENCY REPORT NUMBER TR 551		
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution of this document is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) (see title page)				
14. SUBJECT TERMS redundant manipulators; dextrous manipulation; telemanipulation; grasping			15. NUMBER OF PAGES 25 pages	
			16. PRICE CODE free to sponsors; else \$2.00	
17. SECURITY CLASSIFICATION OF REPORT unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT unclassified	20. LIMITATION OF ABSTRACT UL	

1 Introduction

1.1 Redundant Manipulators

The appendages that animals use to manipulate the physical world typically possess many degrees of freedom and are redundant. More precisely, they possess more degrees of mechanical freedom than are controlled in the objects that are manipulated. The most familiar example is the human arm/wrist system, which has seven primary degrees of freedom and is typically used to position and orient objects in space, that is to say, in six degrees of freedom. In this case there is a single redundant degree of freedom. When both arms are used together on a single object, a common situation, the degree of redundancy is even higher; but people appear to have no more difficulty in figuring out how to manipulate rigid objects with two hands than with one. If anything, tasks are subjectively easier. The human hand is typically used to perform manipulations that require controlling only a few degrees of freedom simultaneously, for example squeezing, twisting, even writing, yet it has twenty plus degrees of mechanical freedom (the exact number depending on what reference is cited). Of course people can perform tasks, such as playing the piano, that appear, at least on the surface, to involve controlling many more degrees of freedom at once, but in terms of the activities required for day-to-day survival, these appear to be the exception.

Flexible manipulators are not limited to humans or primates. The trunk of an elephant, the tentacles of a squid, and the body of a snake are all systems with a large number of mechanical degrees of freedom that are used to interact with physical objects. Spiders use several multi-jointed legs cooperatively to manipulate and wrap prey. The locomotory system of walking animals utilize a large number of degrees of freedom, ranging from fourteen plus in humans, to hundreds in centipedes to accomplish a positioning task with at most six degrees of freedom. An interesting point is that the sophistication of the controller appears to be almost independent of the actual number of mechanical degrees of freedom. Spiders walk with eight legs, humans with two, but the spiders don't need a bigger brain to do it.

In contrast, the systems currently used in artificial robotics applications are generally constructed with as few interacting degrees of freedom as possible, and are rarely redundant. Most commercially available robot arms, for example, have four, five, or six degrees of freedom. A couple have seven, but generally no more. In many applications, a considerable amount of effort is invested to devise strategies that allow a system to be positioned in a multi-dimensional space using fewer controls than the number of dimensions of the space. A familiar example is a car, which is positioned in three dimensions (x, y , and orientation) using only two controls (steering and acceleration). Spacecraft are commonly positioned in a ten-dimensional space ($x, y, z, v_x, v_y, v_z, \omega_x, \omega_y, \omega_z, t$) using only three control parameters (two rotational and one linear acceleration). (Note that although these ten parameters are independently settable, they are not independently controllable). Recently a few robot devices having relatively many controllable degrees of mechanical freedom have become available as research tools, most notably dextrous grippers such as the Utah/MIT hand. However, there is, as yet, no generally accepted strategy for using them.

1.2 Why Use Redundancy?

A system with redundant degrees of freedom has certain theoretical advantages associated with the fact that extra degrees of freedom typically mean that there are (generally uncountably) many solutions to a particular problem. One advantage is increased robustness and flexibility. A simple demonstration of this involves loss of control over one or more of the controlled parameters. In a minimal system, this usually spells catastrophic failure; consider for example, the consequences of loss of control of one joint on a traditional 6 joint manipulator. A system with redundant degrees of freedom however, can often continue to perform adequately, depending of course on the exact mode of failure. More generally, a system with redundant degrees of freedom can cope with additional and unexpected constraints, that would cause a minimal system to fail. For example, a traditional 6DOF manipulator can position itself to pick up an object at any point and any orientation in its workspace. If there is clutter in the workspace however, the manipulator can be proximally blocked from assuming the unique configuration required to perform the task. An extra degree of freedom or two could very well allow the manipulator to position itself around the obstacle. In the case of a gripper, the presence of extra degrees of freedom in a device like a hand permits different types of grasps, appropriate for different objects or tasks to be used, rather than struggling to make a single, low degree of freedom grasp strategy work for all applications.

The case of the gripper also illustrates another advantage of mechanical redundancy, this one having to do with physical scaling. Grasping an object with a traditional parallel jaw or three point gripper exhausts the freedoms of the device. If fine manipulation is now needed, as in an insertion task, this motion must be provided by of the primary positioning system, demanding very high precision relative to the size of the device. Making such adjustments can also induce disproportionately large motions of portions of the manipulator, such as those that can occur in remote center rotations. Extra degrees of freedom in the gripper itself would allow tuning movements to be carried out at a more appropriate scale, reducing the required proportional precision and the possibility of mechanical thrashing. Finally, the presence of extra degrees of freedom in a system can allow it to avoid ill-conditioned control regions, or singularities during operation. In practice, this is a special case of dealing with additional constraints, but traditionally it rates a line of its own.

1.3 Previous Approaches to Redundant Control

At present, high degree of freedom systems are not much utilized. There are two main reasons for this. The first is that sets of non-linear, high-dimensional equations are, in general, very difficult to solve, especially under the real-time constraints involved in robot manipulation tasks. The second is the difficulty in selecting between the multiplicity of solutions. There has, however, been a considerable amount of theoretical work on the topic of redundant control. A general way of handling the multiple solution problem is by imposing enough additional constraints to make the problem solution unique. The problem with this is that it is not generally clear what constraints should be used. The most commonly used approach is to define a quality measure and attempt to find the

solution that maximizes it. Unfortunately, solving such optimization problems tends to be computationally intensive, making use in real-time control difficult.

An early formulation of the inverse kinematics of a redundant manipulator as an optimization problem was given by Liegeois [Liegeois, 1977]. The basic criterion was to keep the joint variables far from their physical limits. A somewhat different quality measure was described by Yoshikawa [Yoshikawa, 1990], who proposed the square root of the determinant of JJ^T , where J is the Jacobian matrix, as a quantitative measure of the "manipulability" of a manipulator (redundant or not) in a given configuration. The solution to the inverse kinematics is then defined as the set of joint angles (or velocities, or accelerations) that achieves the goal while maximizing JJ^T . This measure can also be used to avoid singularities. Another approach is to maximize the distance between the manipulator and various obstacles while still satisfying the goal conditions [Maciejewski and Klein, 1985; Yoshikawa, 1990]. An optimization method based on joint torques is described by Hollerbach and Suh [Hollerbach and Suh, 1985].

A few authors have discussed other aspects of redundant manipulators. Maciejewski [Maciejewski, 1990] discusses the design of redundant manipulators that can tolerate certain joint failures (e.g. a frozen joint). Several researchers have addressed the idea of using the extra degrees of freedom to accomplish several goals at once. One idea is to divide a task into subtasks according to priorities. Joint motions can then be determined so that subtasks with lower priority can be performed utilizing redundancy freedoms of subtasks with higher priority [Nakamura *et al.*, 1987; Yoshikawa, 1990]. Other authors have discussed combining multiple criteria by linear means [Cleary, 1990; Pamanes and Zeghloul, 1991], and probabilistically [McGhee *et al.*, 1994].

There has been a certain amount of work on the control of particular redundant manipulators, principally robot hands and snake-like devices. The redundant nature of hands has been addressed by Iberall *et al.* [Iberall *et al.*, 1988; Iberall, 1987], Bekey *et al.* [Bekey *et al.*, 1993], [Caselli *et al.*, 1993], and Stansfield [Stansfield, 1991], all of whom have used heuristics derived from human grasping to choose among several possible grasps. Snake-like manipulators are discussed by Chirikjian and Burdick [Chirikjian and Burdick, 1994] who use a notion of a "backbone curve" to describe the (time-varying) macroscopic configuration of the robot; actual joint angles are then determined by fitting the physical manipulator to the backbone curve. Hirose [Hirose, 1993] mentions the idea of using the additional degrees of freedom in a snake to allow the same mechanism to work as a gripper, manipulator or locomotor. Both these approaches have some elements in common with the notion of bundling the degrees of freedom together with task-derived constraints that we develop in this paper.

A novel approach to the problem of the complexity of the optimization techniques typically used for redundancy resolution was described by Davidor, who proposed using genetic algorithms. [Davidor, 1991]. This paper noted that, in the interests of efficiency, most other approaches deal only with instantaneous positions and yield suboptimal results over whole trajectories. The standard techniques also have trouble with local minima and non-differentiable objective functions. Genetic algorithms offer a potential solution. So far, this work is fairly preliminary.

2 The Virtual Tool Approach

2.1 Virtual Tools

In this paper, we describe a model referred to as a *virtual tool* that permits many of the advantages of high-degree-of-freedom systems to be realized, without incurring excessive computational cost. We then demonstrate the technique by applying it to the problem of dextrous manipulation using a highly redundant device, the Utah/MIT hand.

The basic idea of a virtual tool is drawn from the observation that the most efficient way of performing a task is generally to use a device that is specially designed for it. This is a fundamental principle of hard automation, and a primary reason that highly flexible robots have not been in much demand for mass production operations. The point is that most tasks, even very complicated ones, can be reduced to a sequence of operations that require the control of only a few, carefully engineered degrees of freedom. This suggests that a way to use extra degrees of freedom in a system, is to use them to customize, or “tailor” the system so that it becomes, in essence, a special purpose tool for the task at hand. We call the resulting instantiation a *virtual tool*. Formally, this tailoring takes the form of implementing constraints on the degrees of freedom of the system through a low-level control system. The remaining degrees of freedom appear as control parameters of the virtual tool, and form the interface used by higher-level control processes. Since they are directly linked to the task at hand and of low dimensionality, the control problem is greatly simplified. Furthermore, the device can shift facilely from one virtual tool to another as the system progresses through the stages of a complex task.

A simple, non-robotic example of such use of extra degrees of freedom is an adjustable wrench. An actively controlled system, however, has the advantage that it is possible to impose constraints more general than simply fixing a mechanical degree of freedom. For example, two adjacent joints in a planar chain linkage could be constrained to maintain equal angles, and the two degrees of freedom replaced by a single one specifying the total deviation through the two joints. A more complicated example would be constraining two 6DOF manipulators to maintain a constant distance between manipulator endpoints—a virtual link. The remaining degrees of freedom could then be usefully expressed in terms of the position and direction of the virtual link and three orientations for each arm with respect to the endpoints of the virtual link. This last example illustrates how imposing constraints can have the effect of remapping the “natural” coordinate system associated with the degrees of freedom of the system. This remapping is part of what we are after in terms of creating virtual tools. Not only shape and kinematic properties, but the control interface is tailored to the application of the moment.

The idea of a virtual tool is a generalization of some techniques that have been described previously for dealing with more specific situations. A good example is the notion of virtual fingers introduced by Arbib *et al.* [Arbib *et al.*, 1985]. A grasp involves the use of virtual fingers, which are composed of one or more real fingers working together to perform a task. Iberall [Iberall, 1987] shows how the hand can be used as essentially three different grippers by changing the mapping of virtual fingers to real fingers.

Sutherland and Ullner [Sutherland and Ullner, 1984] introduced the concept of virtual

legs for a six-legged walking robot. He observed that when two-legs act in unison, they can be thought of as a functionally equivalent virtual leg. In his work two legs could be conceptually combined through a series of constraints so that they would look to the rest of the system as a single leg. This allowed him to express and analyze the complex behavior of a machine with six legs in relatively simple terms. Raibert [Raibert, 1986] used the virtual leg concept to derive two and four legged gaits from algorithms originally designed for a one-legged hopping robot. He observed that in biped gaits at any given time only one leg is providing support while the other is elevated and out of the way. Thus each leg in biped walking machine could be controlled individually by essentially the same algorithms used for the one-legged machine.

2.2 Instantiation and Morphing

The principal issues with the virtual tool approach are how to go about instantiating a tool, and how to transition smoothly between virtual tools during execution of a task. Our approach to instantiation is to use sensory information to determine the constraints that define a virtual tool. Transition between virtual tools is accomplished by a process we refer to as *morphing*. The basic idea is to note that since virtual tools are associated with control parameters that change the physical configuration of the robot device, it is possible for two different virtual tools to have the same physical configuration for certain settings of the control parameters. At this point, a seamless transition can be made, and a different set of constraints swapped in. There is no physical discontinuity, but we now have a different set of control parameters, and different constraints. Note that morphing instantiates a new tool at the morph point. In keeping with our policy of using sensory information for instantiation, we use sensory information for the selection of morph points. This turns out to be the case, and the two processes are intimately linked through the sensorium.

Instantiation involves two natural levels of description. At the top level, is a description that corresponds to a class of tools that are used in similar ways. At the lower level, is a set of discrete parameters that describe how the particular tool is fitted to a particular instance of the problem. For example, in the domain of hand tools, a top level description might specify use of a 12 point box wrench (as opposed to a hammer, a punch, or an allen wrench). The lower level description would specify that it is a 5mm wrench. Both these decisions can be made on the basis of sensory information, but rather different levels of reasoning are involved.

The low level is the most straightforward to automate, and is one of the main topics of this paper. In essence, what we do is provide, *a priori*, a mathematical form that the constraints will take, along with a fixed number of free parameters and sensory processes for determining them. To take a simple example, consider a tool to execute a power grasp. In the simplest case this tool might have single associated control parameter that would represent the degree of closure. Position, orientation, and initial closure would represent instantiation parameters to be initialized on the basis of sensory measurements of the object to be grasped. In this example, instantiation places the tool in the correct orientation with respect to the object to be grasped. The closure parameter could then be decreased by the system until sensory input (e.g. tactile sensors) indicated that the object was securely

grasped. At this point, if we want to move the object we have just grasped, and have arm degrees of freedom, the grip tool could then be morphed into a manipulation tool. Such a tool would have constraints that caused it to maintain a secure grasp on the object — instantiated by the grasping process, and control parameters that specify the orientation of the grasped object. As another example, we might specify a planar rotation tool, with the plane of rotation and fitting of fingers as instantiation parameters, perhaps initialized through the use of a precision grasp tool that is morphed into the rotation tool when the grasp is completed.

The higher level is more problematic, and can itself be addressed at two levels. In the simplest application we assume a set of previously defined tool classes (mathematical forms) and the task of the decision system is simply to select the appropriate one (on the basis of history and sensory information) and pass it down to the lower level for full instantiation. This is the only approach used in this paper, and it is reasonably well defined, though the sensory and control systems could become quite complex (e.g. discrete event dynamic systems, conventional planning, Bayes nets, fixed sequence). Essentially, however, the design of the tools for a particular robot and a particular class of problems is in the hands of the programmer. Whether the set of tools is flexible and useful depends primarily on the ingenuity of this person.

A more complex problem would be to automatically determine or learn the mathematical forms themselves. This problem is considerably less well-defined and probably much more difficult. There are a number of approaches that might be tried (reinforcement learning, genetic algorithms, annealed search etc.) Ultimately, such a search could be cast as a parameter determination problem in some space. However, the structure of the space, and the information required to determine what values the parameters should have seem qualitatively different from what is needed for low-level parameter instantiation, and a separate conceptual level thus seems justified.

2.3 Relationship to Behavioral Robotics

The control approach we describe falls under the general topic of behavioral robotics. In behavioral robotics [Maes, 1992; Brooks, 1986; Brooks, 1991], the robotic system is viewed as always being engaged in some particular task or behavior. Rather than looking for control formulations that will permit equally simple specification and execution of all actions the robot is physically capable of, the behavioral approach considers the specific activities the robot will be called upon to perform, and looks for useful common components that can be abstracted out as primitives. Such generally useful activities are referred to as behaviors. Basically, the behavioral approach claims that the notion of a general purpose robot is an untenable concept; even the most flexible proposed systems have strong task assumptions implicitly embedded at all levels of their design. More to the point, the behavioral approach holds that there is much to be gained by making these task requirements explicit, and embedding them into the system at as low a level as possible. This can be seen as a strategy of most commitment.

A principle area of research in behavioral robotics is thus elucidating the requirements and assumptions associated with specific robot tasks, and determining how these can be

realized in terms of easily describable, interacting pieces. In terms of generating scientific knowledge, the hope is that certain of these pieces will turn out to be reusable in other systems, and thus a toolbox of useful, semi-independent structural components can be gradually accumulated. The Holy Grail of behavioral robotics is a general framework for describing such pieces, and for stitching them together.

In this paper, we use the term behavioral to imply that task knowledge is incorporated as an essential part of resolving the problem of control of redundant systems. Such knowledge enters into the process at the highest level, in the design of the virtual tools, at an intermediate level, in determining the selection of and transition between appropriate tool classes, and at the lowest level, in the instantiation of a tool. In the most general sense, imposing a set of dimensionality-reducing constraints on an active system can be viewed as instantiating a behavior. Another way of looking at the approach is as a method of using domain/task knowledge to partition the control effort between high and low level processes.

The justification for a virtual tool approach has an interesting implication concerning application of flexible systems, namely that flexible automation will be useful primarily in situations where special-purpose fixturing is unjustified. An example is one-off manufacturing, where only a few copies of a product will be made, and the cost of a special-purpose assembly line cannot be recovered by volume production. Another example is assembly in exotic environments such as space or underwater, where it is impractical to supply all the fixturing needed for hard automation because of size, time or weight constraints. A third example is in applications such as nuclear cleanup that involve a high degree of uncertainty about the exact nature of the problem, that can only be resolved on site. Similar observations have been made before.

3 Dextrous Manipulation as a Testbed

In the past few years dextrous robot hands, such as the 9DOF Salisbury hand [Salisbury, 1982], and the 16DOF Utah/MIT hand [Jacobsen, 1984; Jacobsen *et al.*, 1986], have become increasingly available in research environments. This has sparked a considerable amount of research on grasping and dextrous manipulation. From one point of view, these problems are specific to robot hands, however, they exhibit many of the basic problems of redundant control. If these problems could be resolved in a flexible fashion, then it is likely that some of the techniques would have more general application. In addition, the existence of a well developed physical device provides an acid test of control techniques. We have consequently taken the problem of controlling the Utah/MIT hand, particularly for fine manipulation and adaptive grasping of general objects, as a practical test of the virtual tool approach.

Dextrous robot hands with many degrees of freedom offer several advantages over conventional grippers. They provide a large degree of versatility for fine motions. They facilitate the partition of manipulation tasks into gross and fine manipulation, where a robot arm performs gross positioning and the hand handles fine manipulation. They allow a better control of gripping forces. Their large number of joints allows them to adapt well to a many different shapes, enabling them to grasp a wider variety of objects than conventional grippers. Also, their compliance and redundancy aid in overcoming uncertainties and errors

in the modeling. These same properties, however, make programming dextrous hands to perform useful tasks a challenging task. Manipulation with multiple cooperating fingers requires more sophisticated sensing, feedback, and planning than are found in conventional robot controllers [Salisbury *et al.*, 1989].

When an object is manipulated with a parallel jaw gripper, the inflexibility of manipulator generally forces the programmer to provide accurate models of the objects to be grasped or manipulated. Partly because of the predominance of rigid manipulators, most research on fine manipulation has relied on the assumption that accurate and complete models of the objects being manipulated are available or can be easily obtained. If, however, we want a manipulator to operate in unknown environments or manipulate a wide variety of objects, such model-based approaches are difficult to use. A dextrous manipulator, on the other hand, can theoretically get by with a much sparser model, since the redundancy and compliance of the hand can be used to achieve stable grasps in the presence of uncertainty. The main challenge is how to take advantage of this redundancy and compliance without being overwhelmed by the inherent additional complexity of these manipulators.

3.1 Previous Work Using Dense Models

A considerable amount of work has been done on the problems of grasping and fine manipulation. Much early work was theoretical, due to the lack of well implemented physical manipulators. Even now, with dextrous hands are available in many research laboratories, theoretical studies remain important. These studies make almost exclusive use of classical deterministic techniques of statics, kinematics, dynamics and control theory. It is generally assumed that complete models of the geometry of the hand and the grasped object are available.

Early work by Hanafusa and Asada [Hanafusa and Asada, 1982] showed that an elastic grasp was stable when the total energy stored in the elastic fingers was (locally) minimal with respect to small translations and rotations of the grasped object. Salisbury investigated grasping models with and without friction using several types of contact between a finger and an object [Salisbury, 1982]. Later, Trinkle, Abel and Paul [Trinkle *et al.*, 1988] studied the mathematics of frictionless grasping. They presented a method for grasping a polygonal object with a two-dimensional hand composed of a palm and two hinged fingers.

Manipulation was considered theoretically by Mishra [Mishra, 1989; Mishra, 1990], who explored the relationship between analytic approaches to dextrous manipulation and combinatorial geometry. Rus [Rus, 1993] presented a linear time algorithm for rotating a polygon in the frictionless limit by sliding one finger along the edge of the polygon, and Omata [Omata, 1994] presented an algorithm for two dimensional object reorientation using a sequence of regrasps.

Some approaches that deal with uncertainty in manipulation tasks have also been proposed. One of the first was published by Lozano-Pérez, Mason and Taylor in 1984 [Lozano-Pérez *et al.*, 1984]. The focus of this work was to realistically model physical processes that required compliant motion, assuming the worst-case bounds on uncertainties in sensing and effecting were known. Classical planning techniques were then used to find paths from initial to goal configurations. Assuming correctness of the model, the predictions can be

proven to be correct and complete [Mason, 1984]. Latombe gives a good overview of this method as well as some of its subsequent modifications [Latombe, 1989].

These theoretical works were, with the exception of [Hanafusa and Asada, 1982], not accompanied by experimental results to test the validity of the assumptions they made. It is therefore, hard to tell whether they have any practical applicability. In general, the requirement for complete and accurate models is a major impediment to practical implementation.

3.2 Previous Work Using Sparse Models

When dextrous manipulators became widely available for experiments, it was realized that the complete models used for theoretical analyses were difficult to obtain in the real world, and that the worst-case analysis of uncertain models was unnecessarily restrictive. It was also difficult to determine the limits of uncertainty. From the theoretical standpoint, it was shown that the general problem of motion planning with uncertainty is NEXP-Hard [Canny, 1987]. Thus, any algorithm for motion planning that guarantees an optimal solution must be hopelessly slow or assume a complete and correct model.

The compliance inherent in some dextrous manipulators allows the use of sparser models [Pook and Ballard, 1992]. By using such models, some of the problems associated with obtaining and utilizing complete models can be avoided. Such models have only recently become popular but their use is increasing, as they seem to work for the control of real devices.

A number of approaches are essentially based on heuristics derived from observation of human hands. Lyons [Lyons, 1985] defined heuristic measures of *firmness* and *precision* and matched them to the specification of a task. Bekey *et al.* [Bekey *et al.*, 1993] developed a knowledge-based grasp planner for the University of Belgrade/USC robotic hand. This system is an advance over previous ones in that it chooses grasp parameters as well as grasp types, and thus is not limited to a small set of pre-stored positions. Casselli *et al.* [Casselli *et al.*, 1993] presented a hybrid system for grasp synthesis that integrates symbolic and neural computations. The symbolic modules encode heuristic knowledge along with simple geometric reasoning, and the neural networks modules establish the more complex relationships between the geometric attributes of the object and the hand kinematics.

There have also been a number of behavior-based approaches inspired by Brooks' subsumption architecture [Brooks, 1986]. Murphy *et al.* [Murphy *et al.*, 1993] presented a simple behavior-based approach to grasping using a three-fingered manipulator. Matsui *et al.* [Matsui *et al.*, 1992] have used a subsumption architecture for the control of the Hyuma robotic hand. Coelho and Grupen [Coelho and Grupen, 1994] have studied the grasping problem viewing it as a control composition problem. They presented behavior-based force and moment controllers that achieve stable grasps on polygonal objects with a four-fingered hand.

Some work has been done in robotic manipulation using machine learning techniques. In general, these approaches have dealt with simple tasks such as grasping with parallel jaw grippers and simple manipulation strategies not involving complex manipulators. Dunn and Segen [Dunn and Segen, 1988] presented a robotic system that learns how to grasp

objects under visual control by trial and error. Working grasps are saved along with the object's shape. The system generalizes to different positions and orientations but not to sizes. Kamon *et al.* [Kamon *et al.*, 1994] presented a robotic system that learned to grasp objects with a parallel-jaw gripper using visual information. Their system learns to choose and to predict the quality of a given grasp. It incrementally improves its performance over the course of a training session. The system used very little information about the target object. In particular, no attempt was made to recover the object's shape.

Finally, there are a number of techniques that attempt to simplify the control problem by reducing the effective number of degrees of freedom through various techniques. Michelman and Allen [Michelman and Allen, 1993] used hybrid control to perform dextrous manipulation with the Utah/MIT hand. In their system, one or two fingers guide the movement of the object while the others keep the grasping force within a safety range. This decoupling of the degrees of freedom greatly reduces the dimensionality of the parameter space of the hand for the task.

Another approach, pioneered by Speeter [Speeter, 1991], uses *motor primitives*, which are sequences of joint position changes encoding small functional motions of the hand. Motion primitives can be added, subtracted, reversed and scaled. Although the combination rules are not mathematically correct, the compliance of the hand compensated for the inaccurate trajectories, and experimental results using the Utah/MIT hand showed this approach to work reasonably well. Nelson [Nelson] extended Speeter's approach by storing separate motor primitives for large and small objects and computing the appropriate primitives at run time by detecting the size of the object being manipulated and linearly interpolating the stored primitives.

4 Virtual Tools for the Utah/MIT Hand

In this section we describe the use application of the virtual tool approach to the problem of grasping and manipulation with the Utah/MIT hand. In particular, we describe the instantiation of two different virtual tools, namely, a gripper and a 6DOF manipulator. The gripper is a conceptually simple device having a single control parameter, while the manipulator is slightly more complex, with six control parameters. We demonstrate the concept of morphing in a task involving grasping and manipulation of novel objects. Specifically, we show how a robust perceptual signal can be used to signal the transition from the grasp tool to the manipulation tool, and how perceptual information derived from the state of the hand at the transition point serves to set the instantiation parameters of the manipulation tool.

4.1 The Grasp Tool

In order to manipulate objects with a dextrous hand, they must first be grasped with what is commonly referred to as a *precision grasp*, that is, one that allows manipulation with the fingers. In order to avoid overconstraining the fingers, such grasps typically involve only fingertip contacts. The virtual tool we use for grasping is based on the observation

that, particularly in the case where we have more than three fingers, and hence redundant stability, grasping can be viewed as a process of positioning and preshaping the hand, and then closing the fingers as a unit until a grasp stability criterion is satisfied. Local processes are associated with each finger to avoid unstable positions and to detect contact.

At a higher level of detail, the grasp tool is actually composed of two similar one-parameter *subtools* instantiated sequentially, and connected by a hidden morphing transition. To start, we select a preshape based on rough information about the object's shape and position the hand about the object. The fingers are then closed as a group, each towards an appropriate goal location under the guidance of a single control parameter until all fingertips touch the object. Local processes operating off tactile sensors freeze each finger as it contacts the object. A morphing transition then takes place, and a second tool is then instantiated that moves setpoints for the fingers inboard of the boundary, in a manner that achieves a secure and stable grasp.

The preshape information, together with the goal positions for each finger constitute the instantiation parameters of the first subtool. The positions at contact instantiate the second subtool. Because both tools are single parameter operations, and since the second can be automatically instantiated at the termination of the first, the morphing transition is effectively hidden and the entire grasp can be viewed as a single one-parameter operation from the point of view of the user process. If the user does not need to control the speed of the grasp, even this parameter can be hidden, and the entire procedure viewed as an atomic operation. This net effect the grouped movement is somewhat similar to Iberall's virtual finger concept, but in a more general setting.

In general, we can select from a set of preshape and goal closures, however, for the class of objects we have dealt with so far, a single preshape has sufficed. When used in this way, the above procedure is not guaranteed to produce the optimally stable grasp, by any particular criterion. However, because we have more fingers than are needed to produce a minimally stable grasp, it produces a good grasp in most situations involving objects having a minimum diameter greater than the diameter of a finger, and a maximum diameter smaller than the length of a finger. It is possible to contrive situations (e.g. slippery wedges) where stable grasps are difficult to find, and this sort of strategy fails; such situations, however, can be reasonably considered to be pathological. Moreover, it is possible to recognize such situations and take corrective action.

For hard situations (e.g. pathological objects, very small objects, or large, light objects with handles) if a description of the optimal grasp be available from some source, then it can be executed by appropriately setting the preshape (i.e. the grasp tool instantiation parameters). We have run experiments with a system that does this, and found some improvement. Obtaining the object model needed to compute the optimal grasp is a laborious procedure, and hence we have not incorporated this function into the current application. Our current inclination is to recognize whether an object belongs to some general class of difficult objects (e.g. long skinny shapes) and invoke a specialized preshape to deal with it.




Figure 1: The contact tetrahedron

4.2 A Manipulation Tool

Once an object has been grasped we want to manipulate it. For rigid objects, this entails re-orientation in a six-dimensional space. A general rigid manipulation tool would thus have six control parameters. Instantiation of further constraints would permit such a basic tool to be cast into more specialized tool, for instance one performing only translation or rotation about a particular axis. We refer to such specializations as *derived tools*.

The method presented here does not rely on *a priori* models. Instead, we show how a six parameter manipulation tool can be instantiated on the basis of sensory information available at the termination of the grasping process. The basic process involves a morphing transition from the grasp tool to the manipulation tool, that takes place on the completion of the grasp action. The instantiation parameters of the manipulation tool are derived from the commanded fingertip positions known as *setpoints* at the termination of the grasp. These values are directly available from the hand's sensors.

The basic idea is based on the observation that the setpoints of the fingers define a rigid object in three space, in this case a (possibly degenerate) tetrahedron. We refer to this object as the *grasp tetrahedron*. The contact points of the fingers on the object define another object that we refer to as the *contact tetrahedron*. Because of the compliant control system, we can model the statics of the situation by regarding each vertex of the contact tetrahedron as being attached to the corresponding vertex of the grasp tetrahedron by a virtual spring. In the case of an object free to move, net wrench is constrained to be zero. Moreover, for a well conditioned grasp, the wrench zero coincides with a deep local minimum of the spring energy function on manifold defined by the rigid displacements of the grasp tetrahedron with respect to the contact tetrahedron. This means that if we treat the fingertip contacts as fixed points on the object, then the object can be rotated and translated by executing the desired rigid transformation on the grasp tetrahedron. In other words, displace the grasp tetrahedron and the object will follow. Although, due to non-zero finger size and other affects, the assumption of fixed contact points is not strictly correct, our experimental results show that the compliance of the Utah/MIT hand compensates for the errors introduced by the use of this assumption. Somewhat similar approaches have been suggested in [Okada, 1982] and [Salisbury *et al.*, 1989].

In more detail, the formulation is as follows. Let ${}^H p_0, {}^H p_1, {}^H p_2$ and ${}^H p_3$ be the fingertip to object contact points, where ${}^H p_i$ denotes that p_i is being measured with respect to a




Figure 2: The grasp tetrahedron




Figure 3: The object-centered frame of reference C and the hand-centered reference frame H . The y axis is perpendicular to the page and pointing outward.

hand-centered reference frame H . Let ${}^Hq_0, {}^Hq_1, {}^Hq_2$ and Hq_3 be the commanded fingertip positions. The tetrahedron ${}^Hp_0, {}^Hp_1, {}^Hp_2, {}^Hp_3$ is the contact tetrahedron. ${}^Hq_0, {}^Hq_1, {}^Hq_2, {}^Hq_3$ is the grasp tetrahedron. The contact and grasp tetrahedra for a sample grasp are shown in figures 1 and 2. If the object is stably grasped, the grasp tetrahedron will generally be completely contained in the contact tetrahedron. The locations of the grasp and the contact tetrahedra can be obtained respectively from the measurements of the commanded and actual joint angles and the forward kinematics of the hand.

As noted above, the force applied to the object at any of the contact points is approximately proportional to the difference between the fingertip's commanded and actual positions. If the fingertip positions with respect to the object remain constant, we can apply arbitrary rotations and translations to the object while keeping a constant force applied in every contact point by just rotating and translating the grasp tetrahedron. Formally, we describe the situation as follows.

Let ${}^Hc = [c_x, c_y, c_z]^T$ be the centroid of the grasp tetrahedron with respect to frame H . We define a coordinate system C parallel to the hand-based reference frame H . This is illustrated in figures 3 and 4. Following Craig [Craig, 1989], we will denote the transformation describing a coordinate frame A with respect to another coordinate frame B by ${}^B_A T$.




Figure 4: The object-centered frame of reference C . The z axis is perpendicular to the page and pointing inward.

Using this notation,

$${}^H_C T = \begin{bmatrix} 1 & 0 & 0 & c_x \\ 0 & 1 & 0 & c_y \\ 0 & 0 & 1 & c_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

We then define an object-centered frame of reference D that will describe the position and orientation of the object with respect to its original position and orientation. If the object is not moved, then $C = D$ and ${}^D_C T = I$, where I is the identity matrix. The coordinates of a grasp point ${}^D p_i$ with respect to the object-centered frame are given by

$${}^D q_i = {}^D_C T {}^C_H T {}^H q_i$$

If no movement has taken place, ${}^D_C T = I$, and

$${}^D q_i = {}^C_H T {}^H q_i = {}^H_C T^{-1} {}^H q_i$$

We can now describe arbitrary translations and rotations of the grasp tetrahedron ${}^H q_0 {}^H q_1 {}^H q_2 {}^H q_3$ by letting ${}^D_C T$ be an arbitrary rigid transformation matrix instead of the identity matrix.

5 Experimental Results

We performed three sets of experiments on the application of our method using the Utah/MIT dextrous hand. In the first set, we tested a straight-forward implementation of the manipulation tool described in the previous section, observing its open-loop performance and measuring the position errors in the execution of example manipulation sequences. In the second set of experiments, we added sensory input in the form of visual feedback. Using point tracking, we implemented a closed-loop controller to obtain more accurate positioning of the object. Again we measured the performance. In the third set of experiments we used the open-loop strategy of the first experiment for dextrous telemanipulation. With this system, a human operator can control the position and orientation of an object using a Polhemus position/orientation sensing device.




Figure 5: Translation along the x axis




Figure 6: Translation along the y axis




Figure 7: Translation along the z axis




Figure 8: Rotation around the x axis




Figure 9: Rotation around the y axis




Figure 10: Rotation around the z axis




Figure 11: Position errors for a translation along the x axis. Positions and errors are given in inches




Figure 12: Visual Feedback System

5.1 Open-Loop Strategy

The first set of experiments consisted of a straightforward implementation of the manipulation tool. Figures 5 to 10 show the Utah/MIT manipulating a wooden block. The figures show translations along and rotations around the three main axes. Although we only show the results of translations and rotations orthogonal to the main axes, translations and rotations can be performed with respect to any arbitrary axis.

While the manipulations were performed, we took images with a camera with known calibration to measure the errors in positioning. Figure 11 shows the errors in position measured while moving the block along the x axis. It can be seen that the accuracy is very good for displacements ranging from about -2 inches to 1.4 inches with respect to the original object position. Outside of that range, the accuracy decreases quickly because the commanded position starts to be outside of the range of the fingers. As the table shows, our method is remarkably accurate considering its simplicity. A video tape of a sequence of translations and rotations of several objects is available from the authors upon request

5.2 Closed-Loop Strategy Using Visual Feedback




Figure 13: Position errors using visual feedback for a translation along the x axis. Positions and errors are given in inches




Figure 14: Schematic Diagram of the Dextrous Telemanipulation System

Although the accuracy of the open-loop strategy is remarkably good, we tried using sensor information to improve the performance of the system. We used the same camera setting as in the previous experiment, but this time we used the information obtained from the camera as the input signal for a controller. From the image we determined the error in position and used a simple proportional control strategy to try to reduce the errors, as shown in the schematic diagram in figure 12. Figure 13 shows the improvement in performance obtained by adding the closed-loop component. It can be seen from the figure that the errors in position were practically eliminated. When using the control strategy, the precision of the positioning is limited only by the resolution of the image. For every measurement shown in the figure, the error corresponds to less than half of a pixel in image coordinates, so the error is minimal for the image resolution used.

5.3 Dextrous Telemanipulation

Programming a dextrous manipulator such as the Utah/MIT hand to perform fine manipulation is a very difficult task, mainly because of the high dimensionality of its raw control parameter space. If we want to use a dextrous hand for telemanipulation we need to obtain




Figure 15: Telemanipulated insertion

sixteen different parameters from the teleoperator, and each individual parameter is a potential source of error. Another problem is that the Utah/MIT hand is only approximately anthropomorphic, thus there is no straightforward way to reliably map the joint angles of the teleoperator's hand to those of the robotic hand.

An alternative approach is to view the hand as a virtual 6DOF manipulator. In this case we need only specify the desired position and orientation of the object being manipulated. This can be done with a device such as the Polhemus position and orientation sensor, as shown in figure 14. Figure 15 shows a telemanipulated insertion operation using this approach.

From the initial position and orientation of the Polhemus sensor we obtain a base reference frame that is put into correspondence with a frame with origin at the centroid of the grasp and aligned with the hand's base coordinate system. Movements are then commanded by varying the position and orientation of the Polhemus sensor with respect to this reference frame. This corresponds to reading from the sensor the rigid transformation matrix ${}^D_C T$, as explained in Section 4.

6 Conclusions

This work has introduced the concepts of instantiation and morphing of virtual tools as a mathematically powerful and conceptually intuitive way to utilize redundant degrees of freedom in a manipulator. We have demonstrated the efficacy of the method with experiments using the Utah/MIT hand. We have also shown how the precision of the manipulations

can be increased by adding visual feedback. Finally, we have shown that the application of this method to fine telemanipulation allows using a task specification language that is higher-level and more reliable than joint angle specifications.

References

- [Arbib *et al.*, 1985] M. Arbib, T. Iberall, and D. Lyons, "Coordinated Control Programs for Movements of the Hand," Technical report, Department of Computer and Information Science, University of Massachusetts at Amherst, Amherst, Massachusetts, 1985.
- [Bekey *et al.*, 1993] G. A. Bekey, H. L. Liu, R. Tomovic, and W. J. Karplus, "Knowledge-Based Control of Grasping in Robot Hands Using Heuristics from Human Motor Skills," *IEEE Transactions on Robotics and Automation*, 9(6):709-722, December 1993.
- [Brooks, 1986] Rodney A. Brooks, "A Robust Layered Control System for a Mobile Robot," *IEEE Journal of Robotics and Automation*, pages 14-22, April 1986.
- [Brooks, 1991] Rodney A. Brooks, "Intelligence Without Reason," Technical Report 1293, MIT Artificial Intelligence Lab, Cambridge, Massachusetts, April 1991.
- [Canny, 1987] John F. Canny, *The Complexity of Robot Motion Planning*, The MIT Press, Cambridge, Massachusetts, 1987. ACM Doctoral Dissertation Award.
- [Caselli *et al.*, 1993] Stefano Caselli, Eugenio Faldella, Bruno Fringuelli, and Francesco Zanichelli, "A Hybrid System for Knowledge-Based Synthesis of Robot Grasps," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1575-1581, Yokohama, Japan, 1993.
- [Chirikjian and Burdick, 1994] Gregory S. Chirikjian and Joel W. Burdick, "A Modal Approach to Hyper-Redundant Manipulator Kinematics," *IEEE Transactions on Robotics and Automation*, 10(3):343-354, June 1994.
- [Cleary, 1990] K. Cleary, "Incorporating Multiple Criteria in the Operation of Redundant Manipulators," in *Proceedings of the 1990 IEEE International Conference on Robotics and Automation*, pages 618-624, Cincinnati, Ohio, 1990.
- [Coelho and Grupen, 1994] Jefferson A. Coelho and Roderic A. Grupen, "Effective Multi-fingered Grasp Synthesis," Technical Report 94-12, Laboratory for Perceptual Robotics, Department of Computer and Information Science, University of Massachusetts at Amherst, Amherst, Massachusetts, 1994.
- [Craig, 1989] John J. Craig, *Introduction to Robotics: Mechanics and Control*, Addison-Wesley Publishing Company, Reading, Massachusetts, 1989. 2nd Edition.
- [Davidor, 1991] Yuval Davidor, *Genetic Algorithms and Robotics - A Heuristic Strategy for Optimization*, World Scientific Publishing Co., Singapore, 1991.
- [Dunn and Segen, 1988] G. B. Dunn and J. Segen, "Automatic Discovery of Robotic Grasping Configurations," in *Proceedings of the 1988 IEEE International Conference on Robotics and Automation*, pages 396-401, 1988.
- [Hanafusa and Asada, 1982] H. Hanafusa and H. Asada, "Stable Prehension by a Robot Hand with Elastic Fingers," in M. Brady, editor, *Robot Motion*, pages 323-335. MIT Press, Cambridge, MA, 1982.

- [Hirose, 1993] Shigeo Hirose, *Biologically Inspired Robots*, Oxford University Press, New York, 1993.
- [Hollerbach and Suh, 1985] J. M. Hollerbach and K. C. Suh, "Redundancy Resolution of Manipulators Through Torque Optimization," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, pages 1016-1022, Nice, France, 1985.
- [Iberall, 1987] T. Iberall, "The nature of human prehension: Three dextrous hands in one," in *Proceedings of the 1987 IEEE International Conference on Robotics and Automation*, pages 396-401, Raleigh, North Carolina, 1987.
- [Iberall et al., 1988] Thea Iberall, Joe Jackson, Liz Labbe, and Ralph Zampano, "Knowledge-Based Human Prehension," in *Proceedings of the 1988 IEEE International Conference on Robotics and Automation*, pages 82-87, 1988.
- [Jacobsen et al., 1986] Steve Jacobsen, E. Iversen, D. Knutti, R. Johnson, and K. Bigger, "Design of the Utah/MIT Dextrous Hand," in *Proceedings of the 1986 IEEE International Conference on Robotics and Automation*, pages 96-102, 1986.
- [Jacobsen, 1984] Steve C. Jacobsen, "The Utah/MIT hand: Work in Progress," *IEEE Journal of Robotics and Automation*, 3(4), 1984.
- [Kamon et al., 1994] Ishay Kamon, Tamar Flash, and Shimon Edelman, "Learning to Grasp Using Visual Information," Technical report, The Weizmann Institute of Science, Rehovot, Israel, March 1994.
- [Latombe, 1989] Jean-Claude Latombe, "Motion Planning with Uncertainty: On the Preimage Backchaining Approach," in O. Khatib, J. J. Craig, and T. Lozano-Pérez, editors, *Robotics Review I*. The MIT Press, Cambridge, MA, 1989.
- [Liegeois, 1977] Alain Liegeois, "Automatic Supervisory Control of the Configuration and Behavior of Multibody Mechanisms," *IEEE Transactions on Systems, Man and Cybernetics*, 7(12):868-871, December 1977.
- [Lozano-Pérez et al., 1984] Tomás Lozano-Pérez, Matthew T. Mason, and R. H. Taylor, "Automatic Synthesis of Fine-Motion Strategies for Robots," *IEEE Journal of Robotics Research*, 3(1), 1984.
- [Lyons, 1985] Damian Lyons, "A Simple Set of Grasps for a Dextrous Hand," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, pages 544-548, St. Louis, Missouri, 1985.
- [Maciejewski, 1990] A. A. Maciejewski, "Fault Tolerant Properties of Kinematically Redundant manipulators," in *Proceedings of the 1990 IEEE International Conference on Robotics and Automation*, pages 638-642, Cincinnati, Ohio, 1990.
- [Maciejewski and Klein, 1985] A. A. Maciejewski and C. A. Klein, "Obstacle Avoidance for Kinematically Redundant Manipulators in Dynamically Varying Environments," *The International Journal of Robotics Research*, 4(3):109-117, 1985.

- [Maes, 1992] Patty Maes, "Behavior-based Artificial Intelligence," in J. A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 2: Proceedings of the second International Conference of the simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, December 1992.
- [Mason, 1984] Matthew T. Mason, "Automatic Planning of Fine Motion," in *Proceedings of the 1984 IEEE International Conference on Robotics and Automation*, pages 492–503, 1984.
- [Matsui *et al.*, 1992] Toshimoro Matsui, Tooru Omata, and Yasuo Kuniyoshi, "Multi-Agent Architecture for Controlling a Multi-Fingered Robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 182–186, Raleigh, North Carolina, 1992.
- [McGhee *et al.*, 1994] Scott McGhee, Tan F. Chan, Rajiv V. Dubey, and Reid L. Kress, "Probability-Based Weighting of Performance Criteria for a Redundant Manipulator," in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 1887–1894, San Diego, California, 1994.
- [Michelman and Allen, 1993] Paul Michelman and P. Allen, "Compliant Manipulation with a Dexterous Robot Hand," in *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, pages 711–716, 1993.
- [Mishra, 1989] Bhubaneswar Mishra, "Dexterous Manipulation: A Geometric Approach," in S. Stifter and J. Lenarčič, editors, *Advances in Robot Kinematics*. Springer-Verlag, 1989.
- [Mishra, 1990] Bhubaneswar Mishra, "Dexterous Manipulation: A Geometric Approach," Technical Report 524, New York University, Courant Institute of Mathematical Sciences, New York, New York, September 1990.
- [Murphy *et al.*, 1993] T. G. Murphy, D. Lyons, and A. J. Hendricks, "Stable Grasping with a Multi-Fingered Robot Hand: A Behavior-Based Approach," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 867–874, Yokohama, Japan, 1993.
- [Nakamura *et al.*, 1987] Yoshihiko Nakamura, Hideo Hanafusa, and Tsuneo Yoshikawa, "Task-Priority Based redundancy Control of Robot Manipulators," *The International Journal of Robotics Research*, 6(2):3–15, 1987.
- [Nelson] Randal C. Nelson. Personal Communication.
- [Okada, 1982] Tokuji Okada, "Computer Control of Multijointed Finger System for Precise Object-Handling," *IEEE Transactions on Systems, Man, and Cybernetics*, 12(3), 1982.
- [Omata, 1994] Toru Omata, "Planning Reorientation of an Object with a Multifingered hand," in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, San Diego, California, 1994. to appear.

- [Pamanes and Zeghloul, 1991] J. A. Pamanes and S. Zeghloul, "Optimal Placement of Robotic Manipulators Using Multiple Kinematic Criteria," in *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*, pages 933-938, Sacramento, California, 1991.
- [Pook and Ballard, 1992] Polly K. Pook and Dana H. Ballard, "Contextually Dependent Control Strategies for Manipulation," Technical Report 416, University of Rochester, Computer Science Department, April 1992.
- [Raibert, 1986] Marc H. Raibert, *Legged Robots That Balance*, The MIT Press, Cambridge, Massachusetts, 1986.
- [Rus, 1993] Daniela Rus, "Coordinated Manipulation of Polygonal Objects," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 106-112, Yokohama, Japan, 1993.
- [Salisbury, 1982] K. Salisbury, *Kinematic and Force Analysis of Articulated Hands*, PhD thesis, Stanford University, Department of Mechanical Engineering, Stanford, California, 1982.
- [Salisbury et al., 1989] Kenneth Salisbury, David Brock, and Patrick O'Donnell, "Using an Articulated Hand to Manipulate Objects," in Michael Brady, editor, *Robotics Science*. MIT Press, 1989.
- [Speeter, 1991] Thomas H. Speeter, "Primitive Based Control of the Utah/MIT Dextrous Hand," in *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 866-877, Sacramento, California, 1991.
- [Stansfield, 1991] S. Stansfield, "Robotic Grasping of Unknown Objects: A Knowledge-based Approach," *International Journal of Robotics Research*, 1991.
- [Sutherland and Ullner, 1984] Ivan E. Sutherland and Michael K. Ullner, "Footprints in the Asphalt," *The International Journal of Robotics Research*, 3(2):29-36, 1984.
- [Trinkle et al., 1988] Jeffrey C. Trinkle, Jacob M. Abel, and Richard P. Paul, "An Investigation of Frictionless Enveloping Grasping in the Plane," *The International Journal of Robotics Research*, 7(3):33-51, June 1988.
- [Yoshikawa, 1990] Tsuneo Yoshikawa, *Foundations of robotics : analysis and control*, The MIT Press, Cambridge, Massachusetts, 1990.