Annual Review of Control, Robotics, and Autonomous Systems

Toward Robotic Manipulation

Matthew T. Mason

Robotics Institute and Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA; email: matt.mason@cs.cmu.edu

Annu. Rev. Control Robot. Auton. Syst. 2018. 1:1–28

First published as a Review in Advance on March 1, 2018

The *Annual Review of Control, Robotics, and Autonomous Systems* is online at control.annualreviews.org

https://doi.org/10.1146/annurev-control-060117-104848

Keywords

robot, manipulation, evolution, engineering

Abstract

This article surveys manipulation, including both biological and robotic manipulation. Biology inspires robotics and demonstrates aspects of manipulation that are far in the future of robotics. Robotics develops concepts and principles that become evident only in the creative process. Robotics also provides a test of our understanding. As Richard Feynman put it: "What I cannot create, I do not understand."

1. INTRODUCTION

This article surveys manipulation with a broad perspective, including all agents, whether natural (animals, including humans), artificial (robots), or a mix (telerobots).

Manipulation is surprisingly widespread among animals. The chief advantage of biology over robotics is the intelligence and perception of the individual, enabling breadth, robustness, and adaptability at levels well beyond the state of the art in robotics. Manipulation played a significant role in the evolution of humans, starting with tool use (most notably throwing missiles), followed by tool fabrication. Human performance differs from robotic performance in many ways. For example, humans grasp skillfully and change grasps efficiently, yet are less dependent on grasping than robots.

The chief advantage of robotics over biology is the intelligence in the design process. Engineers often design systems that outperform animals in specific tasks. Examples include the extraordinary precision and repeatability of robotic spray painting, handling of heavy loads, and working in extreme environments.

What have we learned from robotics? The number-one lesson is that manipulation is hard. This is contrary to our individual experience, since almost every human is a skilled manipulator. There are also many more lessons—a long list of concepts and principles arising from the attempt to close the gap with biological manipulation.

The plan for this article is to survey and compare biological and robotic manipulation in more detail, focusing on the similarities, differences, and lessons learned, with the ultimate goal of identifying productive directions for future robotics research. We begin by defining terms.

2. DEFINING MANIPULATION

Very few definitions of manipulation appear in the robotics literature. A European research road map defined manipulation as "the function of utilising the characteristics of a grasped object to achieve a task" (1, p. 38). A NASA road-mapping effort yields the following: "Manipulation pertains to making an intentional change in the environment or to objects that are being manipulated" (2, p. 13). My own earlier attempt at defining manipulation was "using one's hands to rearrange one's environment" (3, p. 1). Rather than sorting the pros and cons of those definitions, let us apply the shotgun method and identify every approach that we can.

2.1. Etymological

First, we can appeal to the origins of the term.

Definition 1 (etymological). Manipulation refers to the activities performed by hands.

The etymological definition gives little guidance in determining whether an activity is manipulation or not, although it captures a shared understanding and point of departure for the discussion. It helps us ask the right question: What are the intrinsic defining criteria of manipulative activities?

2.2. Genus/Differentia, Ends Only

One way to characterize the hands' activities is to employ the genus/differentia approach: To define a given term, we identify a superclass (genus) of the term and then specify criteria (differentia) to distinguish the term. Informally, the genus/differentia approach might be described as top down.

This approach nicely fits the way we often think of robotics—as divided into locomotion (moving oneself) and manipulation (moving everything else).

> **Definition 2 (genus/differentia, ends only).** Manipulation is when an agent moves things other than itself.

What if gravity causes the agent's hand to fall on something while the agent sleeps? Does that count as an agent activity? Does the answer depend on whether the agent was conscious and desirous of the falling hand's effects? Fortunately, philosophers are already wrestling with this issue. Davidson (4, p. 42) provides a similar example: "I might have turned off the light by inadvertently brushing against the switch; would it then have been my deed, or even something that I did?" I propose that we leave this problem to the philosophers and rely upon them to determine what "agent activity" means. Or, if you wish, take "agent" to mean "human, animal, or robot."

2.3. Genus/Differentia, Ends and Means

The ends-only definition might be a bit too broad for some tastes. For example, when an automatic door swings open, do we consider that to be an act of manipulation? The essence of manipulation might be that effectors can interact selectively with many different things to attain many different outcomes.

> **Definition 3 (genus/differentia, ends and means).** Manipulation is when an agent moves things other than itself through selective contact.

2.4. Bottom Up

Instead of the top-down genus/differentia approach, we could go bottom up, defining manipulation as the union of several different types of manipulation. For example, if we select the classification of manipulation favored by roboticists, we obtain a constantly evolving definition, a definition that expands as robotic manipulation progresses.

> **Definition 4 (bottom up).** Manipulation is pick-and-place manipulation plus in-hand manipulation plus mechanical assembly plus. . . .

With the bottom-up approach, every taxonomy of manipulation gives rise to a definition. In particular, if robotics research is described as a sequence of demonstrated capabilities, then the union over that sequence serves as an evolving definition of manipulation. We are literally defining manipulation as we go.

2.5. Boundaries

The biggest problem in applying these definitions is to determine the boundaries. If manipulation means moving things other than the agent, then we must determine the boundary between the agent and the task. This often seems simple enough, but consider the process of combing your hair or painting your nails. What about when one hand is manipulating the other hand? As the saying goes, "one hand washes the other." In such cases, a part of the agent seems to cross the boundary to become part of the task domain.

Tools present a similar challenge. Is the tool part of the agent and thereby doing the manipulating? Or is it part of the task domain, being manipulated?

Likewise, it can be challenging to determine the boundary between agents. In factories, it is common to refer to each arm as a robot and to everything else as the environment. But why not count all the arms as one robot? And why should an actuated fixture or vise, engineered right along with the arms and grippers, be considered part of the environment rather than part of the robot?

2.6. Duality

Distinguishing between manipulation and locomotion also presents an interesting challenge. We say that locomotion refers to agent motion and that manipulation refers to nonagent motion. But that difference might be nothing more than a choice of reference frame. Newtonian relativity suggests that it is not a meaningful difference. Thus arises a duality between manipulation and locomotion: Locomotion of an agent relative to its support can be viewed as manipulation of that support relative to the agent. At human scale, to view walking as manipulation of the Earth seems a bit absurd, but the idea gains some force at the smaller scale of ants and spiders, or the cilia of single-celled organisms.

Duality is a useful source of research ideas. For each lesson learned in locomotion, we can look for a dual lesson in manipulation, and vice versa. For example, manipulation is often framed in terms of finding a motion to achieve a goal state, whereas locomotion is often framed in terms of gaits—periodic cycles of motion. Duality suggests that we should also explore the idea of periodic gaits in manipulation, which leads to juggling (5, 6).

2.7. Beyond Motion

Both of the top-down genus/differentia definitions share one failing: They refer only to motion. Do we exclude the process of polishing a surface, or do we consider that to be motion at a microscopic scale? Do we exclude operations that prevent motion, such as placing one's hand on a stack of papers to keep them from blowing away? What about rubbing one's hands together to warm them? The NASA definition refers to a "change in the environment," which is broader than just motion, but perhaps too broad. For the purposes of this article, we will adapt the genus/differentia ends-and-means definition.

> **Definition 5.** Manipulation refers to an agent's control of its environment through selective contact.

We will consider the placement of the boundary between agent and environment to be a choice that may vary with context. We will interpret "control" broadly enough to include motion, as well as processes such as wiping, polishing, painting, and burnishing, with specific choices again varying as desired.

3. ANIMAL AND HUMAN MANIPULATION

This section sketches the history of animal and human manipulation, examines a few examples, and discusses some of the features and advantages with particular relevance to robotics.

3.1. History of Biological Manipulation

When did manipulation begin? At a very small scale, flagella and cilia are employed for both locomotion and manipulation. Study of the cytoskeletons of different families of eukaryotes suggests that flagella evolved before eukaryotes—before the last common ancestor of all plants, animals, and fungi (7). Fossils identified as likely eukaryotes are as old as 1,250–1,600 million years. Thus, at least at a very small scale, it is reasonable to say that manipulation is universal among animals and appeared over a billion years ago.

Manipulation at a human scale and at human skill levels is more recent. The ability to throw with power and accuracy appeared in our ancestors starting about three million years ago. Throwing missiles is an early example of tool use, but surely not the earliest. More rudimentary tool use is readily observed among many animals. Over the next several hundred thousand years, our ancestors evolved more refined manipulation skills, including the use of tools to make better tools, such as stone tools.

We know of the evolution of human throwing, tool use, and toolmaking skills from the study of the fossil record, comparative anatomy of modern primates, and biomechanical studies (8–11). Lucy (*Australopithecus afarensis*, 3.2 million years ago) (12) had some but not all of the features associated with throwing and tool use (8, 13). The earliest evidence of stone tool manufacture appeared at about the same time, although evidence of more frequent and systematic tool manufacture did not appear until later, about 2.7 million years ago, with the more sophisticated Acheulian tools appearing about 1.7 million years ago (10).

Humans excel at throwing. This will come as a surprise to some readers. The human ability to throw with power and accuracy was noted almost 150 years ago by Darwin (14, p. 140), who contrasted human throwing with that of some other animals, observing that other animals "roll down stones or throw them at their enemies; nevertheless, they perform these various actions clumsily, and they are quite unable, as I have myself seen, to throw a stone with precision." Human throwing performance is dramatically better, sometimes exceeding 160 km/h at the highest levels of baseball and cricket (15, 16).

It may also come as a surprise to learn just how effective rock throwing can be. There are documented incidents of an individual human prevailing over a pack of wild dogs or baboons by throwing rocks (17, 18). There are also historical incidents where rock-throwing, club-swinging indigenous groups defeated gun-toting explorers and invaders (19).

What will surprise no one is that modern society places an extraordinary value on throwing and other manipulation skills. There is a notable interest in all types of manipulation, reflected in our fascination with juggling, sleight of hand, and even cup stacking, but most especially with throwing. Many of us encourage our children to develop their throwing skills. When the most skilled professional athletes perform, we fill arenas in the thousands, we watch them on video in the millions, and we reward them with worldwide fame and fortune.

In summary, manipulation has been present over the entire history of animals, but the rise and evolution of humankind was accompanied by dramatic developments in manipulation. Studies show that human morphology is adapted specifically to improve our performance in throwing missiles and fabricating better tools. Humans evolved to use naturally occurring tools and to fabricate better tools, and then coevolved with those tools.

3.2. Animal Manipulation

We have already remarked on the smaller-scale manipulation exhibited by flagella and cilia and the larger-scale manipulation exhibited by humans. At an intermediate scale, it is hard to ignore the impressive manipulation demonstrated by spiders and ants. Spiders exhibit a wide range of manipulation skills. They weave a great variety of webs, they wrap prey, some of them cast nets to capture prey (20), and others even cut rivals' webs to steal their prey (21). Ants are, if anything, even more impressive (**Figure 1*a***), using a broad range of manipulative behaviors that evolved starting about 20 million years ago (22).

From the perspective of a roboticist, the variety and refinement of ant manipulation skills can be startling and humbling. How can something so tiny possess the intelligence implied by these behaviors? Perhaps Darwin (14, p. 145) said it best: "The brain of an ant is one of the most marvellous atoms of matter in the world, perhaps more marvellous than the brain of man." Have we ever built a robot as capable as an ant, at any scale?

Another example of insect manipulation is provided by the dung beetle (**Figure 1*b***), which uses an interesting gait to roll a dung ball along the ground. Its two front legs are placed on the ground, walking backwards. The four rear legs are placed on the dung ball, taking on the role of arms, although their motions are the same as if they were walking on the dung ball. Besides providing a nice example of mobile manipulation, the dung beetle's gait provides an elegant demonstration of the duality principle described in Section 2.6: a behavior that might be regarded as either locomotion or manipulation, depending upon the reference frame. The dung beetle illustrates duality by simultaneously employing two nearly identical behaviors, one locomotive and one manipulative. Or we could say that it is manipulating two balls, one small and one large.

Observations of tool use in animals have been widely reported, including use of naturally occurring tools as well as fabricated tools. Chimpanzee and bonobo termite fishing (**Figure 1*c***) is not as simple as it looks. The stem must be selected and prepared properly, and some skill is required to insert the stem, jiggle it in such a way that the termites' defensive instincts cause them to bite the stem, and then withdraw the stem in such a way that the termites are not scraped off. The technique is passed from chimp to chimp, and in one instance was also passed from chimp to a human naturalist who crashed the chimps' termite-fishing school (24).

Birds also provide many examples of impressive manipulative behaviors. Some of the most dramatic are the nest-building behaviors of weaverbirds (**Figure 1*d***). Some birds use tools and fabricate tools, including a remarkable example in which a New Caledonian crow was observed fabricating a hook from a piece of wire (25, 26) (**Figure 1*e***).

Not all studies focus on tool use. Measurements of the ratio of thumb length to index finger length suggest that gorillas might be closest to humans in dexterous manipulation. Byrne et al. (27) studied mountain gorillas and identified 190 different elemental manipulative actions used in preparing thistle plants for consumption.

These are just a few dramatic examples of animal manipulation. Manipulation is commonplace in nature, from the self-grooming of a housefly to the foraging behavior of apes. It seems that almost any animal would provide sufficient inspiration to keep roboticists busy for quite a while, but for reasons good or bad, we usually focus on humans.

3.3. Human Manipulation

Humans exhibit a wide variety of manipulative behaviors beyond throwing, but I will start by examining throwing in greater detail.

Powerful throwing is not a product of just the arm and hand; it involves the whole body. The body can be viewed as a serial chain of links coupled by muscles, tendons, and ligaments. The most powerful throws occur by loading and then releasing energy in this linkage, as if cracking a whip (13). Straightforward back-of-the-envelope calculations based on high-speed video (**Figure 2*a***) show that during the critical phase of the motion the baseball sustains an acceleration well above 100 g.

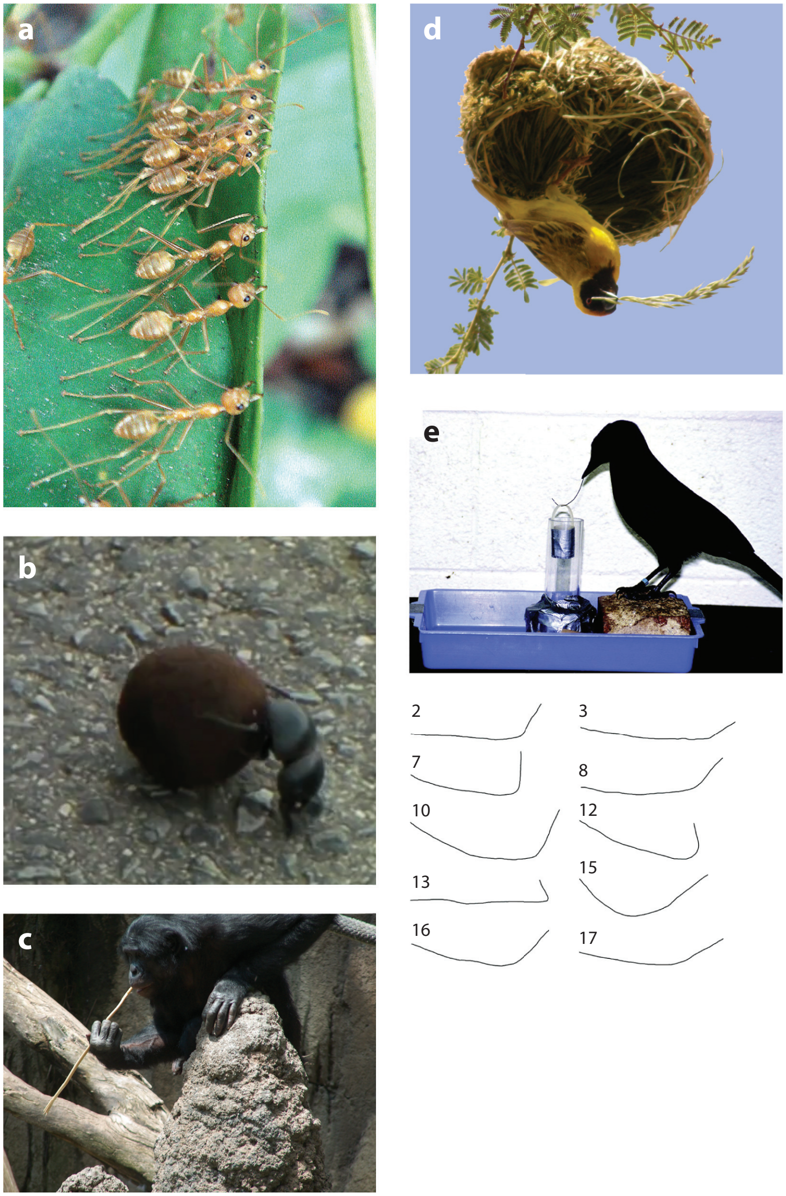

Figure 1

Examples of animal manipulation. (*a*) Weaver ants pulling two leaves together. (*b*) Dung beetle rolling a ball of dung. (*c*) Bonobo termite fishing. (*d*) Weaverbird building a nest. (*e*) New Caledonian crow crafting a hook to retrieve a morsel, along with some of the hooks fashioned. Panel *a* by Wikimedia user PHGCOM (**https://commons.wikimedia.org/wiki/File:AntsStitchingLeave.jpg**). Panel *b* from video by Wikipedia user NJR ZA (**https://en.wikipedia.org/wiki/File:Dung_beetle_working-001.ogv**). Panel *c* by Wikimedia user Mike Richey (**https://commons.wikimedia.org/wiki/File:A_Bonobo_at_the_San_Diego_Zoo_"fishing"_for_termites.jpg**). Panel *d* by Wikimedia user Gossipguy (**https://commons.wikimedia.org/wiki/File:Tisserin_Etosha.jpg**). Panels *a–d* reproduced under the Creative Commons Attribution-ShareAlike 3.0 Unported license (**https://creativecommons.org/licenses/by-sa/3.0**). Panel *e* adapted from Reference 23 with permission from AAAS.

The human manipulation skill that seems to attract the most attention from roboticists is our ability to grasp a wide range of objects, and several studies have proceeded with the goal of classifying the different grasps of the human hand (29–32). Often, these taxonomies distinguish precision grips, involving the fingers alone, from power grips, which also involve the palm.

Another human manipulative behavior that draws roboticists' attention is the turning and shifting of objects in the hand through the motion of the palm and fingers, referred to as dexterous manipulation, in-hand manipulation, or intrinsic movements (33). Fabrication of stone tools (**Figure 2*b***) is credited with helping to drive the evolution of fine manipulation skills. Marzke

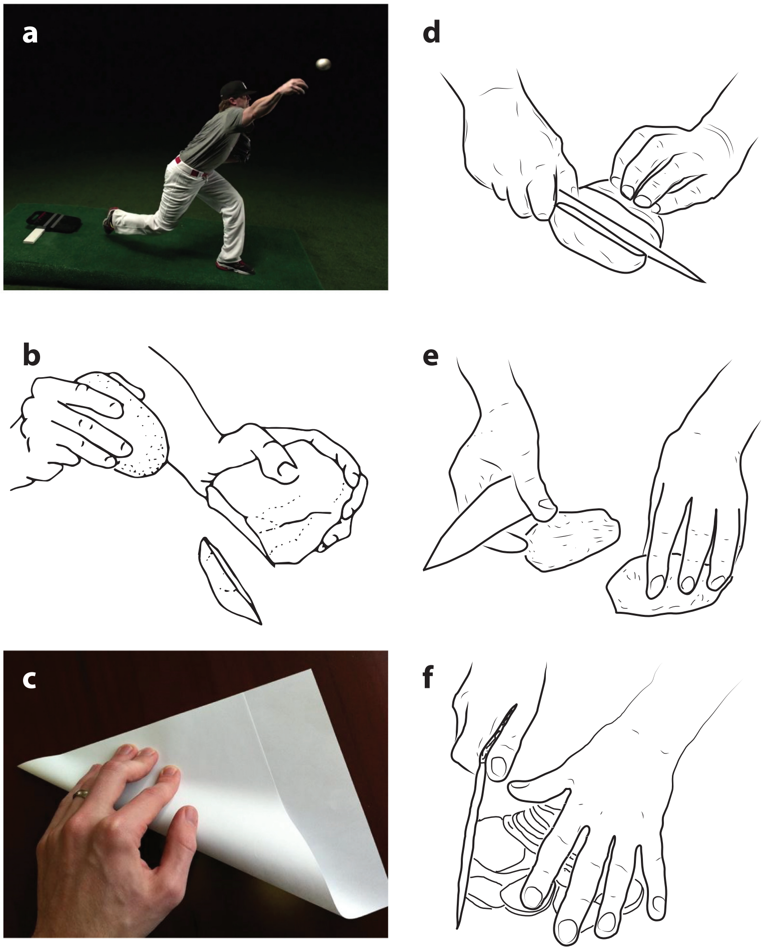

Figure 2

Examples of human manipulation. (*a*) Throwing a baseball. (*b*) Knapping a stone tool. (*c*) Folding origami. (*d*) Cutting a potato. (*e*) Bimanual manipulation of a potato while the knife is parked in an ulnar grasp. (*f*) Pushing potato slices with a knife and spread fingers. Panel *a* from video (**https://youtu.be/jZKvJY6gDfg**) by Power Drive Performance (**http://www.pitcherspowerdrive.com**), reproduced with permission. Panel *b* by Helen Beare (**https://australianmuseum.net.au/image/stone-tools-initial-reduction-flaking**), reproduced with permission from the Australian Museum. Panel *c* from video by YouTube user kiwiwhispers ASMR (**https://youtu.be/SNfLEnnP6Nc**), reproduced with permission. Panels *d*–*f* adapted from frames of *The French Chef* (28).

(34, p. 107) concluded that, while stone tool fabrication is possible with power grips alone, the evolution of forceful precision grips is "essential to habitual and effective stone tool manipulation" during fabrication.

Humans also manipulate objects without grasping them, referred to as nonprehensile manipulation. Surgery (**Figure 3*a***) is a prominent application of telerobotics and provides numerous examples of manipulation of soft materials, often without grasping or with only simple rigid grasps of the surgical instrument.

To observe more examples of nonprehensile manipulation, you might examine an instructional cooking video frame by frame. **Figure 2*d–f*** shows three frames taken from a 14-second clip of *The French Chef* (28) in which the human (Julia Child) manipulates more than 40 distinct objects—knife, potato, 2 potato halves, and 40 potato slices—and yet grasps just one object: the knife.

This simple clip challenges some ideas and inspires others. Ms. Child never puts the knife down. Rather, she shifts the knife back and forth, parking it at the ulnar side of the hand when she needs to use her thumb and index and middle fingers for other tasks. Her manipulation of the potato and potato halves employs the work surface. You might even view the work surface as an extension of the agent, an idea sometimes called whole-body manipulation or whole-world manipulation.

4. ROBOTIC AND TELEROBOTIC MANIPULATION

This section begins with a quick sketch of robotic and telerobotic manipulation and then surveys the main concepts and principles.

Robotic manipulation is making inroads in manufacturing, logistics, health care, and other sectors. A thorough survey might not be possible, but an excellent place to start is the *Springer Handbook of Robotics* (35). The present article is limited to a few examples chosen to illustrate a broad picture and support some comparisons with biological manipulation.

4.1. History

Three historical threads began somewhat independently and then came together to define robotic manipulation as we know it (36). The earliest thread to consider is teleoperators, which evolved from simple tongs to the first remotely operated master–slave manipulators in 1945. The original motive for teleoperated manipulators was to handle nuclear materials with human operators located safely behind a radiation-absorbing barrier (37, 38). The late stages of that evolution, including a transition from purely mechanical pantograph-style devices to electrically coupled fly-by-wire devices, were described in several papers by Goertz and colleagues (39–43).

Telerobotic manipulation has continued to evolve (44) with additional applications, most notably in underwater vehicles, space vehicles, bomb disposal, and surgery. Surgery and bomb disposal provide compelling practical applications of telerobotics. In particular, bomb disposal robots have saved thousands of lives (45).

Present teleoperation technology is well represented by teleoperated surgical systems (**Figure 3*a***) and bomb disposal robots (**Figure 3*b***). Teleoperated systems typically combine human perception and intelligence with robotic hardware. With this arrangement, we can explore the capabilities of one component, such as an effector, unhindered by limitations of other components, such as perception and intelligence. Other arrangements are possible, even using a human as the end effector on a telerobotic arm (46). In surgical teleoperation, the human–robot interface and thorough instrumentation provide an ideal context for the study of surgical techniques (47, 48).

Figure 3

Examples of robotic and telerobotic manipulation. (*a*) Teleoperated surgery. (*b*) Teleoperated bomb disposal robot. (*c*) Articulated serial-chain industrial robot. (*d*) Line of Delta-type parallel industrial robots. (*e*) Shakey. Panel *a* © 2018 Intuitive Surgical Inc., reproduced with permission. Panel *b* by Wikimedia user JePe (**https://commons.wikimedia.org/wiki/File:Explosieven_Opruimingsdienst.jpg**). Panel *c* by KUKA Roboter GmbH (**https://commons.wikimedia.org/wiki/File:Automation_of_foundry_with_robot.jpg**). Panel *d* by Wikimedia user Humanrobo (**https://commons.wikimedia.org/wiki/File:TOSY_Parallel_Robot.JPG**). Panel *e* by SRI International (**https://commons.wikimedia.org/wiki/File:SRI_Shakey_with_callouts.jpg**). Panels *b*–*e* reproduced under the Creative Commons Attribution-ShareAlike 3.0 Unported license (**https://creativecommons.org/licenses/by-sa/3.0**).

The second thread is industrial robotics, which started when advances in electronics and computing led to numerical control of machine tools, such as lathes and mills, inspiring George Devol's invention of a numerically controlled manipulator (49). An oft-repeated account places the beginning of industrial robotics at the 1956 meeting of George Devol and Joe Engelberger (50), leading to the founding of Unimation and the first industrial robot, the Unimate.

Industrial robotics has grown over the years to exceed annual sales of US$10 billion (51), and that figure counts just the arms. If you include the additional components required for a complete robotic system and the cost of integrating that system into a factory, the cost is several times higher. Industrial robots also have many applications outside of factory automation, including artistic applications (52, 53), movement of cameras, and movement of medical imaging and radiation devices (54).

The third thread is academic robotics research, starting in the 1960s, which developed kinematic models, dynamic models, controls, programming systems, and automatic planning for manipulators (55–57). Academic research has played a role in the growth of industrial robotics and telerobotics described above but has expanded to a much broader scope, including manipulation at micro- and nanoscales (58), swarm manipulation (59, 60), and manipulation in anthropic environments with mobile manipulators (61–65).

Over the years there have been several notable demonstrations, platforms, and challenges that help to convey the progress of the field. Early demonstrations of manipulation included systems that manipulated blocks (66, 67). Often neglected in surveys of manipulation are mobile manipulation platforms without anthropomorphic manipulators, including the milestone AI platform Shakey (68) (**Figure 3*e***), the RoboCup soccer competitions (69), and the mobile manipulation platforms developed for anthropic environments (61, 63–65). Recent challenges and demonstrations include the Defense Advanced Research Projects Agency (DARPA) Autonomous Robotic Manipulation (ARM) program (70), the DARPA Robotics Challenge (71), and the Amazon Picking Challenge (72), each of which seems to convey the same message: Robotic manipulation has come a long way, yet still has a long way to go.

4.2. Elements of Robotic Manipulation

This section is a narrative, starting with key capabilities, putting them together to achieve higher-level capabilities, and proceeding to key research issues and unsolved problems.

4.2.1. Programmed motion. Programmed motion is a key capability with many uses both inside and outside of manipulation. We define it as the ability to impart a specified motion to any load bolted to the robot's mechanical interface, with limits on the workspace, speed, acceleration, and load weight and with specified repeatability and accuracy. Programmed motion is the primary function of industrial robots. It rests on developments in motors, transmissions, encoders, kinematics, mechanism design, dynamic modeling, and control, which together constitute a substantial fraction of the literature on robotic manipulation.

Under the assumption that the load is a rigid body, programmed motion requires a robot with at least six degrees of freedom (DOFs). An arm with six DOFs is sometimes called general, and an arm with more than six DOFs is sometimes called redundant. Despite the terminology, redundant freedoms are sometimes useful to extend the workspace, avoid clutter, or provide fine motions, augmenting the coarse motions provided by the arm joints (73). Many tasks, such as moving pieces from one horizontal surface to another, or some product assembly tasks, can be accomplished with fewer than six DOFs, and arm designs such as the Selective Compliance Assembly Robot Arm (SCARA) (74, 75) and Delta (76) types exhibit improved performance for such tasks.

Originally, industrial robots were developed for the task of tending other machines in factories, for example, moving workpieces in and out of ovens (77). Welding and spray painting are also ideally suited to programmed motion. Spray painting is an especially compelling example. Automated path planning software can use a geometric model of an auto body and a process model of paint deposition to produce efficient paths for a paint sprayer (78). The path planning software can communicate the path to a robot, and the robot can execute the path repeatably and precisely, exceeding human performance.

4.2.2. Compliant motion. Manipulation usually implies contact. For example, suppose a robot is writing on a blackboard with a piece of chalk. The programmed motion idea is unlikely to perform well. If the board is slightly closer than expected, the chalk will break. If the board is slightly farther away than expected, the chalk will not even make contact, and the robot will write in air. It would be more sensible to control force, rather than position, along the direction normal to the plane of the blackboard.

So, in some tasks, it makes sense to divide the freedoms into two sets—the freedoms to be determined by the task and the freedoms to be determined by the robot—and handle these two sets differently. The compliant motion problem was first discussed in the context of teleoperation (40) and motivated the implementation of compliant control even in very early robotic systems (79). A geometrical perspective was articulated in my own first paper (80) and further refined and implemented by Raibert & Craig (81) in a framework called hybrid position/force control. At about the same time, insights from dynamic systems modeling were applied by Hogan (82) to obtain impedance control. Continued refinement of these ideas has led to capable compliant control systems (83, 84). For a more thorough history and overview, see Reference 85.

Compliant control is not the only way to address the compliant motion problem. There are many possible sources of compliance. Imagine that a robot moves to grasp a block on a table, but owing to imprecise estimates of the block's pose, the table's pose, and the robot's control, the gripper makes contact prematurely. Compliance will occur, whether graceful or not. Possible sources of compliance include the following:

- Compliant control (described above).
- The motion of the block (sliding, tipping, or deformation).
- Deformation of manipulator links. Robot manufacturers strive for stiff mechanisms to obtain precise motions, but some deformation is unavoidable.
- Passive compliance at arm joints. Robot controllers are limited in the torques they produce when servoing. Sometimes springs are introduced in series with actuators (86).
- Passive compliance of gripper mechanism. Some grippers are built with unactuated passive motion freedoms (87, 88).
- A compliant wrist. The best-known example of an engineered passive compliance is the remote center compliance (RCC) wrist described in Section 4.2.5.

4.2.3. Structured pick-and-place manipulation. Many applications involve moving a sequence of objects one at time from one place to another. If the objects are identical and the motion is repetitive, gripper design and motion programming can be addressed offline by human engineers. The motions do not have to be perfectly identical. Simple variations in output pose are easily incorporated, as when loosely packing items into a box. Variations in input pose are also possible, provided the pose can be determined accurately enough, typically with a vision system.

The approach depends on the ability to engineer the task environment—eliminating clutter and arranging for all work pieces to be in predictable positions. A task environment that has been engineered to simplify a task is called a structured environment.

The approach also depends on engineers' ability to design grippers for the item in question. There are many types of grippers. Grippers resembling tongs [sometimes called impactive (89) grippers] often have fingertips shaped to match the workpiece and are usually actuated by a fast and simple pneumatic cylinder (90). According to Monkman (91), vacuum cups were the second-most-common type 20 years ago, but they have boomed since that report (G. Monkman, private communication) and may well be in first place now.

4.2.4. Unstructured pick-and-place manipulation. Structured pick-and-place leverages the intelligence and expertise of human engineers to solve gripper design and motion planning problems. The limitations of structured pick-and-place manipulation are obvious: It depends on having a single object or very few objects and on highly repetitive motions. To eliminate these limitations, on top of a programmable arm, we need four more elements:

1. Path planning software to produce arm motions.
2. Grippers that can handle a broad range of objects.
3. Grasp pose planning to determine the stable grasp poses for a given gripper and object. (We will avoid the ambiguous phrase "grasp planning" and use either "grasp pose planning" or "grasp process planning" for clarity.)
4. Stable placement pose planning, if the robot needs to put an item down somewhere other than a given goal.

4.2.4.1. Path planning. Suppose that a robot has just grasped a block on a table and wishes to move it to a different spot on the table. How might the robot choose a motion that departs gracefully from the initial pose and moves to the final pose, without attempting to pass through the table or crashing into any other obstacles that might be present?

Progress in path planning is marked by two milestones. The first was the introduction of configuration space by Lozano-Pérez & Wesley (92). Configuration space is a construction well known in applied mechanics (93). For an n-DOF robot arm, a configuration is an n-tuple $\mathbf{q} = \{q_1, \ldots, q_n\}$, specifying the position of each joint of the robot. Then the motion of the arm corresponds to a curve in the configuration space. The free space is the subset of the configuration space comprising all configurations $\mathbf{q}$ that do not place the arm in an obstacle. Then the path planning problem is equivalent to finding a curve in the free space that connects the start configuration with the goal configuration.

The second milestone was the introduction of sampling-based planning by Kavraki et al. (94). Sampling-based planning flips the logic of path planning: Rather than identifying all free configurations and searching in the free space, the idea is to generate paths in the configuration space and test for collisions as the path is generated. There are numerous variations and refinements through which sampling-based planning has proven to be practical and efficient for many problems. For a more detailed survey of path planning, see Reference 95.

4.2.4.2. General-purpose grippers. A general-purpose robot would need a general-purpose effector. If we assume pick-and-place manipulation, we need an effector that grasps—a gripper. (In Section 4.2.7, we will flip this assumption and consider nonprehensile manipulation.) Alternatively, we need a set of grippers and a tool-changing mechanism, so that the robot can select and mount the best gripper for the immediate need.

One idea is to copy the human hand—anthropomorphism. The human hand has the broadest scope of any known effector. Besides, if the robot has an anthropomorphic gripper, it can share our tools. There is a long history of anthropomorphic designs, starting with prosthetics and continuing into teleoperation and robotics (88, 96–102). Because of the complexity of the human hand, many designs have reduced the number of fingers from five to four or three and reduced the number of motors, sometimes even down to a single motor, by coupling multiple motion freedoms to a single motor—underactuation.

As an alternative, rather than copying human hands, we might copy human tools—an approach we might call toolomorphism. Consider the common vise, for example: Two parallel plates can hold a wide variety of objects. Many robotic grippers are similar in design (77, 89).

4.2.4.3. Grasp and placement pose planning. The first question is, what would be a good grasp pose? Given a model of the object shape, what would be a good set of finger locations? The most common criterion used is force closure, meaning that the grasp contact forces can absorb any arbitrary combination of disturbance forces, subject, of course, to several assumptions. A variety of metrics have been defined to support optimization of grasps, which might be characterized as measures of how close the grasp comes to a failure of force closure or as the maximum finger forces required to balance an assumed set of disturbance forces. There is a substantial literature on grasp analysis and synthesis; for a sample, see References 103–112.

Placement pose planning is required when the robot needs to put an item down in an intermediate pose. The robot may need to shift its grasp, reorient the object, or just move the object out of the way. The problem is usually simpler than grasp planning, requiring only a static stability analysis.

4.2.4.4. Putting it all together. With these five pieces—a programmable arm, a gripper capable of grasping the desired objects, a path planner, a grasp pose planner system, and a placement pose planner—one can build a pick-and-place system in an unstructured environment. Experimental pick-and-place systems have been tested, and there are even a nice set of theoretical results addressing pick-and-place manipulation (113–115).

4.2.5. Mechanical assembly and task mechanics. In the simplest cases, manipulation can be modeled as just a kinematic process, and only path planning is required. For example, consider a classic laboratory demonstration: pick-and-place manipulation of an identical set of blocks using a parallel-jaw gripper. Grasp planning is trivial—align and center the gripper on the block, and squeeze. Placement planning is also trivial—put the block anywhere on the table.

However, in general, such a simple approach will not work, and the robot must consider a more detailed model of the task. The best example is the peg-in-hole problem (116, 117). Even if the peg fits into the hole with room to spare, and even if the robot is producing a motion consistent with the contacts between the peg and hole, it is still possible for the peg to get stuck owing to the interaction of frictional forces. By modeling the frictional contacts that can occur, one can choose a suitable compliant control or even design a passively compliant wrist, the RCC, which provides a compliant behavior ideal for the peg-in-hole problem (117, 118).

Even in the simplest cases, task mechanics can be an issue. Upon careful examination, our simple example of squeezing a block with a parallel-jaw gripper turns out to be not so simple. Intuitively, we expect that squeezing a block between two planes will align the block squarely with the gripper, but it can get stuck, cocked between the two fingers, owing to unforeseen interactions of frictional forces. Task mechanics plays an important role in almost every task, and a considerable amount of work is focused on modeling and simulating the dynamic processes of manipulation.

4.2.6. In-hand manipulation. Pick-and-place adopts a narrow view of the hand. It assumes that the hand grasps an object rigidly. In this section, we consider the motion of a grasped object in the hand, and in the next section we consider the motion of an object that is not even grasped.

Early experiments with pick-and-place systems (67, 113) revealed a serious shortcoming: the need for regrasp. Consider the act of picking up a pencil. To write with the pencil, you need to wrap your fingers around it, near the point. But when you pick up the pencil from a table, you cannot grasp it that way. The table is in the way, and the preferred grasp is not available. Humans are untroubled by this because they can quickly shift the pencil after picking it up. A pick-and-place system must employ a different strategy—for example, it might put the pencil back down, with the point projecting beyond of the edge of the table, and then grasp it a second time using the preferred grasp.

Regrasp imposes a performance hit that is often unacceptable. Factory automation engineers avoid regrasps whenever possible. Eliminating a regrasp can halve a cycle time. But regrasp is not just a factory automation issue. Humans adjust grasp poses frequently and efficiently without putting the object down for a regrasp. Doing everything with regrasps would be incredibly tedious, at best. In-hand manipulation is also useful to produce fine motions for mechanical assembly or other compliance tasks.

Roboticists have identified several ways to produce in-hand manipulation. The most studied technique is to design a hand with enough motion freedoms to wiggle the fingers while maintaining fixed or rolling contact (119). Salisbury's "dexterous hand" (97, 120) was an example. It is even possible to obtain large rotations in some cases. In fact, global controllability of object orientation is possible, subject to assumptions on the shape of the object, the mobility of the gripper, and the nature of the contacts (121). Other techniques include finger-gaiting, controlled slip in either translation or rotation, and using the object's momentum to throw it from one grasp pose to another (122–128).

4.2.7. Nonprehensile manipulation. The role of the hand is a central issue in robotic manipulation. We started by assuming rigid grasps. We then relaxed this assumption to consider motion in the hand. We now take one further step and consider motion of ungrasped objects: nonprehensile manipulation, or graspless manipulation.

We have already seen examples of nonprehensile manipulation. **Figure 2*d–f*** depicts a sequence of manipulation behaviors involving almost no grasping at all. Most of the manipulation uses pushing by either the fingers or the knife, a technique that has frequently occupied this author's attention and that of many others (129–132).

Nonprehensile manipulation is an essential part of the grasp process. Humans often move objects as preparation for grasping them (133). Nonprehensile motion is often unavoidable during grasping and sometimes can be designed to eliminate pose uncertainty and produce robust grasps (129, 134). Nonprehensile manipulation is also important in dealing with clutter (131). Some experimental systems have been demonstrated using only nonprehensile manipulation (135), including mobile manipulators, such as Shakey (68) (**Figure 3*e***).

Nonprehensile manipulation is also essential for parts-orienting systems, which are vital in factory automation. It is common for parts to arrive in a disorganized jumble. In some cases, a vision system suffices to treat the problem as a pick-and-place problem. More often, parts-orienting and parts-feeding machinery is required to present the part in a repeatable position. Some nonprehensile techniques that have been demonstrated include a tilting rectangular tray (136), a sequence of squeezes by a parallel-jaw gripper (134), a sequence of obstacles suspended above a conveyor (137), grasping with a simple three-jaw gripper (138), and, in theory, orienting parts by dropping them through a cleverly designed shape under the action of gravity alone (139).

4.2.8. Whole-X manipulation. We often assume that manipulation occurs through the interaction of an object with the fingers of a hand. But humans use other parts of their bodies: elbowing a door open, using a hip to hold it open, tucking things under the elbow or under the chin. From the perspective of a kinematic and dynamic model, these are all parts of the agent's body that can be controlled to manipulate the external world. This perspective was originally called whole-arm manipulation (140) and has come to be known also as whole-body manipulation.

We can take the term a step further: whole-world manipulation. Consider the potato-cutting operation of **Figure 2*d–f***. Just before the cut, the potato is held between four fingers of the hand and the table. The table is playing the same role as one of the fingers, and the techniques for analyzing stable grasping and in-hand manipulation are readily adapted to the situation.

5. COMPARING ANIMALS AND ROBOTS

5.1. Breadth Versus Specialization

The most obvious difference between animals and robots is the breadth and robustness of animals versus the specialization and brittleness of robots. Humans exhibit an astonishing breadth and variety of behaviors—cooking, medical procedures, sculpting, juggling, musical performance, personal grooming, caressing . . . a seemingly endless list.

Animal and human manipulative behaviors are robust and adaptable, employing their superior perception and intelligence to adjust quickly to minor variations in a task, and in some cases even to make substantial changes as appropriate. Robots, on the other hand, often work in carefully engineered and controlled environments—the structured environments described in Section 4.2.3.

Robots are highly specialized, as are many machines. The degree of a robot's specialization is often difficult to appreciate when the limitations are behavioral rather than morphological. When you look at a sewing machine, you can see that it is only going to sew, and, further, that it is going to sew only if the work is presented in a highly specific manner. But when you look at an industrial robot or a humanoid, you imagine that it is capable of a broad range of activities. Unfortunately, that is not truly the case, because of the limitations in perception and intelligence. Engineers are limited in the decisions that can be delegated to online software. The behaviors are so narrowly defined that seemingly minor variations in the task will defeat the robot—the behavior is brittle.

The other side of the coin is the performance advantages that come with specialization. Compliant structures, such as the SCARA robot and the RCC (described above), can be selected for a specific task. A robot can be selected and configured to make a select class of motions at high speed. We lament the brittleness, but we profit, literally, from the performance.

Sometimes animals and humans are at a disadvantage in speed and precision, but these comparisons might not be fair since they often involve robots that are highly specialized. When we consider a task for which humans are at least somewhat specialized, such as throwing, the >100-g acceleration cited above comes to mind. Further, deficiencies in open-loop speed and precision may be offset by the dexterity with which a human can accomplish very-high-precision tasks.

There are many other differences between animals and robots. One surprising behavioral advantage is the speed with which a human can shift grasps, which enables very fast tool changing. For example, the shifting of the knife in **Figure 2*d–f*** happens in a fraction of a second. This behavior is reminiscent of Marzke's hypothesis that fine manipulation evolved in part to support agile shifting of grasps during toolmaking (34).

5.2. Perception, Intelligence, and Communication

Humans have the advantage in visual perception, although robotics is making great strides. Robots also make good use of force/torque sensors and contact sensors. An especially notable advantage of humans and animals is tactile perception. A human can inspect a gym bag for a gun in a few moments or pull a coin from a pocket while relying largely on tactile perception.

Humans have an enormous advantage in intelligence, but there are gaps. Some difficulties with spatial reasoning can be surprising. Determining whether two shapes are actually the same shape in different orientations can be surprisingly slow (141), and we have all enjoyed the frustrating difficulty of mechanical puzzles, such as the alpha puzzle. Even though the alpha puzzle involves only a short path of one rigid object relative to another, it can take seconds or even minutes to solve.

There are enormous differences in communication between humans and robots. Above, I mentioned the spray-painting example, where a motion can be communicated easily and precisely to a robot. In that case, and in many cases, the robot is more biddable than a human, but in the broader sense the human is more biddable simply because of superior intelligence.

5.3. Soft Things

Humans manipulate soft things with ease, and while robots have reproduced a few such behaviors, such as folding origami (142), suturing (47), and folding laundry (143), there is a large gap between robot and human performance. Modeling and controlling soft things presents many challenges. The typical assumptions of robotic manipulation do not apply. For example, when tying one's shoelaces, what is the goal? It certainly is not a specific precisely determined pose of the laces.

5.4. Expense

The final and most obvious difference is expense. Humans come in discrete units. With robots, you can buy a single arm, or even buy an arm with just two or three joints. With robots, you pay for what you need. With humans, you pay for an entire unit, regardless of need.

6. EVOLUTION VERSUS INTELLIGENT DESIGN

It is the nature of robotics to compare robots with the animals that inspired them, that is, to compare engineers' creations with evolution's creations. Instead, in this section, I compare the creative processes directly.

As a foundation for the comparison, let us create a timeline for every choice involved in an agent's actions. For example, consider the choices behind a factory robot's reach and grasp action (**Figure 4*a***). It is not practical to list all the choices, but they would include the robot morphology—the arm and effector design—as well as behavioral choices, such as what motion to attempt and what motor torques to apply. Teasing decisions apart takes some care. For example, at first glance it seems that the choice of motor torque occurs online as frequently as every millisecond, but actually the torque was largely determined offline, when the controller was designed. The online process only puts the finishing touches on the choice, based on sensory feedback. Likewise, the choice of motion is largely offline, although it may vary a bit, based on online sensory data. In many cases, the choice of motion is made entirely offline.

Now consider a similar timeline for a human (**Figure 4*b***). To keep things simple, let us neglect culture and focus on just an isolated individual. Suppose a human throws a rock at a bird, and consider all the decisions, behavioral and morphological, behind that act. The big decisions include bipedalism, the high mobility of the torso and of the wrist, the choice of stance and target, the neural structures behind muscle programs and muscle control, and the decision that the human should hunt with stones rather than claws.

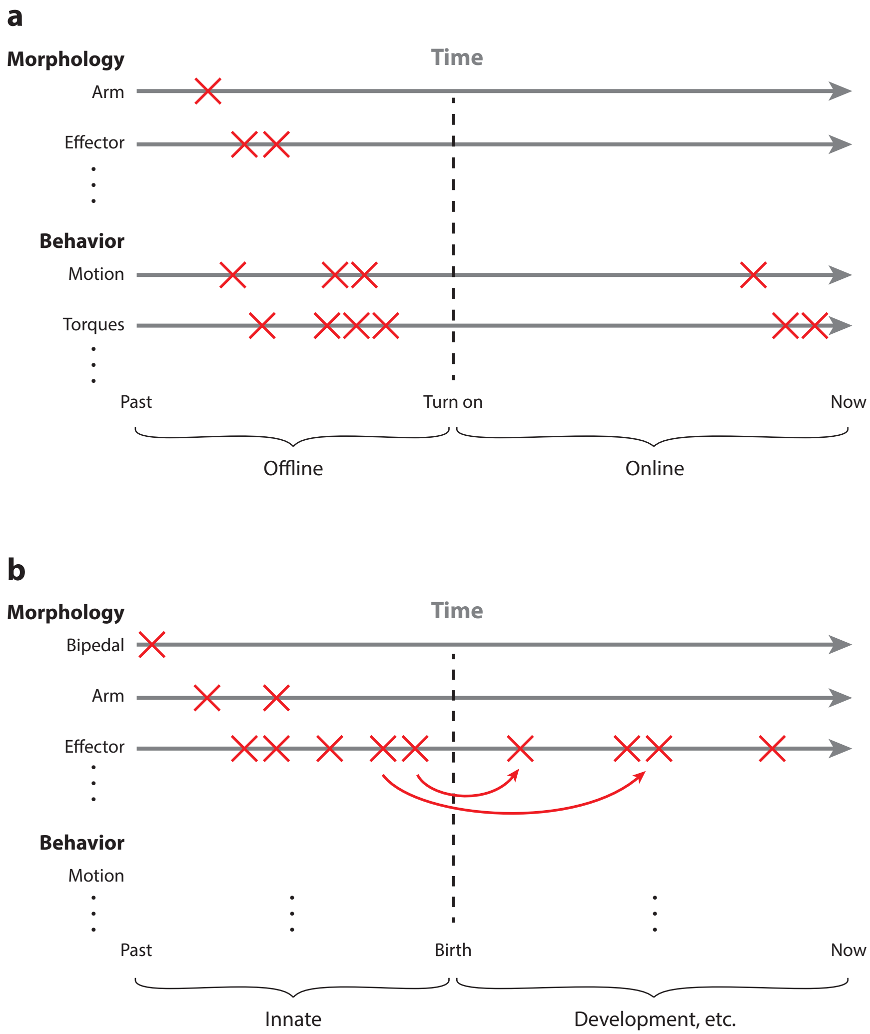

Figure 4

Timelines for all choices involved in an agent's actions. (*a*) Factory robot. (*b*) Human.

Some choices are innate, occurring through the process of evolution before the individual is born, such as bipedalism. Other decisions are in real time, even as the throw is in progress, such as aborting the throw if the bird takes flight. Still others occur earlier in the individual's lifetime, such as the morphological development that occurs through physical training and the skill development that occurs through learning.

Now, consider the choice of whether to equip the human body with claws or some different weapon. Evolution chooses not to choose. Evolution delegates the decision to the individual human. Each of us has the choice of hunting with a knife (an artificial claw) or some other weapon.

Thus, the human hand is less adapted or specialized than some other animals' effectors. Anthropologists say that the human hand is comparatively primitive (144). Sometimes this choice is described in negative terms: Nature has given the human neither armor nor weapons. But, to view the situation in a more positive light, nature has delegated the choice to us, so that we can adapt more intelligently, according to the needs of the moment. When cutting is called for, we pick up a knife. When throwing is called for, we pick up a stone.

Engineers have the same choice when choosing a robot effector. Offline, they can determine the single best effector for the job. Or they can equip the robot with a tool-changing mechanism and software to change tools as appropriate. However, given the present state of robotics, the engineer must make almost all the decisions. Engineers are considerably smarter than their creations.

On the other hand, suppose the task to be performed, or the variation in that task, strays beyond what evolution has delegated to the individual. Even if evolution can adapt to these changes, it will be slow. Engineers can respond more quickly, adapting to new needs in days rather than eons.

In short, the strength of robotics is the speed and consequent adaptability of the creative process, as well as the greater performance enabled by building highly specialized machines.

Robots have another advantage. Evolution has the option of making the choice offline or delegating the choice to the individual. Engineers have a third option: delegation to themselves. Unlike evolution, an engineer has the option to meddle with a robot later on. **Figure 4** shows a single "turn on" event in the life of a robot, but the reality is quite different. The engineer can deploy the robot in one configuration, then turn it off, revise the morphology and behavior, and turn it on again.

Engineers go quite a bit further in the direction of modular robot mechanisms. Automation technology is highly modular. Roboticists are also exploring self-reconfigurable modular mechanisms: robots that autonomously disassemble and reassemble to create the mechanism suited to the present need (145).

7. MANIPULATION IS HARD

Manipulation is hard. Robotics in general is hard. If you ask a roboticist to state the number-one lesson we have learned from robotics, the likely answer is "Robotics is hard!" We have a name for this: Moravec's paradox (146). The usual example is chess. Chess was one of the great challenge problems of AI. AI researchers eventually developed world-champion-level chess players, except that the computers still need human beings to do the actual moving of the chess pieces. Yes, robots can move chess pieces, but not nearly as well as humans. Championship chess, attained by only the most gifted minds, is actually easier than moving the pieces, which almost every one of us can easily do.

7.1. Why Is Manipulation So Hard?

We should not be surprised by the difficulty of manipulation. The creation of autonomous robotic manipulation is surely one of the most challenging engineering problems we face. It encompasses many difficult problems, some of which are described below.

7.1.1. Mechanisms. We want devices that can apply enough force to grip and move significant loads, have sufficient motion freedoms, can move precisely but also quickly, and are compact enough to work in tight spaces. Further, all the surfaces of the manipulator should be sensitive to touch and serve as effector surfaces as needed, and many of these surfaces should be soft. Finally, they need to wear well, be easily replaced, and be inexpensive.

7.1.2. Perception. We need perception to provide a broad appreciation of the scene but also to provide high-resolution information, including knowledge of contact locations and forces exchanged. Research systems (especially machine learning) benefit from copious amounts of precise ground truth. Vision cannot provide sufficient resolution in spatial data or force data. The most critical need for high-resolution data is when contact is imminent, at which point occlusion is inevitable. That shortcoming would be mitigated to some extent by tactile sensing, but neither the devices nor the methods are ready to fill the gap.

7.1.3. Modeling and control. We need models of the manipulation process so that we can analyze, simulate, plan, and control our manipulation systems. Many of the underlying phenomena are challenging—unilateral contacts, frictional contacts, impact, and deformation—even when considered individually. Further, the workpieces are themselves system components that are not known in advance, with shapes, materials, and mass distributions that will not be known precisely even at run time. We need controllers for mechanisms that are seriously underspecified, using feedback signals that are variable in quality and sometimes nonexistent.

7.1.4. Planning. Several results have demonstrated the computational intractability of manipulation planning, starting with a paper by Reif (147). The reduction of manipulation problems to well-known intractable computational problems is straightforward, or even trivial in some cases, such as the bin-packing problem (148). Fortunately, intractability results are not the final word if one is willing to accept algorithms that produce less-than-optimal outcomes in some cases. Nonetheless, intractability results do confirm the fundamental underlying complexity of the problem. To get an intuitive notion of the daunting combinatorics of manipulation problems, it is only necessary to count the number of orderings of operations when cleaning the kitchen, or to count the number of different functions from m sensory signals to n motor signals, or, addressing planning for uncertainty, to count the number of probability density functions over an n-dimensional state space.

The daunting combinatorics, combined with theoretical intractability results, suggest that developing the best technique for a given task can be a lengthy process. Human experience confirms it. When a human performs a manipulation task quickly, it is by drawing on an enormous repertoire of behaviors, developed by evolution and passed to us genetically, or developed by our ancestors and passed to us by demonstration and tutorial. An attempt at a task outside that repertoire might take a long time or even fail.

7.1.5. Uncertainty. Uncertainty presents challenges that go beyond the combinatorics issue mentioned above. Goertz (39) was the first to note the importance of uncertainty. He presented two examples that have been prominent in manipulation research ever since:

- Turning a crank. The motion required to turn a crank is simple: a circle. The problem is, where exactly should the robot produce the circle? We can do our best to estimate the crank's position, but we will never get it exactly right. If the robot and the crank disagree, then the result can be large forces, failed grasps, or worse. We need compliant motion: the use of controls and/or mechanisms that allow us to specify motion while complying with contact constraints.
- Peg-in-hole. Putting a peg into a hole requires a precision comparable to the difference in their diameters, which can be extremely small—smaller than we are likely to obtain with computer vision, and also smaller than the positioning accuracy of our robot. We need

planning systems that can reason out a compliant motion that uses contact to guide the peg into the hole.

Uncertainty also affects grasping. We need grasping motions that are robust to the inevitable errors in object pose and robot motion.

7.2. The Thinker and the Doer

How do we explain Moravec's paradox? Given the various challenges detailed in the previous section, one might suppose that robotic manipulation is one of the greatest engineering challenges ever. Why would we be surprised to discover that manipulation is hard?

Imagine a clean division between a conscious thinking part of the brain and a subconscious manipulating part of the brain. Let us call them the thinker and the doer. It is the thinker who chooses the chess moves, and the doer who actually moves the pieces. The thinker also observes and discusses. A thinker wrote this article, and another thinker is reading it.

It is the thinker who studies and theorizes about manipulation. Unfortunately, the thinker cannot directly observe the inner mechanisms of the doer. The thinker's usual experience is that it need only think of a behavior, and the behavior happens. Subjectively, manipulation is effortless, and there is no reason for the thinker to doubt its understanding of manipulation.

How does a thinker achieve an objective perspective on manipulation and an assessment of its understanding? Feynman famously wrote, "What I cannot create, I do not understand" (149, p. 83). Feynman would suggest that the limitations of our robots reflect the limitations of our understanding. Robotics is the objective test of our understanding. We can claim to understand certain stylized types of manipulation: structured pick-and-place manipulation, for example. Of the vast repertoire of manipulation behaviors exhibited by humans, we understand little.

8. DIRECTIONS

8.1. Machine Learning

Comparing the robots' repertoire of behaviors with the animals', and considering the rate at which we expand the robots' repertoire, it seems that robotic manipulation will lag biological manipulation for a very long time. Clearly we need to address the vast repertoire collectively, not one behavior at a time.

One approach to expanding the breadth of robotics is machine learning. Interest in machine learning is deep and long-standing, sustained by a continuing stream of exciting results, such as machines that learn to fold laundry (143) and grasp (150), and reinforced by a long history of related techniques, including system identification, adaptive control, stochastic learning automata, and iterative learning control. You might view evolution as an existence proof that machine learning can solve robotics, and in fact evolution is the direct inspiration for one type of machine learning: genetic algorithms (151). On the other hand, you might view machine learning as a threat, if you think it might produce the robotic technology we want without advancing our understanding of manipulation, and render irrelevant thousands of papers on the principles and engineering of robotic manipulation.

The more likely scenario is that machine learning provides techniques and ideas that accelerate our understanding as well as the technology. Fears are misplaced. Even if machine learning does deliver great technology without advancing our understanding, it would give us a technology entirely different from biotechnology, which would be an invaluable new class of subjects to study.

8.2. Measuring and Comparing

While it seems obvious that a human's manipulation skills are superior to an ant's, such a comparison is challenging. Ants do things that humans cannot, and humans do things that ants cannot. Perhaps ants pity humans for our poor nest-building skills.

We do have ways to compare industrial robots: workspace, speed, maximum load, and repeatability. Likewise, we have ways to compare humans: baseball-playing ability, number of clubs one can juggle, and so on. But all of these are intraspecies comparisons; none of them suffice to compare ants with humans or robots. Ants, humans, and robots do not play the same games.

There is one game that both robots and humans play: tapping. Psychologists use tapping to explore human capabilities (152). Subjects move back and forth between two points, tapping each point as accurately as possible. Performance is measured by the separation distance, accuracy, and speed. Robots play a similar game. Some robots are compared in terms of a pick-and-place cycle, which is to move 25 mm up, then 300 mm sideways, then 25 mm down, back and forth. A 1-second cycle time was a very good time for articulated or SCARA robots in the 1980s (153). Delta-type robots are considerably faster, with repeatability better than 0.1 mm (154).

Another approach is to examine the component resources of the agent: number of DOFs, amount of muscle mass, number of neurons, and so on. Counting neurons or synapses might seem particularly relevant, but that is only one aspect of an agent's manipulation assets. There are trade-offs between morphology and intelligence. Roboticists have a variety of names for this: mechanical intelligence (99), intelligence by mechanics (155), and morphological intelligence.

We need some way of boiling morphology, mechanical structure, signal processing, and computation down to one quantity. Information-theoretic approaches seem promising, to measure either the complexity of an agent's behavior (156) or the complexity of a task (157, 158).

SUMMARY POINTS

1. Manipulation is hard.
2. Animals exhibit outstanding breadth, robustness, and adaptability.
3. Robots exhibit specialization and consequent performance advantages.
4. Engineering's chief advantage over evolution is quick adaptability based on intelligent design.
5. Robotics has led to a thorough understanding of a few styles of manipulation, which collectively constitute a small fraction of the behaviors exhibited by animals.

FUTURE ISSUES

1. Is there a fundamental and precise metric for comparing manipulative behaviors, or for comparing tasks, that would provide a basis for measuring progress in the field?
2. How can we best take advantage of advances in machine learning to advance our understanding and improve our technology?
3. How do we develop the adaptability, robustness, and breadth of behaviors exhibited by animals and humans?

DISCLOSURE STATEMENT

I have received funding from the National Science Foundation, DARPA, the US Army Research Laboratory, and Foxconn and am employed at Berkshire Grey.

ACKNOWLEDGMENTS

This article has benefited from numerous discussions with friends and colleagues. In particular, I would like to thank Dan Koditschek, Marc Raibert, Mike Erdmann, Tomás Lozano-Pérez, Oliver Brock, Nancy Pollard, Alberto Rodriguez, Kevin Lynch, Ken Goldberg, Byron Boots, Siddhartha Srinivasa, Howie Choset, Chris Atkeson, Tucker Balch, Vijay Kumar, Richard Byrne, Cynthia Gelb, and the current and past members of the Carnegie Mellon Manipulation Lab. This material is based upon work supported by the National Science Foundation under grants IIS1409003 and IIS1637908. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

LITERATURE CITED

1. euRobotics. 2013. *Robotics 2020: strategic research agenda for robotics in Europe*. SRA Doc., euRobotics, Brussels. **https://ec.europa.eu/research/industrial_technologies/pdf/robotics-ppp-roadmapn.pdf**
2. Miller D. 2015. *2015 NASA technology roadmaps: TA4: robotics and autonomous systems*. Roadmap Doc., NASA, Washington, DC. Available at **https://www.nasa.gov/offices/oct/home/roadmaps**
3. Mason MT. 2001. *Mechanics of Robotic Manipulation*. Cambridge, MA: MIT Press
4. Davidson D. 2001. *Essays on Actions and Events*. Oxford, UK: Clarendon. 2nd ed.
5. Buehler M, Koditschek DE, Kindlmann PJ. 1994. Planning and control of robotic juggling and catching tasks. *Int. J. Robot. Res.* 13:101–18
6. Rizzi AA, Koditschek DE. 1993. Further progress in robot juggling: the spatial two-juggle. In *1993 IEEE International Conference on Robotics and Automation*, Vol. 3, pp. 919–24. New York: IEEE
7. Yubuki N, Leander BS. 2013. Evolution of microtubule organizing centers across the tree of eukaryotes. *Plant J.* 75:230–44
8. Marzke M. 1983. Joint functions and grips of the *Australopithecus afarensis* hand, with special reference to the region of the capitate. *J. Hum. Evol.* 12:197–211
9. Marzke MW. 1996. Evolution of the hand and bipedality. In *Handbook of Human Symbolic Evolution*, ed. A Lock, CR Peters, pp. 126–47. Oxford, UK: Clarendon
10. Richmond BG, Roach NT, Ostrofsky KR. 2016. Evolution of the early hominin hand. In *The Evolution of the Primate Hand: Anatomical, Developmental, Functional, and Paleontological Evidence*, ed. TL Kivell, P Lemelin, BG Richmond, D Schmitt, pp. 515–43. New York: Springer
11. Williams-Hatala EM. 2016. Biomechanics of the human hand: from stone tools to computer keyboards. In *The Evolution of the Primate Hand: Anatomical, Developmental, Functional, and Paleontological Evidence*, ed. TL Kivell, P Lemelin, BG Richmond, D Schmitt, pp. 285–312. New York: Springer
12. Johanson DC, Taieb M. 1976. Plio-pleistocene hominid discoveries in Hadar, Ethiopia. *Nature* 260:293–97
13. Roach NT, Venkadesan M, Rainbow MJ, Lieberman DE. 2013. Elastic energy storage in the shoulder and the evolution of high-speed throwing in homo. *Nature* 498:483–86
14. Darwin C. 1871. *The Descent of Man, and Selection in Relation to Sex*. London: John Murray
15. Guinness. 2017. *Fastest delivery of a cricket ball (male)*. **http://www.guinnessworldrecords.com/world-records/fastest-bowl-of-a-cricket-ball**
16. Arndt J. 2016. How Aroldis Chapman threw the fastest pitch ever. *Huffington Post*, Aug. 28. **https://www.huffingtonpost.com/entry/how-aroldis-chapman-threw-the-fastest-pitch-ever_us_57c38390e4b06384eb4066ef**

17. Savage-Rumbaugh ES, Lewin R. 1994. *Kanzi: The Ape at the Brink of the Human Mind*. New York: Wiley
18. WildNet Afr. 1997. Baboons stone highway motorists. *WildNet Africa News Archive*, Jan. 8. **http://www.cs.cmu.edu/~mason/baboons08011997.pdf**
19. Isaac B. 1987. Throwing and human evolution. *Afr. Archaeol. Rev.* 5:3–17
20. Coddington J, Sobrevila C. 1987. Web manipulation and two stereotyped attack behaviors in the ogre-faced spider *Deinopis spinosus* Marx (Araneae, Deinopidae). *J. Arachnol.* 15:213–25
21. Kullman E. 1959. Beobachtungen und betrachtungen zum verhalten der theridiidae *Conopistha argyrodes*. *Mitt. Zool. Mus. Berlin* 35:275–92
22. Hölldobler B, Wilson EO. 2010. *The Leafcutter Ants: Civilization by Instinct*. New York: Norton
23. Weir AAS, Chappell J, Kacelnik A. 2002. Shaping of hooks in New Caledonian crows. *Science* 297:981
24. Teleki G. 1974. Chimpanzee subsistence technology: materials and skills. *J. Hum. Evol.* 3:575–94
25. Auersperg AMI, von Bayern AMP, Gajdon GK, Huber L, Kacelnik A. 2011. Flexibility in problem solving and tool use of kea and New Caledonian crows in a multi access box paradigm. *PLOS ONE* 6:e20231
26. Klump BC, Sugasawa S, St Clair JJH, Rutz C. 2015. Hook tool manufacture in New Caledonian crows: behavioural variation and the influence of raw materials. *BMC Biol.* 13:97
27. Byrne R, Corp N, Byrne J. 2001. Manual dexterity in the gorilla: bimanual and digit role differentiation in a natural task. *Anim. Cogn.* 4:347–61
28. Child J, host. 1963. The potato show. *The French Chef*, season 1, episode 22, aired June 22
29. Schlesinger G. 1919. Der mechanische Aufbau der künstlichen Glieder. In *Ersatzglieder und Arbeitshilfen*, ed. M Borchardt, K Hartmann, RR Leymann, G Schlesinger, H Schwiening, pp. 321–661. Berlin: Springer
30. Cutkosky M, Wright P. 1986. Modeling manufacturing grips and correlations with the design of robotic hands. In *1986 IEEE International Conference on Robotics and Automation*, Vol. 3, pp. 1533–39. New York: IEEE
31. Napier J, Tuttle RH. 1993. *Hands*. Princeton, NJ: Princeton Univ. Press
32. Feix T, Pawlik R, Schmiedmayer H, Romero J, Kragić D. 2009. *The generation of a comprehensive grasp taxonomy*. Presented at Robot. Sci. Syst. Conf., Workshop Underst. Hum. Hand Adv. Robot. Manip., Seattle, WA, June 28–July 1. **http://www.csc.kth.se/~danik/taxonomyGRASP.pdf**
33. Elliott J, Connolly K. 1984. A classification of manipulative hand movements. *Dev. Med. Child Neurol.* 26:283–96
34. Marzke M. 1997. Precision grips, hand morphology, and tools. *Am. J. Phys. Anthropol.* 102:91–110
35. Siciliano B, Khatib O, eds. 2008. *Springer Handbook of Robotics*. New York: Springer
36. Mason MT. 2012. Creation myths: the beginnings of robotics research. *IEEE Robot. Autom. Mag.* 19:72–77
37. Sheridan T. 1989. Telerobotics. *Automatica* 25:487–507
38. Hokayem PF, Spong MW. 2006. Bilateral teleoperation: an historical survey. *Automatica* 42:2035–57
39. Goertz R, Bevilacqua F. 1952. A force-reflecting positional servomechanism. *Nucleonics* 10:43–45
40. Goertz R. 1952. Fundamentals of general-purpose remote manipulators. *Nucleonics* 10:36–42
41. Goertz R, Thompson W. 1954. Electronically controlled manipulator. *Nucleonics* 12:46–47
42. Goertz R. 1954. Mechanical master-slave manipulator. *Nucleonics* 12:45–46
43. Goertz R. 1963. Manipulators used for handling radioactive materials. In *Human Factors in Technology*, ed. E Bennett, J Degan, J Spiegel, pp. 425–43. New York: McGraw-Hill
44. Niemeyer G, Preusche C, Stramigioli S, Lee D. 2016. Telerobotics. See Ref. 35, pp. 1085–108
45. Shachtman N. 2009. Military gears up for bomb-bot 2.0. *Wired*, Feb. 25. **https://www.wired.com/2009/02/new-bomb-dispos**
46. NASA. 1993. *Wisoff on the arm*. Photo, NASA, Washington, DC. **https://upload.wikimedia.org/wikipedia/commons/d/d0/Wisoff_on_the_Arm_-_GPN-2000-001069.jpg**
47. van den Berg J, Miller S, Duckworth D, Hu H, Wan A, et al. 2010. Superhuman performance of surgical tasks by robots using iterative learning from human-guided demonstrations. In *2010 IEEE International Conference on Robotics and Automation*, pp. 2074–81. New York: IEEE
48. Zappella L, Béjar B, Hager G, Vidal R. 2013. Surgical gesture classification from video and kinematic data. *Med. Image Anal.* 17:732–45

49. Devol GC Jr. 1961. *Programmed article transfer*. US Patent No. 2988237
50. Pierce J. 2011. George C. Devol, inventor of robot arm, dies at 99. *New York Times*, Aug. 15, p. B16
51. Int. Fed. Robot. 2016. *World robotics report 2016*. Rep., Int. Fed. Robot., Frankfurt am Main, Ger.
52. Byrne K, Proto J, Kruysman B, Bitterman M. 2014. The power of engineering, the invention of artists. In *Robotic Fabrication in Architecture, Art and Design*, ed. W McGee, M Ponce de Leon, pp. 399–405. Cham, Switz.: Springer
53. Goldberg K, ed. 2001. *The Robot in the Garden: Telerobotics and Telepistemology in the Age of the Internet*. Cambridge, MA: MIT Press
54. Adler JR Jr., Chang SD, Murphy MJ, Doty J, Geis P, Hancock SL. 1997. The CyberKnife: a frameless robotic system for radiosurgery. *Stereotact. Funct. Neurosurg.* 69:124–28
55. Paul RP. 1981. *Robot Manipulators: Mathematics, Programming, and Control: The Computer Control of Robot Manipulators*. Cambridge, MA: MIT Press
56. Whitney DE. 1972. The mathematics of coordinated control of prosthetic arms and manipulators. *J. Dyn. Syst. Meas. Control* 94:303–9
57. Vukobratovic M, Potkonjak V. 1982. *Dynamics of Manipulation Robots: Theory and Application*. Berlin: Springer
58. Sitti M. 2007. Microscale and nanoscale robotics systems. *IEEE Robot. Autom. Mag.* 14:53–60
59. Kube CR, Zhang H. 1993. Collective robotics: from social insects to robots. *Adapt. Behav.* 2:189–218
60. Werfel J, Petersen K, Nagpal R. 2014. Designing collective behavior in a termite-inspired robot construction team. *Science* 343:754–58
61. Khatib O, Yokoi K, Brock O, Chang K, Casal A. 1999. Robots in human environments: basic autonomous capabilities. *Int. J. Robot. Res.* 18:684–96
62. Brock O, Park J, Toussaint M. 2016. Mobility and manipulation. See Ref. 35, pp. 1007–36
63. Srinivasa SS, Ferguson D, Helfrich CJ, Berenson D, Collet A, et al. 2010. HERB: a home exploring robotic butler. *Auton. Robots* 28:5–20
64. Bohren J, Rusu RB, Jones EG, Marder-Eppstein E, Pantofaru C, et al. 2011. Towards autonomous robotic butlers: lessons learned with the PR2. In *2011 IEEE International Conference on Robotics and Automation*, pp. 5568–75. New York: IEEE
65. Borst C, Wimbock T, Schmidt F, Fuchs M, Brunner B, et al. 2009. Rollin' Justin—mobile platform with variable base. In *2009 IEEE International Conference on Robotics and Automation*, pp. 1597–98. New York: IEEE
66. Ernst HA. 1962. MH-1, a computer-operated mechanical hand. In *AIEE-IRE '62 (Spring): Proceedings of the May 1–3, 1962, Spring Joint Computer Conference*, pp. 39–51. New York: ACM
67. Paul R, Pingle K, Feldman J, Kay A. 1971. *Instant insanity - computer vision & robotics*. Film recorded at Artif. Intell. Lab., Stanford Univ., Stanford, CA. **https://www.youtube.com/watch?v=O1oJzUSlTeY**
68. Nilsson NJ, ed. 1984. *Shakey the robot*. Tech. Note 323, SRI Int., Menlo Park, CA
69. Kitano H, Asada M, Kuniyoshi Y, Noda I, Osawai E, Matsubara H. 1998. RoboCup: a challenge problem for AI and robotics. In *RoboCup-97: Robot Soccer World Cup I*, ed. H Kitano, pp. 1–19. Berlin: Springer
70. Hackett D, Pippine J, Watson A, Sullivan C, Pratt G. 2013. An overview of the DARPA Autonomous Robotic Manipulation (ARM) program. *J. Robot. Soc. Jpn.* 31:326–29
71. Pratt G, Manzo J. 2013. The DARPA robotics challenge. *IEEE Robot. Autom. Mag.* 20:10–12
72. Correll N, Bekris KE, Berenson D, Brock O, Causo A, et al. 2018. Analysis and observations from the first Amazon picking challenge. *IEEE Trans. Autom. Sci. Eng.* 15:172–88
73. Hollis RL, Salcudean SE, Allan AP. 1991. A six-degree-of-freedom magnetically levitated variable compliance fine-motion wrist: design, modeling, and control. *IEEE Trans. Robot. Autom.* 7:320–32
74. Makino H. 2014. Development of the SCARA. *J. Robot. Mechatron.* 26:5–8
75. Makino H, Furuya N. 1980. Selective compliance assembly robot arm. In *Proceedings of the First International Conference on Assembly Automation*, pp. 77–86. Bedford, UK: IFS
76. Rey L, Clavel R. 1999. The Delta parallel robot. In *Parallel Kinematic Machines*, ed. CR Boër, L Molinari-Tosatti, KS Smith, pp. 401–17. London: Springer
77. Engelberger J, Lock D, Willis K. 1980. *Robotics in Practice: Management and Applications of Industrial Robots*. New York: AMACOM

78. Atkar PN, Greenfield A, Conner DC, Choset H, Rizzi AA. 2005. Uniform coverage of automotive surface patches. *Int. J. Robot. Res.* 24:883–98
79. Inoue H. 1971. Computer controlled bilateral manipulator. *Bull. JSME* 14:199–207
80. Mason MT. 1981. Compliance and force control for computer-controlled manipulators. *IEEE Trans. Syst. Man Cybernet.* 11:418–32
81. Raibert MH, Craig JJ. 1981. Hybrid position/force control of manipulators. *J. Dyn. Syst. Meas. Control* 102:126–33
82. Hogan N. 1980. Mechanical impedance control in assistive devices and manipulators. In *Proceedings of the 1980 Joint Automatic Control Conference*, pap. TA-10-B. Dayton, OH: Autom. Control Counc.
83. Khatib O. 1987. A unified approach for motion and force control of robot manipulators: the operational space formulation. *IEEE J. Robot. Autom.* 3:43–53
84. Bischoff R, Kurth J, Schreiber G, Koeppe R, Albu-Schaeffer A, et al. 2010. The KUKA-DLR lightweight robot arm - a new reference platform for robotics research and manufacturing. In *2010 41st International Symposium on Robotics and 2010 6th German Conference on Robotics (ROBOTIK)*, pp. 741–48. Berlin: VDE
85. Villani L, De Schutter J. 2016. Force control. See Ref. 35, pp. 195–220
86. Pratt GA, Williamson MM. 1995. Series elastic actuators. In *1995 IEEE/RSJ International Conference on Intelligent Robots and Systems: Human Robot Interaction and Cooperative Robots*, Vol. 1, pp. 399–406. New York: IEEE
87. Laliberté T, Gosselin C. 1998. Simulation and design of underactuated mechanical hands. *Mech. Mach. Theory* 33:39–57
88. Dollar A, Howe R. 2010. The highly adaptive SDM hand: design and performance evaluation. *Int. J. Robot. Res.* 29:585–97
89. Monkman G, Hesse S, Steinmann R. 2007. *Robot Grippers*. New York: Wiley
90. Birglen L, Schlicht T. 2018. A statistical review of industrial robotic grippers. *Robot. Comput.-Integr. Manuf.* 49:88–97
91. Monkman G. 1997. An analysis of astrictive prehension. *Int. J. Robot. Res.* 16:1–10
92. Lozano-Pérez T, Wesley M. 1979. An algorithm for planning collision-free paths among polyhedral obstacles. *Commun. ACM* 22:560–70
93. Lanczos C. 1970. *The Variational Principles of Mechanics*. Toronto: Univ. Toronto Press
94. Kavraki L, Svestka P, Latombe J, Overmars M. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. Autom.* 12:566–80
95. Kavraki LE, LaValle SM. 2016. Motion planning. See Ref. 35, pp. 139–62
96. Tomovic R, Boni G. 1962. An adaptive artificial hand. *IRE Trans. Autom. Control* 7:3–10
97. Salisbury JK Jr. 1982. *Kinematic and force analysis of articulated hands*. PhD Thesis, Dep. Mech. Eng., Stanford Univ., Stanford, CA
98. Jacobsen S, Wood J, Knutti D, Biggers K. 1984. The Utah/MIT dextrous hand: work in progress. *Int. J. Robot. Res.* 3:21–50
99. Ulrich NT. 1989. *Grasping with mechanical intelligence*. Tech. Rep. MS-CIS-89-51, Univ. Pa., Philadelphia
100. Butterfass J, Grebenstein M, Liu H, Hirzinger G. 2001. DLR-Hand II: next generation of a dextrous robot hand. *2001 IEEE International Conference on Robotics and Automation*, Vol. 1, pp. 109–14. New York: IEEE
101. Catalano M, Grioli G, Farnioli E, Serio A, Piazza C, Bicchi A. 2014. Adaptive synergies for the design and control of the Pisa/IIT softhand. *Int. J. Robot. Res.* 33:768–82
102. Birglen L, Laliberté T, Gosselin C. 2008. *Underactuated Robotic Hands*. Berlin: Springer
103. Hanafusa H, Asada H. 1977. A robot hand with elastic fingers and its application to assembly process. In *IFAC Proc. Vol*. 10:127–38
104. Lakshminarayana I. 1978. *The mechanics of form closure*. Tech. Pap. 78-DET 32, ASME, New York
105. Asada H, By AB. 1985. Kinematic analysis of workpart fixturing for flexible assembly with automatically reconfigurable fixtures. *IEEE J. Robot. Autom.* 1:86–94
106. Mishra B, Schwartz JT, Sharir M. 1987. On the existence and synthesis of multifinger positive grips. *Algorithmica* 2:541–58

107. Rimon E, Burdick JW. 1995. New bounds on the number of frictionless fingers required to immobilize planar objects. *J. Robot. Syst.* 12:433–51
108. Bicchi A. 1995. On the closure properties of robotic grasping. *Int. J. Robot. Res.* 14:319–34
109. Bicchi A, Kumar V. 2000. Robotic grasping and contact: a review. In *2000 ICRA: International Conference Robotics and Automation*, pp. 348–53. New York: IEEE
110. Bicchi A. 2000. Hands for dexterous manipulation and robust grasping: a difficult road toward simplicity. *IEEE Trans. Robot. Autom.* 16:652–62
111. Miller A, Allen P. 2004. Graspit! A versatile simulator for robotic grasping. *IEEE Robot. Autom. Mag.* 11:110–22
112. Ciocarlie M, Allen P. 2009. Hand posture subspaces for dexterous robotic grasping. *Int. J. Robot. Res.* 28:851–67
113. Lozano-Pérez T, Jones J, Mazer E, O'Donnell P, Grimson W, et al. 1987. Handey: a robot system that recognizes, plans, and manipulates. In *1987 IEEE International Conference on Robotics and Automation*, Vol. 4, pp. 843–49. New York: IEEE
114. Stilman M, Schamburek JU, Kuffner J, Asfour T. 2007. Manipulation planning among movable obstacles. In *2007 IEEE International Conference on Robotics and Automation*, pp. 3327–32. New York: IEEE
115. Alami R, Simeon T, Laumond JP. 1989. A geometrical approach to planning manipulation tasks. The case of discrete placements and grasps. In *The Fifth International Symposium on Robotics Research*, pp. 453–63. Cambridge, MA: MIT Press
116. Simunovic S. 1975. Force information in assembly processes. In *Proceedings of the 5th International Symposium on Industrial Robots*, pp. 415–31. Dearborn, MI: SME
117. Whitney DE. 1983. Quasi-static assembly of compliantly supported rigid parts. *ASME J. Dyn. Syst. Meas. Control* 104:65–77
118. Whitney DE, Nevins JL. 1979. What is the remote center compliance and what can it do? In *Proceedings of 9th International Symposium and Exposition on Industrial Robots (ISIR)*, pp. 135–52. Dearborn, MI: SME
119. Murray RM, Li Z, Sastry SS. 1994. *A Mathematical Introduction to Robotic Manipulation*. Boca Raton, FL: CRC
120. Mason MT, Salisbury JK Jr. 1985. *Robot Hands and the Mechanics of Manipulation*. Cambridge, MA: MIT Press
121. Marigo A, Bicchi A. 2000. Rolling bodies with regular surface: controllability theory and applications. *IEEE Trans. Autom. Control* 45:1586–99
122. Trinkle J, Paul R. 1990. Planning for dexterous manipulation with sliding contacts. *Int. J. Robot. Res.* 9:24–48
123. Rus D. 1999. In-hand dexterous manipulation of piecewise-smooth 3-D objects. *Int. J. Robot. Res.* 18:355–81
124. Chavan-Dafle N, Rodriguez A, Paolini R, Tang B, Srinivasa S, et al. 2014. Extrinsic dexterity: in-hand manipulation with external forces. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1578–85. New York: IEEE
125. Viña B FE, Karayiannidis Y, Pauwels K, Smith C, Kragić D. 2015. In-hand manipulation using gravity and controlled slip. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5636–41. New York: IEEE
126. Hou Y, Jia Z, Johnson A, Mason MT. 2016. *Robust planar dynamic pivoting by regulating inertial and grip forces*. Presented at Workshop Algorithmic Found. Robot., 12th, San Francisco, Dec. 18–20
127. Shi J, Woodruff JZ, Umbanhowar PB, Lynch KM. 2017. Dynamic in-hand sliding manipulation. *IEEE Trans. Robot.* 33:778–95
128. Chavan-Dafle N, Rodriguez A. 2017. Sampling-based planning of in-hand manipulation with external pushes. arXiv:1707.00318
129. Mason MT. 1986. Mechanics and planning of manipulator pushing operations. *Int. J. Robot. Res.* 5:53–71
130. Lynch KM, Mason MT. 1995. Stable pushing: mechanics, controllability, and planning. In *Algorithmic Foundations of Robotics*, ed. K Goldberg, D Halperin, J-C Latombe, R Wilson, pp. 239–62. Boston, MA: A.K. Peters
131. Dogar MR, Srinivasa SS. 2011. A framework for push-grasping in clutter. In *Robotics: Science and Systems VII*, ed. H Durrant-Whyte, N Roy, P Abbeel, pp. 65–72. Cambridge, MA: MIT Press

132. Yu KT, Bauza M, Fazeli N, Rodriguez A. 2016. More than a million ways to be pushed. A high-fidelity experimental dataset of planar pushing. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 30–37. New York: IEEE
133. Chang LY, Zeglin GJ, Pollard NS. 2008. Preparatory object rotation as a human-inspired grasping strategy. In *Humanoids 2008 – 8th IEEE-RAS International Conference on Humanoid Robots*, pp. 527–34. New York: IEEE
134. Goldberg KY. 1993. Orienting polygonal parts without sensors. *Algorithmica* 10:201–25
135. Erdmann MA. 1998. An exploration of nonprehensile two-palm manipulation. *Int. J. Robot. Res.* 17:485–503
136. Erdmann MA, Mason MT. 1988. An exploration of sensorless manipulation. *IEEE J. Robot. Autom.* 4:369–79
137. Peshkin MA, Sanderson AC. 1988. Planning robotic manipulation strategies for workpieces that slide. *IEEE J. Robot. Autom.* 4:524–31
138. Mason MT, Rodriguez A, Srinivasa SS, Vazquez AS. 2012. Autonomous manipulation with a general-purpose simple hand. *Int. J. Robot. Res.* 31:688–703
139. Berretty RP, Overmars MH, van der Stappen A. 2002. Orienting polyhedral parts by pushing. *Comput. Geom.* 21:21–38
140. Salisbury JK. 1987. Whole arm manipulation. In *Proceedings of the 4th International Symposium on Robotics Research*, ed. R Bolles, B Roth, pp. 183–89. Cambridge, MA: MIT Press
141. Shepard RN, Metzler J. 1971. Mental rotation of three-dimensional objects. *Science* 171:701–3
142. Balkcom DJ, Mason MT. 2008. Robotic origami folding. *Int. J. Robot. Res.* 27:613–27
143. Maitin-Shepard J, Cusumano-Towner M, Lei J, Abbeel P. 2010. Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *2010 IEEE International Conference on Robotics and Automation*, pp. 2308–15. New York: IEEE
144. Lemelin P, Schmitt D. 2016. On primitiveness, prehensility, and opposability of the primate hand: the contributions of Frederic Wood Jones and John Russell Napier. In *The Evolution of the Primate Hand: Anatomical, Developmental, Functional, and Paleontological Evidence*, ed. TL Kivell, P Lemelin, BG Richmond, D Schmitt, pp. 5–13. New York: Springer
145. Chen IM, Yim M. 2016. Modular robotics. See Ref. 35, pp. 531–42
146. Moravec H. 1988. *Mind Children: The Future of Robot and Human Intelligence*. Cambridge, MA: Harvard Univ. Press
147. Reif J. 1979. Complexity of the mover's problem and generalizations (extended abstract). In *20th Annual IEEE Conference on Foundations of Computer Science*, pp. 421–27. New York: IEEE
148. Johnson DS. 1973. *Near-optimal bin packing algorithms*. PhD Thesis, Dep. Math., Mass. Inst. Technol., Cambridge, MA
149. Hawking S. 2001. *The Universe in a Nutshell*. New York: Bantam
150. Pinto L, Gupta A. 2017. Learning to push by grasping: using multiple tasks for effective learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2161–68. New York: IEEE
151. Goldberg DE, Holland JH. 1988. Genetic algorithms and machine learning. *Mach. Learn.* 3:95–99
152. Fitts PM. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *J. Exp. Psychol.* 47:381–91
153. Mortimer J, Rooks B. 1987. *The International Robot Industry Report*. Berlin: Springer
154. Omron Adept Technol. 2017. *Adept Quattro s650HS*. Data Sheet, Omron Adept Technol., San Ramon, CA. **https://www.adept.com/downloads/doc_details/341-datasheet-adept-quattro-s650hs-robot**
155. Blickhan R, Seyfarth A, Geyer H, Grimmer S, Wagner H, Günther M. 2007. Intelligence by mechanics. *Philos. Trans. R. Soc. Lond. A* 365:199–220
156. Byrne RW, Corp N, Byrne JME. 2001. Estimating the complexity of animal behaviour: how mountain gorillas eat thistles. *Behaviour* 138:525–57
157. Lee K, Moses M, Chirikjian GS. 2008. Robotic self-replication in structured environments: physical demonstrations and complexity measures. *Int. J. Robot. Res.* 27:387–401
158. Sanderson A. 1984. Parts entropy methods for robotic assembly system design. In *1984 IEEE International Conference on Robotics and Automation*, pp. 600–8. New York: IEEE