A Control Paradigm for General Purpose Manipulation Systems

Roderic A. Grupen and Thomas C. Henderson

Department of Computer Science
University of Utah
Salt Lake City, Utah 84112

Abstract

Mechanical end effectors capable of dextrous manipulation are now a reality. Articulated, multi-fingered hands can perform a wide variety of tasks: sensing the environment, and behaving alternately as powerful and delicate manipulation devices. This compliant nature of multi-fingered hands allows constraints derived from the object and the task to take precedence when selecting a manipulation strategy. The goal of this work is the creation of a framework within which independent objectives of a manipulation process may be used to direct a manipulation strategy. The set of contacts that are applied to a task can be partitioned into subsets with independent objectives. A system of this sort is flexible enough to manage large numbers of contacts and to address manipulation tasks which require the removal and replacement of fingers. A simulator has been constructed and results of its application to position synthesis for initial grasps is presented. A discussion of the manipulation testbed under construction at the University of Utah employing the Utah/MIT Dextrous Hand is presented.

1 Introduction

Manipulators have been developed with a variety of configurations and with correspondingly diverse capabilities[6,13,15,16]. The method decribed in this paper will be applied to the Utah/MIT Dextrous Hand[7], but may be applied to any manipulator. The premise of our work is that the value of the human hand is largely its ability to perform as a *general purpose manipulation* device. A general purpose manipulator, in conjunction with other sensory modalities, can be used to: measure the stiffness of the environment[16], support object recognition and localization procedures[1], and to perform alternately as a delicate and powerful manipulator. The last of these characteristics requires that hand adapt to a wide variety of tasks and objects. The work presented in this paper is focused on the development of a control structure for position synthesis which is flexible enough to support general purpose manipulation.

Researchers have enumerated grasps employed by humans with the hope of discovering the intrinsic qualities of successful hand-object interactions[3]. Such a manipulation *syntax* could be used to define an operational paradigm for each object and task. Arbib *et al.* proposed the concept of *virtual fingers*[2] as an effective means of reducing the complexity of the human hand and matching the hand's capabilities with the object's surface in light of the task. Tomavic *et al.* are developing the Belgrade hand and the concept of reflex control[17] in order to reduce the dimensionality of the control problem.

The approach employed in this paper is the enumeration of all the mechanical impedances imposed on the object by virtue of its contact with each finger. To completely constrain an object, the union over the contact set must span the object's six degrees of freedom. Kobayashi posed this problem in terms of hybrid position and force controllers[11]. Salisbury expresses the constraint criteria in terms of screw systems[16]. The Grip Transform relates command wrench intensities to the net forces applied to the object. The solution for the internal force magnitudes can be selected to optimize the grasp robustness[9], or to minimize the magnitude of the internal forces[5].

2 General Purpose Manipulation

We propose a functional integration of the hand and the object which recognizes the independent nature of the forces that the object's surface may transmit, and the manipulator's ability to generate those forces. The approach described here enforces functional priorities among independent prerogatives: the contact surface elements must be capable of transmitting forces that span the space occupied by the task, the hand must be properly configured to generate those forces.

2.1 Control Hierarchy

A perspective on the ideas that support the control structure presented in this section is presented by Minsky in his discussions of the Society of Mind[14]. The Mind is viewed as a society of separate, independent agents, each with its private agenda, contending for the resources available.

The flexibilities of the society structure can be used to implement general purpose manipulation. The *Manipulation Society* is illustrated in Figure 1. The top level of this society focuses

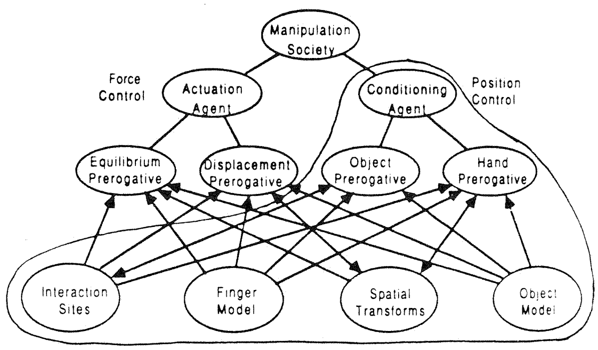

Figure 1: The Manipulation Society - the agency supporting position synthesis is indicated

CH2555-1/88/0000/0710$01.00 © 1988 IEE

attention on individual agents. The *Actuation* agent is indulged until the system violates the task defined conditioning envelope, at which time, the *Conditioning* agent is recognized.

The prerogative level of this control structure expresses various independent interests of the manipulation society. The compound objectives of manipulation are distilled into independent tasks which are prioritized and managed by the agents above them. The bottom level of this structure is a state model for the manipulation system. It represents the geometry of the hand-object system. The arrows in Figure 1 designate which agent uses the data at the prerogative level, and define the agent responsible for changing the dynamic state of the system at the bottom level. We will restrict our discussion to the synthesis of a geometry for the hand-object interaction, the components of which are indicated in of Figure 1.

2.2 Position Synthesis

The *Conditioning* agent is responsible for developing an interaction geometry for the hand-object system. To do this, it examines two agents: the *Object Prerogative* agent and the *Hand Prerogative* agent. The Object Prerogative agent seeks reachable positions on the surface of the object through which forces may be transmitted which collectively span the task. The Hand Prerogative agent drives the hand frame to positions and orientations which best address the generation of task define forces and velocities. For general purpose manipulation, it is the policy of the Conditioning agent to comply with the Object's Prerogative whenever possible, and to compromise the hand's conditioning if necessary to accomplish the task. Crosstalk between the two takes place through the changing state of the system. The following text describes: the models which support these processes, a means of expressing the task, the computation of the state error with respect to this task, and the expression of the object and hand prerogatives to reduce this error.

2.2.1 The Object Model

The object model uses a point contact with friction model to determine the family of wrenches that may be transmitted through a surface element. Figure 2 (a) depicts a set of contact forces that may transmitted to the object using this friction model. The forces, F_1 through F_5, are not independent, however, since the tangentially applied forces are dependent on the normally applied force, F_1. The object model represents a surface element using a unit normal force and scaled (by the static coefficient of friction) tangential forces.

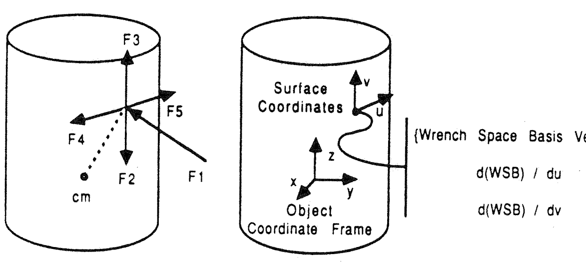

Figure 2: The Object Model Representing a Point Contact with Friction

The proximity of the contact to the object's center of mass allows the force system to be replaced by a set of wrench space basis vectors. By so doing, we have combined local surface properties, and contact friction properties into a six dimensional subspace representing the degree of constraint due to this contact from the perspective of the object. The object model used to support manipulation consists of an orthonormal basis (with associated magnitudes) for this wrench space and the net change in the contact wrench with respect to changes in the surface coordinates, u and v as illustrated in Figure 2 (b).

2.2.2 The Kinematic Models

Several metrics have been suggested to describe the conditioning of a redundant manipulator[10,12,16]. We characterize the manipulator's ability to control forces and velocities in terms of a modified form of the *manipulability ellipsoid*[18].

The manipulability ellipsoid is defined by examining the singular value decomposition of the manipulator Jacobian[4].

$$\begin{gathered}\text{If } J \in R^{M\times N},\\ \text{Then } \exists\ U \in R^{M\times M} \text{ and } V \in R^{N\times N},\\ \text{Such that, } J = U\ \Sigma\ V^T, \text{ where:}\end{gathered}$$

$$\Sigma = \left[\begin{array}{cccccc|c} \sigma_1 & & & & & & \\ & \sigma_2 & & & & & \\ & & \sigma_3 & & & & \\ & & & \cdot & & & 0 \\ & & & & \cdot & & \\ & & & & & \cdot & \\ & & & & & \sigma_m & \end{array}\right] \in R^{M\times N},\ with,$$

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_m \geq 0.$$

These σ_i represent the m singular values of J. The *Principally Conditioned Axes* (PCA's) are defined:

$$[\sigma_1 u_1,\ \sigma_2 u_2,\ \ldots,\ \sigma_m u_m]$$

where the $u_i \in R^m$ are the m column vectors of the U matrix above, weighted by the corresponding singular values. The PCA's describe the kinematic effectiveness of the transformation from finger joint space to the fingertip Cartesian space.

While the manipulator geometry may be capable of generating instantaneous velocities, it may not be able to do so in the region immediately surrounding its current state. The PCA's are weighted to reflect the physical joint limits of the manipulator. This is accomplished by defining a weighting coefficient which penalizes extreme joint positions. For the four degree of freedom finger:

$$W(\vec{\theta}) = \prod_{i=0}^{3} \cos\left\{2\pi\left(\frac{\theta - \theta_{nom}}{\theta_{max} - \theta_{min}}\right)_i^2\right\}$$

The results of applying such a weighting coefficient to varying configurations of the Utah/MIT finger is presented in Figure 3. A finger is viewed from the side and from behind to illustrate the weighted PCA's for a variety of finger geometries.

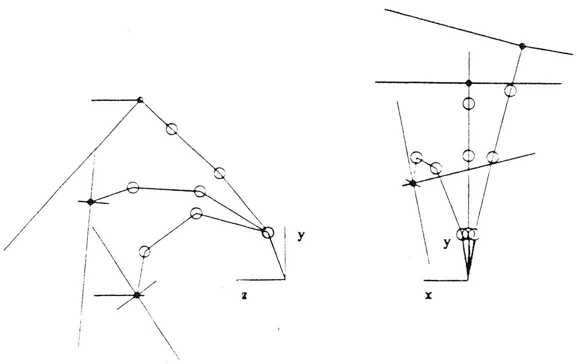

Figure 3: Weighted Principally Conditioned Axes for the Utah/MIT Finger

The weighted PCA's represent the directional character of the kinematic conditioning of the hand. The proximity of this contact force system to the object's center of rotation defines the transformation into a six dimensional wrench subspace. The union over all contacts of these wrench subspaces must span the task.

This characterization of the manipulator is used directly to apply the manipulator to the task. Any kinematic subchain of the finger can be used to define a family of PCA's which are applied to the task through the contact position.

We define an inverse kinematic model containing joint space solutions which maximize the product of the weighting coefficient and the singular values for all fingertip positions within the workspace. Physically, the metric is proportional to the volume spanned by the Principally Conditioned Axes. This is similar to the manipulability index suggested by Yoshikawa[18], but now expresses the physical joint limits of each finger. The finger's workspace is discretized into $0.25in$ bins. The joint configuration which maximizes this conditioning metric, and the maximal value of the metric are recorded in the model. This model allows the planner to efficiently assess the ability of the hand to comply to contact position trajectories.

2.2.3 The Task and the State Error

Screw systems are a natural means of expressing the six degrees of freedom of an object[16]. We present the stability and external force components of a task as the sum of two wrench (generalized force) commands.

$$\overline{T} = [\Omega_{R1}\ \Omega_{R2}\ \Omega_{R3}\ \Omega_{R4}\ \Omega_{R5}\ \Omega_{R6}] + [\Omega_{T1}\ \Omega_{T2}\ \Omega_{T3}\ \Omega_{T4}\ \Omega_{T5}\ \Omega_{T6}]$$

where:

Ω_{Ri} = a magnitude goal for forces in the i^{th} degree of freedom,

Ω_{Ti} = terms used to anticipate the task, inertial or external loads.

The system is not permitted to be stable only with respect to infinitesimal perturbations, but is compelled to seek contact topologies which are predictably robust with respect to expected perturbations.

A system is quantified by computing the error of the current state with respect to the task. If we represent the wrench space spanned by a contact system using a set of n orthonormal six dimensional basis vectors, then we may express the error of this system relative to the task as follows:

$$\vec{E} = \vec{T} - \sum_{1}^{n} mag_i \times \hat{b}_i$$

where:

$\vec{E}$ = Contact system error relative to the task

$\vec{T}$ = task vector

$\hat{b}_i$ = an orthonormal basis vector for the contact wrench space

mag_i = min(the contact system capability along $\hat{b}_i$, the projection of $\vec{T}$ onto $\hat{b}_i$)

The equation above removes the components of the task which project onto the contact wrench space and are within the magnitude limitations of the contact system. The residual vector is, therefore, the deficiency of the current contact system relative to the task.

2.2.4 The Object Prerogative

Given a set of positions on the object's surface through which interactions take place, it is possible to produce an incremental improvement in the system state by defining migrations of these sites over the object's surface. It is the object's prerogative to select features of the object and interaction topologies particularly well suited to the task. Unions over sets of contact wrench systems, like those described in Section 2.2.1, span a subspace of the six degrees of freedom of the object. We require as our criterion of stability, that a basis for the union of the wrench subspaces over all contacts must be of rank 6.

We may compute a state error for the contact wrench basis as was described in Section 2.2.3. In order to reduce this object state error, the planner interrogates the object model to determine the value of the derivative wrench systems with respect to orthogonal surface migrations. If the error is expressed as a linear combination of the derivative wrenches, a migration of the interaction sites in surface coordinates can be defined which improve the state of the system.

A demonstration is presented in Figure 4 for a system a cylinder 4 *inches* in diameter and two contacts. Here, a migrating interaction site is directed to improve the net wrench space defined by it and a fixed interaction site. The task is a uniform wrench. The migrating interaction site seeks a position on the object's

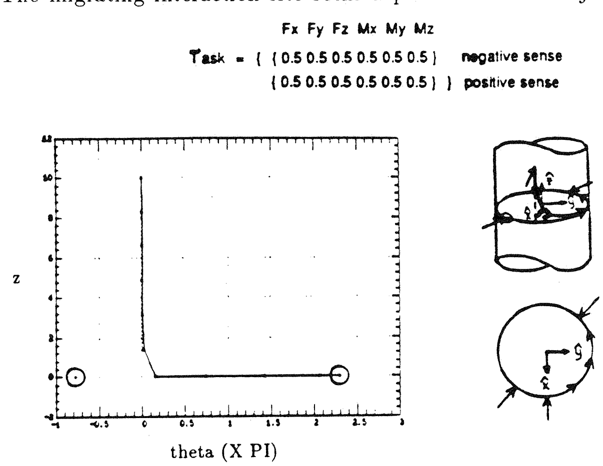

Figure 4: The Directed Migration of a Two Interaction Site Contact System

surface where a wrench subspace complementary to that of the fixed interaction site can be generated. It is in essence a two dimensional gradient following across a six dimensional manifold. The task is never completely satisfied by this contact system, but the state error with respect to the task is minimized.

When coupled with an agent representing the manipulator, the object prerogative is constrained by the ability of the hand to accommodate the surface trajectory. Therefore, only that portion of the trajectory that is reachable by the hand in its current configuration is used. The contact position which most effectively addresses the system deficiencies can be identified and a trajectory of this site toward the stabiltity robustness goal can be computed.

2.2.5 The Hand Prerogative

We remarked earlier, that the hand is required to comply to the task and to the geometry of the object. To accommodate the trajectories over the surface of the object defined by the Object Prerogative, the hand coordinate frame must move to positions for which all fingers are isotropically conditioned. The responsibility of the hand to comply is supported by the workspace model for the fingers. This inverse kinematic model contains the weighted manipulability index at discretized positions within the workspace. This index defines a smooth continuous scalar field with a single maximum which can be used to define a gradient space toward well conditioned configurations. It is critical to the conditioning process that the scalar field has 'hese properties, they constrain the range of candidate weighting functions we may employ to express the joint limits of the manipulator. During the phase of the conditioning process when the Object Prerogative is not yet satisfied, this manipulability gradient is used to direct the hand into configurations which are uniformly suited to arbitrary displacement tasks.

Following the convergence of the Object Prerogative, the hand must comply to the requirements of the task. The Principally Conditioned Axes, weighted to reflect joint limits, were developed in the Finger Model section to reflect the conditioning of the finger anywhere within its workspace. When several fingers contribute to a contact system, we may use the PCA's to compute the wrench space capabilities of this configuration relative to the object's center of rotation. In this way, the wrench subspace spanned by the contact system may be used to define an error with respect to the task as described in Section 2.2.3. This error is resolved by computing a migration of the hand frame which optimally applies these principally conditioned axes to the object in light of the task.

To accomplish this objective, the transform representing the position and orientation of the hand frame relative to the object is changed by a virtual displacement. The error resulting from this hypothetical state is compared with that resulting from other virtual displacements and the current state. The trajectory that reduces the error by the greatest amount is selected. The process continues until no further improvement in the hand state is possible with the current set of contact sites.

2.3 Experiments with Position Synthesis

Figure 5 illustrates the result of submitting a task to the Conditioning Agent. The figure illustrates the top view of a system consisting of a cylindrical object four inches in diameter and a two fingered Utah/MIT hand. The actual fingertip position is

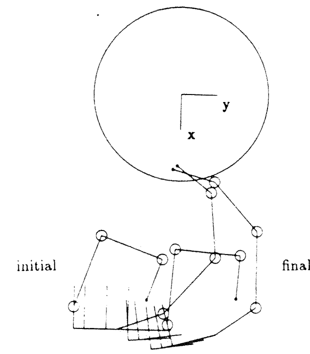

Figure 5: The Conditioning of One Finger for an Isotropic Wrench

displaced slightly from the correct surface position. This is due to the coarseness of the inverse kinematic solution. In practice, the planner must use a virtual object which is slightly larger than the real object to ensure that no contact actually takes place while the grasp is evolving. The Conditioning agent integrates the Hand's Prerogative and the Object's Prerogative. The task illustrated is to generate an isotropic wrench space, that is a uniform wrench in both plus and minus directions for all six degrees of freedom, using only the index finger of the hand. The state of the system cannot be improved from the object's perspective, but the hand frame migrates to a position relative to the object for which the index finger is more isotropically conditioned.

Figure 6 presents the results of modifying the task presented to the Conditioning Agent. In addition to the isotropic wrench space, a preferred force in the negative $\hat{y}$ direction is requested. The Object Prerogative is not immediately satisfied in this case, it directs the migration illustrated in the figure. The hand frame complies to the Object's Prerogative producing a markedly different behavior in the system. Prior to the time that the Object Prerogative is satisfied, the Hand Prerogative seeks to improve the conditioning of the fingers isotropically. This directs the hand to seek well conditioned states, from which it will be better prepared to accommodate subsequent state changes. A wrench/twist space task is submitted to the Hand Prerogative agent only after the Object Prerogative has converged, therefore, the hand frame retreats slightly from its most advanced position.

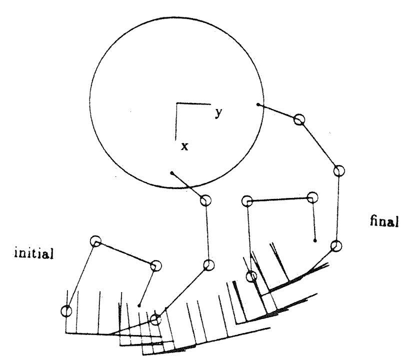

Figure 6: The Conditioning of One Finger for an Isotropic Wrench + F_{-y}

Finally, we have applied the Conditioning Agent to the geometry of the Utah/MIT hand. The resulting manipulator can reproduce the results presented earlier by defining a contact system consisting of the thumb and the index finger. Moreover, we may request three and four fingertip contacts and use arbitrary combinations of fingers. A simple, four fingered grasp of the cylinder is presented in Figure 7. Figure 7(a) presents the top view and side view of the initial hand-object configuration. Figure 7(b) illustrates the intermediate hand coordinate frame positions and the final configuration of the hand. These results demonstrate that the system can be applied to multiple finger contacts. It required 5.6 seconds of CPU time on the Vax 750 to reach the final state. Managing four independent contacts is pushing the real time capabilities of the system. It is conceivable that the concept of virtual fingers [2], identified by experience with a class of objects and tasks, may help alleviate the computational burden on the planner while still allowing the system to respond to new or unexpected circumstances.

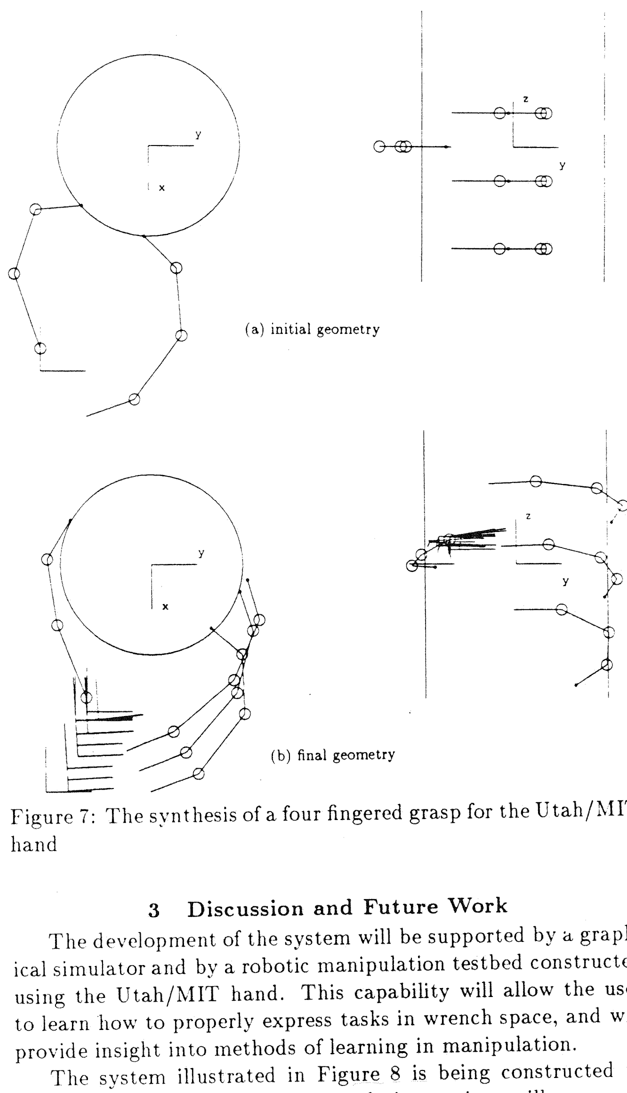

Figure 7: The synthesis of a four fingered grasp for the Utah/MIT hand

3 Discussion and Future Work

The development of the system will be supported by a graphical simulator and by a robotic manipulation testbed constructed using the Utah/MIT hand. This capability will allow the user to learn how to properly express tasks in wrench space, and will provide insight into methods of learning in manipulation.

The system illustrated in Figure 8 is being constructed to conduct experiments. The manipulation society will run on a VAX 750 and create a child process which maintains a command queue. The planner can then proceed at its characteristic rate while the child process submits subtasks to the mechanical systems and waits for the completion of those tasks asynchronously.

The task submitted to the mechanical arm is a homogeneous transform representing the hand frame position and orientation. The transform is translated into a sequence of VAL II[1] expressions which deliver the hand frame to the correct position and orientation.

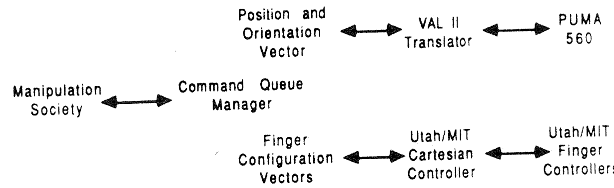

Figure 8: The Control Structure for a Manipulation Test Bed

The tasks submitted to the Utah/MIT hand consist of the position of the contact site, the orientation of the last phalange, and stiffness matrix for every finger in hand frame coordinates. This task is distributed over four Motorola 68020's which act as a Cartesian controller for the hand.

The task may require the system to partition the contact set into cooperating subsets with independent tasks. This allows a system which has become ill conditioned in the process of executing a task to stably constrain the object with a subset of the contact system while reconditioning the complementary subset. The organization of the system lends itself to the partitioning of contact elements into sets which may respond to independent tasks. The ability of the system to accommodate tasks which require finger replacement, or which require multiple tasks within a contact system will be developed.

The Conditioning Agent concludes by defining commands to the arm and hand relative to a virtual object. The actual contact must be achieved by implementing a compliant guarded move. The Utah/MIT hand facilitates this action by combining Cartesian stiffness control with low-level reflexive movements, such as proximal stiffening and distal curling [8]. An unexpected contact on a proximal phalange causes the distal kinematic chain to curl, while more proximal joints stiffen. Taction is not required for this behavior, contacts are sensed by measuring tendon tensions. Reflexive movements can be applied at low levels to improve the application of the fingers to the object while these contact sites are managed by the planner.

A great deal of specific knowledge is recorded in the models of the hand and object; therefore, modeling is a primary concern. The demonstrations presented earlier represent the cylinder analytically, since such an approach works well for primitive shapes. In general, it will be necessary to tag surface models with the surface's wrench space, and to build derivative wrench spaces which highlight graspable surface elements. The surface model will instantiate a *realm of influence* for the derivative wrench spaces, thus accenting surface elements which are especially useful (such as *finger sized* concavities in the surface). The exact nature of the models used to represent general classes of objects and the automatic generation of such models from a CAD representation will be examined further.

We also wish to investigate the extent to which the object model can be used to support learning in manipulation. It is possible to learn which initial hand-object interactions permit the system to migrate to a solution. We would also like to examine the automatic classification of virtual fingers, to exploit the reduction in complexity warranted in certain tasks. Skill refine-

[1] VAL II is the control language for the PUMA 560 robot used in these experiments

ment behaviors based on these ideas will be investigated further.

References

[1] P.K. Allen. Surface Descriptions from Vision and Touch. In *Proceedings of the 1984 Conference on Robotics*, pages 394–397, IEEE, Atlanta, Georgia, March 1984.

[2] M. Arbib, T. Iberall, and D. Lyons. *Coordinated Control Programs for Movements of the Hand.* Technical Report COINS 83-25, University of Massachusetts, 1983.

[3] M.R. Cutkosky and P.K. Wright. Modeling Manufacturing Grips and Correlations with the Design of Robotic Hands. In *Proceedings of 1986 Conference on Robotics and Automation*, pages 1533–1539, IEEE, San Francisco, CA, 1986.

[4] G.H. Golub and C.F. Van Loan. *Matrix Computations.* The Johns Hopkins University Press, Baltimore Maryland, 1983.

[5] R.A. Grupen, K. Biggers, T.C. Henderson, and S. Meek. *Task Defined Internal Grasp Wrenches.* Technical Report UUCS-88-001, Department of Computer Science, University of Utah, 1988.

[6] H. Hanafusa and H. Asada. A Robot Hand with Elastic Fingers and its Application to Assembly Process. In *IFAC Symposium on Information and Control Problems in Manufacturing Technology*, pages 127–138, Tokyo, 1977.

[7] S.C. Jacobsen, E.K. Iverson, D.F. Knutti, R.T. Johnson, and K.B. Biggers. Design of the Utah/MIT Dextrous Hand. In *Proceedings of the 1986 Conference on Robotics and Automation*, pages 1520–1532, IEEE, San Francisco, CA, April 1986.

[8] S.C. Jacobsen, J.E. Wood, D.F. Knutti, and K.B. Biggers. The Utah/MIT Dextrous Hand Work in Progress. *Journal of Robotics Research*, 3(4):21–50, Winter 1984.

[9] J. Kerr and B. Roth. Analysis of Multifingered Hands. *Journal of Robotics Research*, 4(4):3–17, Winter 1986.

[10] C. Klein and B. Blaho. Dexterity Measures for the Design and Control of Kinematically Redundant Manipulators. *Journal of Robotics Research*, 6(2):72–83, Summer 1987.

[11] H. Kobayashi. Control and Geometrical Considerations for an Articulated Robot Hand. *Journal of Robotics Research*, 4(1):3–12, Spring 1985.

[12] Z. Li and S. Sastry. Task Oriented Optimal Grasping by Multifingered Robot Hands. In *Proceedings of 1987 Conference on Robotics and Automation*, pages 389–394, IEEE, Raleigh, NC, April 1987.

[13] D. Lian, S. Peterson, and M. Donath. A Three-Fingered, Articulated Hand. In *Proceedings of the 13th International Symposium on Industrial Robots*, pages 18–91 to 18–101, SME, Chicago, IL, April 1983.

[14] M. Minsky. *The Society of Mind.* Simon and Schuster, New York, NY, 1986.

[15] T. Okada. Object-Handling System for Manual Industry. *IEEE Journal of Systems, Man, and Cybernetics*, 9(2):79–89, February 1979.

[16] J.K. Salisbury. *Kinematic and Force Analysis of Articulated Hands.* PhD thesis, Stanford University, May 1982.

[17] R. Tomovic, G. Bekey, and W. Karplus. A Strategy for Grasp Synthesis with Multifingered Robot Hands. In *Proceedings of the 1987 Conference on Robotics and Automation*, pages 83–89, IEEE, Raleigh, NC, April 1987.

[18] T. Yoshikawa. Manipulability of Robotic Mechanisms. *Journal of Robotics Research*, 4(2):3–9, Summer 1985.