Robot Hands and the Mechanics of Manipulation

Matthew T. Mason
J. Kenneth Salisbury, Jr

Robot Hands and the Mechanics of Manipulation

Matthew T. Mason
J. Kenneth Salisbury, Jr.

The MIT Press
Cambridge, Massachusetts
London, England

PUBLISHER'S NOTE

This format is intended to reduce the cost of publishing certain works in book form and to shorten the gap between editorial preparation and final publication. Detailed editing and composition have been avoided by photographing the text of this book directly from the authors' prepared copy.

This book was set under the direction of the authors using the TEX typesetting system and printed and bound by Halliday Lithograph in the United States of America.

Library of Congress Cataloging in Publication Data

Mason, Matthew T.
Robert hands and the mechanics of manipulation.

(The MIT Press series in artificial intelligence)
Includes bibliographies and index.
1. Robotics. 2. Manipulators (Mechanism)
I. Salisbury, J. Kenneth. II. Title. III. Series.
TJ211.M366 1985 629.8'92 85-3
ISBN 0-262-13205-2

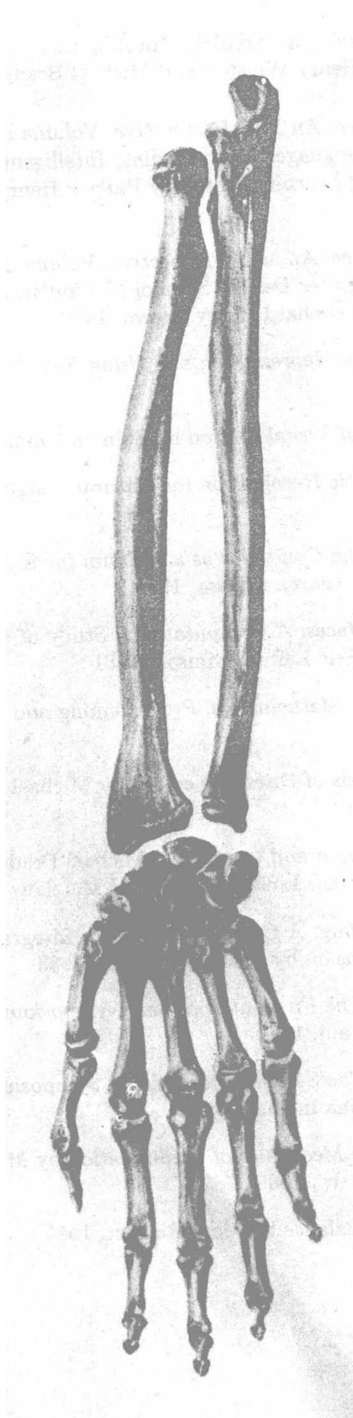

D. Salisbury

The MIT Press Series in Artificial Intelligence
Edited by Patrick Henry Winston and Michael Brady

Artificial Intelligence: An MIT Perspective, Volume I: Expert Problem Solving, Natural Language Understanding, Intelligent Computer Coaches, Representation and Learning edited by Patrick Henry Winston and Richard Henry Brown, 1979

Aritificial Intelligence: An MIT Perspective, Volume II: Understanding Vision, Manipulation, Computer Design, Symbol Manipulation edited by Patrick Henry Winston and Richard Henry Brown, 1979

NETL: A System for Representing and Using Real-World Knowledge by Scott Fahlman, 1979

The Interpretation of Visual Motion by Shimon Ullman, 1979

A Theory of Syntactic Recognition for Natural Language by Mitchell P. Marcus, 1980

Turtle Geometry: The Computer as a Medium for Exploring Mathematics by Harold Abelson and Andrea diSessa, 1981

From Images to Surfaces: A Computational Study of the Human Early Visual System by William Eric Leifur Grimson, 1981

Robot Manipulators: Mathematics, Programming and Control by Richard P. Paul, 1981

Computational Models of Discourse edited by Michael Brady and Robert C. Berwick, 1982

Robot Motion: Planning and Control by Michael Brady, John M. Hollerbach, Timothy Johnson, Tomas Lozano-Pèrez, and Matthew Mason, 1982

In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension by Michael G. Dyer, 1983

Robotics Research: The First International Symposioum edited by Michael Brady and Richard Paul, 1984

Robotics Research: The Second International Symposium edited by Hideo Hanafusa and Hirochika Inoue, 1985

Robot Hands and the Mechanics of Manipulation by Matthew T. Mason and J. Kenneth Salisbury, Jr., 1985

Legged Robots that Balance by Marc Raibert, 1985

SERIES FOREWORD

Artificial intelligence is the study of intelligence using the ideas and methods of computation. Unfortunately, a definition of intelligence seems impossible at the moment because intelligence appears to be an amalgam of so many information-processing and information-representation abilities.

Of course psychology, philosophy, linguistics, and related disciplines offer various perspectives and methodologies for studying intelligence. For the most part, however, the theories proposed in these fields are too incomplete and too vaguely stated to be realized in computational terms. Something more is needed, even though valuable ideas, relationships, and constraints can be gleaned from traditional studies of what are, after all, impressive existence proofs that intelligence is in fact possible.

Artificial intelligence offers a new perspective and a new methodology. Its central goal is to make computers intelligent, both to make them more useful and to understand the principles that make intelligence possible. That intelligent computers will be extremely useful is obvious. The more profound point is that artificial intelligence aims to understand intelligence using the ideas and methods of computation, thus offering a radically new and different basis for theory formation. Most of the people doing artificial intelligence believe that these theories will apply to any intelligent information processor, whether biological or solid state.

There are side effects that deserve attention, too. Any program that will successfully model even a small part of intelligence will be inherently massive and complex. Consequently, artificial intelligence continually confronts the limits of computer science technology. The problems encountered have been hard enough and interesting enough to seduce artificial intelligence people into working on them with enthusiasm. It is natural, then, that there has been a steady flow of ideas from artificial intelligence to computer science, and the flow shows no sign of abating.

Robotics is the part of Artificial Intelligence concerned with the intelligent connection of perception to action. A key issue in Robotics is understanding the mechanics of manipulation as a basis for the successful planning and execution of robot tasks. Currently, we cannot claim to understand manipulation. Our knowledge of robotics is almost entirely empirical. If fundamental principles exist, we have not discovered them. We consider this book to be an important step toward uncovering such principles.

The purpose of this MIT Press Series in Artificial Intelligence is to provide people in many areas, both professionals and students, with timely, detailed information about what is happening on the frontiers in research centers all over the world.

Patrick Henry Winston
Michael Brady

PREFACE

Robotic manipulation began less than twenty–five years ago, when the first computer–controlled manipulator was demonstrated. Since that time, scientists and engineers have designed hundreds of different manipulators, and demonstrated thousands of different applications. Nonetheless, we cannot claim to understand manipulation. If fundamental principles exist, we have not yet discovered them. Our knowledge of robotics is almost entirely empirical—the result of twenty–five years of experimentation. When general observations arise, they have the character of design heuristics, rather than of scientific principles. Theoretical tools can be drawn from related fields, but there is very little theory that bears directly on manipulation. The most important product of this twenty–five years of empiricism is an accumulation of experience, from which we might identify key issues and formulate new research directions. One lesson that has been demonstrated many times is that manipulation involves much more than just moving an arm around. The mechanical interactions between the hand and the task, and the interactions between different objects in the task, are key elements of the manipulation process.

The present state of robotic manipulation is that we can build mechanical arms and we can move them around accurately, but we do not know how to use them effectively. Robotic manipulators are competent only in the simplest application domains. There are many instances of spray painting and welding robots, but these are simple tasks which move only the robot about. Most real manipulation performed by robots is limited to simple packing and stacking operations. There are some cases of complex robotic assembly but these usually involve extensive fixturing and special tooling. When complex relationships between parts arise, such as in assembly tasks, or when the task environment is unpredictable, robots are found wanting. Even the simplest assembly operations such as "getting" and "putting" involve complex sequences of motions and sensory operations. The simplicity with which humans can "get" or "put" something somewhere belies the extreme difficulty in telling a robot how to do the same thing.

To understand manipulation, we have to understand how objects move when we touch them and push on them. If we push a block resting on a surface it may slide in some way or it may not move at all. If we attempt to grasp an object by closing two fingers about it, it may immediately slide away, it may settle into a secure grasp, or it may be held in some way that still allows motion relative to the hand. To predict the outcome in such situations, we must understand how the actions of the robot affect the motions of objects.

This book looks at several aspects of the basic mechanics of grasping, pushing and, in general, manipulating objects. The body of the book is derived from two Ph.D. theses in robotic manipulation. Salisbury's thesis, "Kinematic and Force Analysis of Articulated Hands," was written while he was at Stanford University. Mason's thesis, "Manipulator Grasping and Pushing Operations," was written while he was at the Massachusetts Institute of Technology. The two works have been brought together here because they provide perspectives which are both complementary and mutually reinforcing. Both theses feature theoretical studies of the mechanics of manipulation, and both theses focus on the interactions between the manipulator and the task. A realistic treatment of friction is central to both works. Salisbury includes friction's role as a constraint in his study of grasping and manipulating objects. The results are used to investigate the design and operation of hands. Mason studies characteristics of object motion in the presence of friction, and uses the results to demonstrate automatic planning of pushing operations.

More specifically, Salisbury's thesis, "Kinematic and Force Analysis of Articulated Hands," explores the mechanics of constraint and freedom that arise between rigid bodies in contact. This analysis is used to address the problem of designing and controlling robot hands for manipulation. The analysis aims at identifying "hand" designs which can securely hold and arbitrarily move a grasped object. The thesis derives a precise formulation for these two fundamental hand functions and then applies the formulation as a design criterion for a multi-fingered robot hand.

When contact occurs between two bodies the relative motion between them will depend upon the geometric and friction properties of the contact. Further, the same contact will permit a range of forces and moments to be transmitted through the contact. These forces and moments (or collectively—forces) will be orthogonal to the permitted motions in the sense that no force that can be transmitted through the contact can do any work along any motion allowed by the contact. The mathematical tools necessary for understanding the geometries of these motions and forces are introduced. Thus, a basic introduction to screws, twists, and wrenches is given. The effects of multiple contacts with a body (i.e. fingers and a grasped object) are then studied using both a geometric and an analytical approach. The techniques are applied to a large set of systematically generated candidate hand designs and used to identify acceptable hand kinematics. Further analyses provide insight into the selection of finger joint axis orientations and link dimensions so as to maximize force exertion accuracy.

Salisbury's work also focuses upon several basic issues in hand control. The identification of external, or net, forces and the internal, or grasping, forces follows as an extension of the multiple finger constraint analysis. The result is the introduction of the *grip matrix*. This transformation relates the fingertip force and motion states to the grasped object force and motion states. The fundamental importance of the grip matrix in hand control parallels that of the Jacobian matrix in arm control; both express the differential relationship between input and output motions of the respective systems. The grip matrix is ultimately used in deriving a strategy for controlling the grasped object's motion and apparent stiffness.

One practical application of Salisbury's work has been the development of the Stanford/JPL robot hand. This hand is one embodiment of the design principles presented and provides a suitable mechanism for studying a variety of hand control and dexterity issues. The arrangement of the tendons (cables) to actuate the finger joints, and the use of tendon tension sensors are particularly novel aspects of the hand. Throughout the design a minimization of complexity was stressed. It employs the minimum number of fingers and joints required for securely holding and arbitrarily moving a grasped object. The fingers use the minimum number of tendons required for independent control of motion and force exertion. In the same spirit, the Section I appendix includes an analysis of the minimum number of (force) sensors required for determining the basic local geometry of objects contacting the fingers.

Mason's thesis, "Manipulator Grasping and Pushing Operations," explores the use of pushing in robotic manipulation. The first part of the thesis is a systematic exploration of the mechanics of pushing, to determine how an object moves under the combined influence of the manipulator and the forces of sliding friction. The second part of the thesis applies this theory to some examples, demonstrating verification and automatic planning of some simple manipulator operations.

Since the use of pushing in manipulation is a previously unexplored phenomenon, the first question is whether pushing plays a significant role. Mason lists a number of situations in which pushing is important:

- Aligning a lid on a box, or straightening a pack of cards, is quickly and accurately accomplished by squeezing.

- When placing a grasped object, the mechanics of pushing applies if the object slips in the fingers.
- When grasping an object, pushing can eliminate position and orientation uncertainty.
- Pushing can be used to manipulate several objects simultaneously, as when straightening a pack of cards, or when sweeping a floor.

Pushing is different from simple positioning in a very fundamental way. The motion of an object being pushed is only partially determined by the manipulator. The uncontrollable (and often unpredictable) frictional forces also help to determine the object's trajectory. Consequently, it is not possible to predict, much less control, the trajectory of an object. The most surprising result of Mason's thesis is that this apparent deficiency does not prevent useful manipulation. The partial nature of the control is, in fact, the source of some of the advantages of pushing.

This difference requires a change in our view of robot planning. With conventional approaches to manipulation, we assume that we can grasp an object, and afterward control the position of the object. Usually we can express the goals of a task in terms of desired positions of a set of objects, so that the planning problem is, in principle, straightforward. First we decide in what order to move the objects. Then for a given object we must find a collision-free path from its initial position to its goal position. According to this conventional *pick-and-place* model of manipulation, to plan a manipulation task is, in essence, a kinematic problem.

Pushing departs drastically from the pick-and-place model of manipulation. The robot cannot, by itself, control the trajectory of an object. In order to achieve a particular goal, the planner must consider the input from nature, i.e. the intrinsic mechanics of the task, and then choose actions which will accomplish the goal. The proper choice of an action depends on the nature of the immediate goal (absolute position, relative position, alignment of geometric features, incremental displacement, etc.) and on the mechanics of the task domain. In short, the planner has to know a lot more about the physical world. The payoff is that pushing can be used where positioning cannot, allowing us to perform useful manipulation in situations where robots are now useless.

Thus, we have brought together two perspectives on the basic mechanics of real manipulation. Our point of view is that robot tasks encompass far more than moving the end of the manipulator. This broader view of robot potentials may be a bit unsettling in that it forces us to realize how little we consciously understand about our own manipulative capacities. More than that, it compels us to consider innovative ways for using robots to do useful work.

Consider, for example, all the neglected surfaces on a robot. Only a small percentage of the robot's surface area is available for doing work. Why can't we use parts of the forearm, or any accessible surface, for pushing or holding objects? A human, when faced with carrying a load of logs, does not try to pick them up with his fingertips and daintily place them in the woodbox. More likely, he uses his fingertips to pry between some of the logs, moves them onto his forearm, and holds them against his chest. A person needing to hold a bag while both hands are already occupied may, without much thought, press it against his side with his elbow. Both these examples demonstrate an intuitive understanding of the mechanics of constraint. Given the problem of holding something, we rapidly consider what surfaces are available for the task and select the ones that will suit our need. True, the mechanical design of present day robots precludes such rash behavior but not because of any inherent limitation.

When we are told not to drop an object we know that squeezing it the right amount, supporting it from below, balancing it and moving smoothly will aid in achieving that goal. Again, an intuitive understanding of the mechanics of friction and stable grasping come into play.

A man running a back hoe (hardly a dexterous machine) can still use it to move a concrete block into a precise position by pushing on it in just the right way to make it slide across the ground. The same supposedly clumsy machine, with the proper commands from a skilled operator, can be made to literally climb into its trailer by using its hoe and front scoop as simple legs. In this case, a relatively simple mechanism can be made to accomplish tasks it was never intended to do because of the skill of its operator and his intuitive understanding of the basic task mechanics.

All these examples rely on a variety of perceptual, conceptual and mechanical capabilities on the part of the human. An important question is what are the essential elements of these capabilities and how might they be used to guide us toward more innovative robot design and use. Perhaps even more importantly, we must also consider what non-human capabilities we can introduce to push robot capabilities beyond human capacity.

To illustrate some of the assumptions we use in manipulation, let us look at the simple task of "getting" an object and "putting" it at some desired location. When a robot is programmed to perform the same task, all of the following issues are present either explicitly, as actual robot commands, or implicitly, in the programmer's understanding of the problem.

Consider what must be known before we can begin to "get" the object. Is the object within reach? Do we have to know its location or not? There are a variety of ways that we can exercise control over the object. Assuming we must touch it, to what degree must we assume control over it? Can it simply be pushed to our desired goal without actually grasping it? Perhaps we can propel it toward our goal by striking it. If we must grasp the object, what constitutes acceptable acquisition? Do we have to hold it securely or is some freedom of motion in the grasp permissible or even desirable? If it must be held securely what constitutes the most secure grasp? Should the grasp use structural restraint or is some degree of frictional restraint necessary or useful? Is there an accessible grasping posture compatible with our final placement of the object or will it have to be regrasped at some point during the task? In moving the object, is there a trajectory that is best in some sense such as efficiency, minimum time or obstacle avoidance? Can we move very quickly or are there limitations on power and acceleration?

Assuming we have achieved a degree of control over the object's motions by touching, pushing or grasping, we must still consider how to reach the goal state. The actual goal, "putting" in this case, must be specified in a way that permits planning how to reach it and knowing when we are done. If a robot is programmed to put something in a defined location then the task is simply a matter of driving the robot to these coordinates, if there are no obstacles or tight fits. But, what if the goal is putting an object into a tightly constrained goal position? Is there some strategy that will help us get there without excessive force? Can we use the geometry of the goal to help specify a strategy for the motion? Can the motion be a single strategy such as "move until contact" or move with a specified stiffness or is it necessary to look for distinct states along the way, and change our strategy at certain critical points? How do we recognize what state the task is in? Are simple sensor readings adequate or is a more complex interpretation of sensory information required? Can we use instantaneous sensory information or must we refer to time histories of sensor readings and complex task models? Are there sensing modalities that will require some active sensing such as moving or searching?

How will we know if there are errors occurring and whether they are fatal or not? Is testing whether or not we reached the goal sufficient error checking or are there other errors during task execution that must be considered? If an error occurs how do we determine if and how we can recover? If we can recover from the error will simply repeating the attempt with small modification suffice, or must we try an entirely new strategy?

Finally, how will we know when we have reached the goal? Do we even need to check? If we do, simple goals, such as position states or force states are directly confirmable from basic sensors. More complex goals such as "put this object into that one" or "grind until smooth" will require sophisticated sensors and interpretation.

The question of appropriate accuracy underlies all of the above issues. The precision with which goal statements are made must be compatible with the accuracy with which actions and sensed information are available. Judicious choice of strategy can greatly reduce the need for accuracy, as in Mason's pushing examples. In other cases the accuracy requirements for a strategy in one domain can be exchanged for accuracy requirements in another domain. The accuracy required for sliding an object along a constraint surface under position control is much more stringent than when doing the same task under force control.

The issues discussed in the get and put task were presented in the context of a programmer trying to plan and program the task with a robot. The issues are typical of those concerning a robot user with a given robot and a particular task to accomplish. Often robot users find themselves in the embarrassing situation of having a robot and trying to find a task that they can accomplish with it. While a better understanding of the mechanics of manipulation will certainly improve our ability to use robots, the real value may be in what it tells us, as designers, about how to build robots.

Typically, robots are designed with some intuitive specifications of required reach, speed, and force capabilities. The relative utilities of speed and force, range and resolution, structural stiffness and compliance, as well as bandwidth and efficiency, are not well understood. The attribute of "dexterity" is probably one of the least well understood and most misrepresented concepts in robot design. While mechanical redundancy appears to offer advantages in improved reliability, obstacle avoidance, and efficiency, we do not yet understand the real synergy of redundant systems in improving dexterity. Today's robot arms and hands are designed separately and yet, to achieve the best combination of capabilities, they must be considered simultaneously. The improvement of overall performance of a robot will depend not only on better mechanical systems, but upon the application of control systems that take into account the complete system kinematics and dynamics.

The impact of understanding task mechanics will bear strongly on sensor issues. The realization that full utilization of a robot goes beyond positioning it poses questions about where and what type of sensors will be of most value. The accuracy and range of the sensors must be at least as good as the mechanics of the system they are used to observe. If we plan to move objects by pushing or manipulating them in a hand, our sensors must give us information to allow us to adequately observe the process. The idea that manipulation within a hand may be viewed as controlling contact trajectories across the fingers suggests a need for high bandwidth contact resolving sensors described in the Section I appendix. Sensors must be compatible with the capabilities of the mechanism. A tactile sensor may be rendered useless by a robot unable to control contact forces sufficiently. As tasks become more complex, the quantities being sensed will involve more sensors and more complex interpretation. The efficient monitoring of task progress and comparison with internal models and expectations will require sophisticated ways of interpreting sensory data from a variety of possibly overlapping sources and modalities.

It is clear that the many aspects of robot design and use are intimately related. It is by understanding how we really want the robot to perform, and the mechanical domain it will work in, that we will begin to address the problems of robot design and utilization in an effective way. The task of designing and using fully integrated robot systems is a formidable one; it is our hope that this book will add a few useful pieces to this puzzle for you.

For both of us, this book is the product of many years of pleasant association with the Stanford and MIT Robotics and Artificial Intelligence research communities. In one sense it is a gift to those who have challenged and inspired us to undertake this work. In another sense, it is a challenge to you to continue it and improve upon it.

J. Kenneth Salisbury, Jr.
Matthew T. Mason
Cambridge, Massachusetts
November 1984

ACKNOWLEDGMENTS

We would like to thank the many friends and associates who have advised and inspired our work. Mike Brady, at MIT, suggested this collaboration and encouraged us along the way. Terry Ehling, at the MIT Press, patiently stuck with us through the final manuscript preparation as did Dan Brotsky, at MIT, with numerous feats of TeX magic. Funding for our work has been from a variety of sources including the National Science Foundation, the National Aeronautics and Space Administration, the System Development Foundation, the Advanced Research Projects Agency and the Office of Naval Research.

My most sincere thanks are extended to Bernie Roth at Stanford University for his inspiration and sensitive encouragement of the work presented in Section I. Among the many who have helped me, I am especially grateful to Tom Binford at Stanford and Alison Shafer. Tony Bejczy and Carl Ruoff at Jet Propulsion Labs have been supportive of my design efforts through the years. Finally I am grateful to my parents for their unquestioning support, and my sister, Deborah, for the hand drawing in the front of the book.

J.K.S.

Many of the ideas in Section II were born during discussions with Tomas Lozano-Perez at M.I.T., and likely should be credited to him. I would like to thank Berthold Horn, John Hollerbach, and Patrick Winston at M.I.T. and Marc Raibert at Carnegie-Mellon University, for their guidance and examples. I am also grateful for the support of my parents. Mary Mason drew all sixty figures in record time, and provided both moral and financial support.

M.T.M.

CONTENTS