ANNUAL REVIEWS

*Annual Review of Control, Robotics, and Autonomous Systems*

# What Is Robotics? Why Do We Need It and How Can We Get It?

Daniel E. Koditschek

Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA; email: kod@upenn.edu

Annu. Rev. Control Robot. Auton. Syst. 2021. 4:1–33

First published as a Review in Advance on January 5, 2021

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

https://doi.org/10.1146/annurev-control-080320-011601



## Keywords

robotics, synthetic science, programming work, active materials, fundamental physical limits, hybrid dynamical systems theory, applied category theory, applied type theory, academic disciplines, academic departments, social impact of technology, diversity of scientists

## Abstract

Robotics is an emerging synthetic science concerned with programming work. Robot technologies are quickly advancing beyond the insights of the existing science. More secure intellectual foundations will be required to achieve better, more reliable, and safer capabilities as their penetration into society deepens. Presently missing foundations include the identification of fundamental physical limits, the development of new dynamical systems theory, and the invention of physically grounded programming languages. The new discipline needs a departmental home in the universities, which it can justify both intellectually and by its capacity to attract new diverse populations inspired by the age-old human fascination with robots.

# 1. INTRODUCTION

Robotics is a growing body of technology presently in the early stages of developing its disciplinary foundations. Emerging from decades of collaboration among biologists, computer scientists, and engineers—its looming commercial presence a mere harbinger of the enormous social impact to come—its reach has already begun to outstrip the grasp of its still meager foundations.

This article proposes that robotics is destined to become (eventually understood and explicitly advanced as) a new synthetic science concerned with programming work—the exchange of energy and information between a machine and its environment toward some specified set of goals. After exploring what these terms mean in Section 2, the argument turns to the value of their consideration for advancing the field in Section 3. The article proceeds in Section 4 to offer a speculative view of how this new discipline might come into being, and then concludes in Section 5 with a brief account of why it was written and who might benefit from reading it.[1]

# 2. WHAT IS ROBOTICS?

In contrast to the natural sciences—bodies of theory revealing the empirical realm as it exists—a synthetic science seeks theory that anticipates what could be brought into existence (1). This requires that we go beyond Feynman's oft-quoted statement "What I cannot create, I do not understand,"[2] insisting not merely that we test our intuition and understanding by what we create, but also that our understanding be codified in principles of design that predict its artifacts' empirical properties in advance of their creation.[3] Thus, a discipline of robotics must aspire to a formal body of theory that endows some language of goal specification with the blueprints of material construction for mechanisms that will exchange energy with their environment in a manner that provably achieves the goals. Such specifications must scope types of environments so as to express settings within which the goals must or cannot prevail. In particular, the theory must prevent the specification of goals that cannot prevail in any environment because their achievement would violate fundamental physical limits.

## 2.1. Lessons from Computer Science

Computer science, which is concerned with programming the exchange of information between a machine and its environment (4),[4] emerged over the course of the first half of the twentieth century as a body of theory and practice at the intersection of mathematical logic, linguistics, neuroscience,

Supplemental Material >

[1]The **Supplemental Appendix** offers a supporting (but necessarily abbreviated) sketch of ideas from dynamical systems theory (Section A), some details on models of a robot's interface to its environment (Section B), and further details (Section C) supporting the speculation on how to advance the discipline of robotics (discussed in Section 4 of the main text); references led by these letters (often in combination with punctuated numbers) refer to the corresponding sections in the **Supplemental Appendix**. Footnotes using italicized lowercase letters rather than numerals similarly refer to footnotes in the **Supplemental Appendix**; these footnotes introduce an additional degree of technical detail that depends on notions that are introduced in (and, hence, must be relegated to) the supplement.

[2]Rendering the contrapositive of this quotation as "If I do understand, then I can create," synthetic biologists have also begun to ponder the gap between that statement and its converse, "If I can create, then I do understand" (2)—a central focus of this article.

[3]Here and throughout, such terms as synthesis, design principles, design methodologies, and so on are intended to connote a formal tool set whereby specification might map onto behavior. Addressing the profound and fascinating ascent from the science of tools toward the art of their application pondered over the centuries by thinkers far more capable than this author (3) lies well beyond the scope of this article.

[4]The technological prescience of this late-twentieth-century paradigm shift to concurrent models of computing—specification of the computer–world information interface (5, chap. 41)—is by now

and electrical engineering (8). It took a good part of the second half to win gradual recognition as a discipline (9). Amidst the many controversies about what else it has become (10), the discipline arguably earned the distinction of synthetic science upon Landauer's (11) identification of the fundamental physical cost of information, following which a great deal more has been discovered about the fundamental limits of computing (12).

The lessons of computer science for robotics include the fruitful interplay between our ability to express design goals and the capacities of physical substrates to achieve them (Section 2.1.1); the vital role of modularity and its apotheosis in the form of a program (Section 2.1.2); and the central importance of the design triple—distinguishing between and relating tasks, architectures for accomplishing them, and the environments within which those accomplishments are sought (Section 2.1.3).

**2.1.1. Computational substrates.** Few would contest that relating formally specified architecture (13) to physical substrate (14–16) played a decisive role in the past half century's information technology revolution. The remarkable advances in theory that permit researchers to tie abstract features of logical operation to physical fabrication and design in light of such absolute barriers as Landauer's limit (17) span nearly 20 orders of magnitude—from nanoscale devices ($\sim 10^{-8}$ W) to regional data centers ($\sim 10^{9}$ W) (18)—a staggering testament to the success of computing as a physical synthetic science, underlying its inescapable social impact. Still more compelling from the envious view of a nascent-robotics scientist is the computer engineer's crucial ability to define and operationalize design tradespaces whose dimensions freely mix and compare metrics of information (instruction word width, data versus instruction memory address space, instruction loop length, etc.) against those of energy (19).

Ingenious technological work-arounds that continue to mitigate Landauer's thermodynamic threat to Moore's law (20) do not negate but rather underscore the value of identifying and designing explicitly with respect to such fundamental physical limits. In fact, a key difference between robotics and computer science is illuminated by Bennett's recourse to reversible computation (21). Physically realized decades later (22), reversible computing might promise to free from Landauer's limit machines that merely exchange information with the environment. Regardless, however, the ingenuity of computer engineers has to date enabled users' abstractions to function as if this were the case. By contrast, computing might be cool, but robotics must be hot (Section 2.2.1.1).

**2.1.2. From modules to programs.** Key to this staggering information technology revolution has been the role of programming languages. Mead & Conway's (23) formal design rules that could scale with the computing substrates (15) fueled the development of register transfer languages for the specification of instruction set processors (24). McCarthy (25) recognized that Church's $\lambda$-calculus could be used to represent programs, elevating them to first-class mathematical objects (26), eventually reconceived as affording the full expressive power of constructive mathematics itself (27).

Modular design via kits of interchangeable parts had been practiced for millennia (28, p. 26) before Simon's parable of the two watchmakers enshrined its merit for hierarchical composition (1). Modularization promotes reuse, inviting grammars—general rules of recombination—to accommodate varying tasks. The Chomsky hierarchy calibrates a language, a subcollection of

---

amply attested not only by the preponderance of communications applications (6) but also by the unremitting rack- and chip-level interconnect bottleneck that has begun to encroach upon even algorithmic complexity theory (7).

lexical strings, in terms of its computational model specified by the memory capacity and access of a discrete finite automaton that can recognize or generate it (8). But robotics is initially concerned with a different class of dynamical systems (Section B). Moreover, the imperatives of synthesis demand not merely a language, but a programming language (29)—a grammar of specification (30)—whose fundamental account follows the type-theoretic branch (5)[e] of the Church–Turing hypothesis first taken by McCarthy's turn toward AI (31).

#### 2.1.3. The design triple.

Throughout his book *The Sciences of the Artificial*, Simon (1) calls our attention to a designer's problem triple: the goal sought, by an artifact, in its environment. Yet his unforgettable image of the ant steered by its surroundings seems to have been quickly forgotten within the AI field he helped found. Notwithstanding important exceptions (32), AI has been preoccupied with its architectures and their efficacy in achieving its goals, reserving very little systematic attention for even empirical study of their situated behavior, much less an attempt to theorize about which classes of environments will abet or impede them.[5]

The broader reaches of the discipline came to accept and embrace the centrality to computer science of understanding and formally specifying the interaction of an agent with its environment (4). Yet even the most scholarly, comprehensive, and contemporaneously valuable text in the field of AI (33), while emphasizing this interaction, does not offer a unified formalism for its specification and analysis—a conundrum inherited by robotics (Section B.2).

### 2.2. Problems Awaiting the Discipline of Robotics

The essential job of a robot[6] is to perform work on its environment specified by its user. Thus, the discipline of robotics must address three main problems:

- Problem 1 (the robot): Build a physical body equipped with material resources that can move energy from a (typically chemical) reservoir to the environment directed along the right degrees of freedom at the right time.
- Problem 2 (the robot program): Accept a specification of a goal over the state space of a robot conjoined to its environment and then either declare failure in consequence of some insufficient resource or execute a solution to Problem 1 that achieves the goal.
- Problem 3 (the robot design): Accept a specification of a world model (a class of environments) and a task domain (a class of goals), then return a specification of resources along

[5]It would require a digression too long for this article to discuss the appropriate role of benchmarking—empirical performance in curated environments—in robotics. By itself, a preoccupation with benchmarks cannot be construed as proper empirical science. Surely, testing against benchmarks cannot be mistaken for theorizing (forming precisely stated hypotheses to be tested in new environments of a postulated type) about what features of what types of environments will elicit what sorts of behaviors from an artifact. Alternatively (veering a bit dangerously toward the art of science practice banned in footnote 3 from consideration in this article), it is hard to imagine any novel empirical discovery ever arising from a collection of examples chosen to summarize what is already known to be hard about a design problem. Benchmarking can play a vital role in helping assess the likely performance of an architecture in the presence of environmental conditions or task specifications, which go beyond the scope of its (typically disappointingly conservative) provably sufficient conditions for success along the lines discussed in Section 3.1.

[6]To bound its scope, this article focuses on the specific case of mechanical robots (whose energetic exchange is described by classical mechanics) at the human scale (lengths of roughly $10^{-2}$–$10^{2}$ m). Extending or improving on these ideas toward the promise of power- and information-autonomous robots down below the millimeter scale (34) and of machines that work by chemistry (35) will become the urgent business of the emerging future discipline.

with a solution to Problems 1 and 2 that achieves tasks in those environments given those resources.

Section 2.2.1 sketches what is known about the challenges of Problem 1—realizing the working substrate. Section 2.2.2 addresses Problem 2, introducing models of the robot–environment interface and examining the conceptual gaps impeding the programming and execution of abstract goals. Problem 3 is taken up in Section 4 through the more speculative discussion required to do justice to terms such as task domain and class of environments.[7]

### 2.2.1. Physical resources: substrates of work.

Solutions to Problem 1—limbs and body, and their endowment with actuation, perception, and manipulation capabilities—have exhibited substantial advances since the first modern designs appeared roughly half a century ago (38). But, possibly apart from the benefits of better chips and sensors, there is no hint of accelerating progress comparable to that already evident within the first two decades of the information technology revolution (13, 14). Researchers have only begun to explore the intuitively compelling questions of how and why mechanical circuits—interconnected systems whose assigned tasks require exchange of joules as well as bits—are fundamentally more difficult to modularize, design, and scale than very-large-scale integration (VLSI) (39).[8]

This section proposes that insufficient attention to fundamental physical resources and the way they must interact impedes progress—a key obstacle to accounting for (and, hence, accelerating) its slow pace. The physical setting of robotics is distinguished by three specific resources (energy, information, and bonds); their rates; and, most particularly, the complicated interplay between them that is required to harness and effectively deploy any of them. The sidebar titled The Challenge of Problem 1 summarizes the discussion in the rest of this section that addresses these fundamental physical resources of robotics.

#### *2.2.1.1. Energy and information rates.*

Because there is a premium on getting work done quickly, power—the rate at which actuators can move joules—is a first scarce resource. Prior to

> **THE CHALLENGE OF PROBLEM 1**
>
> Addressing Problem 1 requires materials affording scalable mechanical power and information processing (Section 2.2.1.1) as well as grip (Section 2.2.1.2). Design methodologies that can specify their distribution across the interior volumes and interface surfaces of robot bodies as a function of task–environment pairings require models of the robot–environment work and information interface (Section 2.2.2.1) that incorporate explicit representations of these material resources.

[7] Again, constraints of space and time preclude consideration in this article of agency—the capacity of a system to develop (and perhaps execute) its own goals. Inevitably, gradations of autonomy required to achieve users' goals will begin to encompass increasingly broad decision-making capabilities and motivational dynamics (36) as specifications become more abstract and the environment departs from their designers' anticipated type. The embodied situation of robotics presents an ideal setting for empirically grounded advancement of the science of autonomy (37), deserving of its own dedicated discussion.

[8] One anecdotal but revealing measure of the challenge is presented by the effort of one of VLSI's cofounders (23) to foment a similar revolution for analog circuit design (40)—an exciting and influential [there are now thousands of papers working within this promising framework (41)] but surely not (at least yet) comparably paradigm-shifting development. Section B.2.2 offers some discussion of the crucial role that might be played in robotics by analog computation and its recent advances (e.g., 42).

Huygens's invention of a gunpowder-forced linear actuator in the 1670s, animal muscle was the only source of taskable mechanical work (43). Different arrangements of muscle fibers and skeletal attachment work as motors, brakes, springs, and struts with huge variations in actuation capacity, achieving, for example, up to $10^3$ W/kg at 10 Hz or $10^2$ W/kg at $10^2$ Hz (44), over mass scales spanning seven orders of magnitude, $10^{-4}$–$10^3$ kg (45). Meanwhile, the twentieth century's seedbed of physical principles and mechanisms available for synthetic actuator design (46) has blossomed into a hothouse of active materials energized by a diverse range of physical effects (43). These abundant variations reveal substantially contrasting performance along such dimensions of merit as power density, bandwidth, stress, strain, efficiency, and linearity, many exceeding biological performance over a delimited subset (47). However, no single synthetic approach seems close to matching the pluripotency of animal muscle tissue, much less its capacities of self-assembly, regeneration, and intrinsic adaptation (43).

Whereas computation need not intrinsically entail energetic cost (21), robots must expend energy, not merely when working explicitly on their environment (48) but also to achieve any task formulated in terms of dynamical attractor basins, as advocated below (Section 4.1.2.2).[9] This motivates the question, What are the fundamental limits of specific power? It seems inevitable that dynamical versions of Landauer's limit tying together not simply energy and information but also power and bit rate (Section C.1.1) will play an essential role in any formulation of fundamental physical constraints on robotics. Understanding these relationships has in the last decade become an active area of theoretical inquiry (55) and increasingly practical device design (56) at the nanoscale. This article argues (Section 4.1.1.1) that exploring the implications of such fundamental constraints at the mesoscale of conventional robotics will play a key role in establishing the discipline. Equally important, of course, is the question of how to effectively use whatever specific power a robot has been endowed with (Section 4.1.3.1).

***2.2.1.2. Making and breaking bonds.*** Whether fixed at design time or, crucially, designed to be made and broken through the course of a task along the degrees of freedom relevant to the robot's goals, bonds must be strong enough to withstand the forces they communicate. Apparently, grip—material properties conferring friction and adhesion, their removal, and their facility for higher rates of alternation—is also scarce.

Joining of materials takes its place among the very oldest human technologies (57), and in this article, the term framing cost refers to the added mass incurred in achieving sufficiently strong and robust permanent mechanical bonds between dissimilar materials, as is typically required to integrate components such as actuators and sensors into a targeted morphology. Rapid advances in materials science and engineering anticipate an eventual future wherein fundamental principles of condensed matter physics drive 3D printing of working devices and complex structures from homogeneous ingredients, in analogy to the 2D inkjet printing of arbitrary images from a few colored ink reservoirs (58). Already, algorithmic thinking about the design of pluripotent materials from cutting and folding of homogeneous sheets (59) is achieving tunably compliant and shaped robot limbs (60), and programmable, spatiotemporally complex self-assembly of similarly active materials seems to be on the way as well (61). A major challenge for robotics is the development of

[9] Unlike classical Hamiltonian systems, hybrid compositions of piecewise lossless holonomic systems can sustain stable, partially asymptotically stable attractors (49–51), a phenomenon that has been exploited for inspirationally efficient legged locomotion (52, 53). Basins of ambient volume, as advocated in Section 4.1.2.2, are not possible, however, so the resulting steady-state behaviors are neutrally stable and can be pushed around—perhaps to useful effect by the controller (54), but just as easily by the environment to the detriment of the prescribed goal.

a systematic framework for specifying the distribution of materials' properties across the interior volumes and interface surfaces of the permanently (or, perhaps, developing—or even evolving) integrated body (Section 4.1.1.2).

A contrasting and still greater challenge is the development of materials to interface with the environment that are capable of making and breaking strong bonds, rapidly on command. Difficult though it may be to join with high mass-specific adhesion (normal) and friction (shear) forces, it is truly daunting to arrange for rapid, effortless detachment as well. Multicellular life is enabled by adhesion on many scales; in particular, animals use quickly reversibly forcible grippers to eat, crawl, climb, and capture prey in every terrestrial environment imaginable (62). In turn, impressive feats of manipulation abound in the animal world (63). Unsurprisingly, then, the architecture of grippers has appropriately preoccupied robotics for decades (64). But the physical basis of specific animals' grip has only recently begun to be understood (65), with the first revelation (66) only following millennia of puzzlement. Section 4.1.1.1 examines the prospects for identifying the fundamental limits of grip, touching as well upon the question of how to identify and use whatever grip the environment affords.

**2.2.2. Programming: architectures for tasks undertaken in environments.** Robots are quite different from computers as physical machines; hence, their Lagrangian internal models (67, 68) are different from the discrete finite automaton models of computing machines (8). But a program is a mathematical object as first viewed by McCarthy's (26) $\lambda$-calculus representation, and more generally as a function defined with respect to a theory of types—the specification of available domains and codomains (27). Types, the "central organizing principle of the theory of programming languages" (5, p. xvii), impose constraints on syntax ensuring that its evaluation always yields a valid function—one whose behavior is defined by the resulting domain and codomain provided in the theory. In this view, the meaning of the mathematical objects specified by the syntax is given by the operational semantics of any evaluation step (29).

There are two essential challenges to achieving the physically grounded type theory required by Problem 2. First, prescribing the behavior of a system in its environment in any formal terms presupposes a description of the interface between them relative to which goals can then be specified. Section 2.2.2.1 assesses the availability of interface descriptions that clearly manifest the underlying physical resources so that alternative architectures can be considered relative to their requirements and operation in different environments. Second, since complex systems entail a hierarchy of interfaces (1), a behaviorally interesting robot will need a deep specification, requiring many layers of programming languages with clearly related operational semantics by which the meaning of a specification at any given layer is expressed in terms of the behaviors of the child and sibling components it calls out (30). Section 2.2.2.2 discusses the challenges of identifying useful internal layers and appropriate abstraction barriers to separate them. The sidebar titled The Challenge of Problem 2 summarizes the discussion in the rest of this section that addresses the problems of robot programming.

### THE CHALLENGE OF PROBLEM 2

Addressing Problem 2 requires reasoning about the complete synthesis triple (Section 2.1.3), including the robot's interface to its environment (Section 2.2.2.1) and its architecture in relation to its targeted tasks along with the environment's affordances toward achieving them (Section 2.2.2.2). Developing models of a robot's information interface to match the consensus model of its work interface remains a key challenge (Section B).

#### *2.2.2.1. Describing the robot–environment interface.*

A robot's working interface is specified by the properties of (*a*) its actuarium (a representation of its energy ports modeled by the first equation of Section B.1.1),[10] whose capacity is fundamentally characterized by specific power in W/kg; (*b*) its sensorium [a representation of its information channels (69)], which is characterized by bit rate in B/s; and (*c*) its tenacium (a representation of the speed and strength with which it can grip), which is characterized by rate at magnitude of reversal force ratio in N/s.[11] The brazenly Latinate terms aim to underscore the increasingly rich variety of smart materials (71) that offer new opportunities for distributing these capacities [e.g., W/(kg·m$^3$) or B/(s·m$^2$)] without suffering the framing and other scaling costs incurred by traditional actuators (46) or by integrating local, scalable computation in networked communications channels (72).

One potentially confounding aspect of these specifications is that all resources play a role in both interfaces: There is information (which degrees of freedom get what rates) in the work exchanged through ports, information channels (regarding both computation and communications) inevitably have an energy cost, and grip plays a central role in working on and receiving information about the environment. Notably, proprioceptive devices that participate in both actuation and sensing have long played an important role in robotics (73) and complicate the characterization of capacities as they enrich the behavioral suite of robots.

The consensus view within robotics of the work interface to the environment leads to an internal model over a conjoined robot–environment state, $x_v \in \mathcal{X}_v$, indexed by the mode of contact, $v \in V$ (Section B.1). Although this consensus may well be undercut by the introduction of advanced materials in view of their promise of distributed interfaces (71) and ubiquitous compliance (74), it remains a useful point of departure in considering the specification of robot architectures and their deployment. That caveat in place, it is convenient for purposes of exposition to posit the standard model for the work interface as having ports taking the form of the first equation in Section B.1.1.

Given this work-interface model, power resources are manifest by the properties of the input signal, $\tau$, in that equation—most simply, by its dimension relative to that of the state vector, $x$ (the degree of underactuation), or, more accurately, by a range of increasingly detailed internal dynamical models (e.g., 75–77) [and eventually requiring specific robot–environment modeling (78)], as discussed in Section B.1. By contrast, the immaturity of the discipline is such that it does not yet seem possible to propose a standard model for a robot's information interface (Section B.2). A likely general candidate for such models may be found in the process algebra literature (69); hence, it is convenient, even if merely as a conceptual placeholder,[12] to use the terminology channels (Section B.2.1) when discussing requirements of and prospects for representations of a robot's information resources.

#### *2.2.2.2. Specification: internal layers and models.*

Notwithstanding the central importance of robot perception (81, 82), the absence of a consensus model of channels (Section B.2.1) corresponding to the work interface of ports (Section B.1.1) compounds the challenge of

---

[10] See Section B.1.2 for a brief discussion about the prospects for a more general behavioral representation of this interface.

[11] This very speculative suggestion for appropriate units of grip in terms of time rate and magnitudes of load-to-preload and reversal ratio (70) is contextualizd in Section 4.1.1.1.

[12] It is daunting to contemplate the challenge here, since the eventual conceptual apparatus will need to encompass a vast scope of intricately entangled phenomena, ranging from transparency of drivetrains (Section 4.1.1.2) to the notorious correspondence problem (79) and statistical active perception (80) (Section B.2.1) through representation and use of analog computation (Section B.2.2).

designing a robot's deep layers. This is an intrinsically fraught enterprise because interior processing interfaces—information processing modules and their interrelations—are confusingly underconstrained. Viewed at the work interface modeled by the first equation of Section B.1.1, it is conceptually straightforward (albeit often technically challenging) to develop sufficient conditions for the success of an architecture in achieving a specified task relative to a specified environment since Newtonian mechanics is physically ineluctable. Furthermore, at the resolution of local behavior (Section B.2.2), careful control-theoretic reasoning can yield necessary conditions with generality adequate even to constrain animal architecture in the light of clever experiments (83–85). A central motivation for precision in describing a task domain, as discussed in Section 4.1.2.2, is the possibility of at least demonstrating sufficiency. Deeper conceptual progress in the form of necessary conditions on the internal architecture will require carefully articulated reasoning about generative models (37, 86) to establish how fundamental resource limits constrain it.[13]

These last observations underscore the obviously crucial but still inscrutable role and structure of memory—prior information about the robot and environment and the history of their encounters. The necessary critique of AI's obsession with representation in the designed architecture has unfortunately been clouded by charismatic calls "to use the world as its own model" (87, p. 140).[14] The true disciplinary question concerns representation of architecture: interior interface specifications that elucidate the design triple (Section 2.1.3). This fundamental problem seems to have been traditionally avoided by both the internal-model and reactive-behavior camps of AI. Both traditions shy away from any study of the robot–environment pairing beyond the most cursory level of empirical anecdote (93). But robotics cannot afford this luxury: The key challenge is developing tools for reasoning about the degree to which some internal model of a particular environmental affordance (and the conjoined robot–environment state in relation to it) is necessary or sufficient to achieve a particular goal (94).

There are two tightly related but distinct dimensions of depth in the interfacing layers to be accounted for. The first, manifesting the needs of task specification, arises from the human predilection for abstraction as a means of taming behavioral complexity and is particularly challenged by the signal–symbol divide (Section 4.1.2.2). The second, introduced by the physical resources of information and grip, arises from the ever more spatiotemporally distal aspects of the environment that must be mechanically engaged or perceptually experienced as behavioral complexity increases. Here, the conceptual bottleneck lies in appropriately abstract models of the environment, both endowed by the designer (Section B.2.3.1) and learned from experience (Section B.2.3.2).[15] This

[13]Animal architecture offers a tantalizing source of necessary conditions for any performance model of animals' agency. However, it remains to be explored whether such constraints have any purchase over the competence models of behavior that would be adequate (and likely preferred) for prescribing those properties of a robot's internal architecture required to achieve a specified task in a given environment.

[14]Sophomore control engineers who have contemplated stabilizing a force-controlled mass–spring–damper system with only position feedback understand the necessity of augmenting sensory cues with internal representations of certain environments to achieve certain tasks (88, example 6.2.1). More general formal reasoning reveals that a complete internal model of the relevant environmental disturbance is necessary for any control architecture capable of stably and robustly rejecting it (89). The reinvention of decades-prior hierarchically arranged inner and outer (or minor) loops (90) reveals such "new approaches" to robotics (87) as uninformed by and uncommitted to science. Section B.2.3 briefly reexamines such questions by considering a range of design settings. At the limit of this range in task–environment pairings lies the extreme of "kicking the sensing habit" (91) entirely via open-loop procedures that require no measurement of the world state at all—provably guaranteed and empirically demonstrated to succeed in such structured settings (92).

[15]Given the recent triumphal emergence of computational learning (95), it is particularly unfortunate that limitations of space and time preclude anything close to a consideration of their import for robotics. Their

article seeks to advance the perspective (to be first articulated in Section 4.1.2.2 and pursued with more technical detail in Sections B.2 and C.2) that at whatever level of specification, the daunting challenge of sensorimotor coherence—keeping symbols arising from learned models and their sensed referents relevant to programmed task expressions—can be overcome by grounding them all in the sublevel sets of the energy landscape.[16]

# 3. WHY DO WE NEED IT?

## 3.1. Foundations of Intellectual Progress

Are robots getting better? Certainly their computers and sensors are. If we set up competitions, we can pretty well discern when there is a winner, measure how much progress has been made between iterations (98), and surely recognize technological inadequacies from the post hoc scorn of the lay public (99). But why? Continually relearning that the technology project is very hard (100–103) does not seem to diminish its hundreds of billions of capital inflow (104–106). But neither does merely throwing more money and replacing one test with another test seem to bring technological progress beyond that all too readily ascribable to improved component hardware or advancements in algorithms imported from distinct fields.[17]

Important conceptual progress in robotics has surely accrued and can be roughly charted by the appearance of landmark monographs. The algorithmic foundations of motion planning (109) were greatly enriched and made practicable by the adaptation of Bayesian filtering to navigation and mapping (110), and it can be expected that the huge impact of learning in this domain will in time generate a comparably high-impact summary account. Insight into the mechanics of manipulation has grown dramatically (111, 112). The empirical foundations of dynamical locomotion (113) were greatly strengthened by the introduction of more formal ideas from nonlinear feedback control (114). But it might not be apparent to a student—or even an accomplished practitioner—how these books relate to each other. Indeed, it does not yet seem clear how to build machines that benefit simultaneously from all three traditions of insight. In particular, the physical resources whose scarcity most dramatically obstructs performance seem different in each: information flows for navigation, gripping affordance for manipulation, and power budgets for locomotion (although all three make an appearance in each). How do these different sorts of robotic capabilities fit together? What are the fundamental limits to performance for any or all of them?

It is the job of a scientific discipline to pose carefully and answer such questions. Formal synthesis—a precisely stated hypothesis of what properties must inhere from a design in advance of its construction—is a profound enabler of better technology even if it is construed as merely a debugging tool. If a correct theorem states that a particular architecture must be capable of achieving a specified task in a type of environment, and empirical evidence contradicts that conclusion, then we know that there is some discrepancy between the assumptions in the hypothesis and the prevailing conditions in the physical world. Either the architecture fails to achieve some

---

huge potential for control has been understood for decades (96). Used with precision in architectural design (97), they hold at least comparable value for robotics; see footnote 16 just below and Section B.2.3.2 for brief speculative remarks bearing on the matter.

Supplemental Material >

[16]These terms and concepts related to sensorimotor coherence are given brief technical descriptions in footnotes *m* and *n* in the **Supplemental Appendix** and the text that calls them out. Footnote *k* provides a speculative but succinct general statement of this idea in that more technical context.

[17]Economists have understood for decades that demand-side pressure is inadequate to generate new technology absent appropriate scientific foundations (107, chap. 14). There is at least some empirical evidence that radical invention in the sphere of mechanical devices may be particularly driven by new advances in fundamental knowledge (108).

specific capability listed as necessary or the environment fails to conform to the properties of the assumed type.

This last possibility underscores the driving intellectual importance of clearly posited assumptions and the proofs they enable. Without them, we have no methodical way of drilling down into the details of what makes a synthesis problem hard. By definition, the environmental model is an abstraction that will miss details of the physical setting. Explicitly stating what environmental properties must be assumed in order for the architecture to be appropriate crystallizes the role of that affordance in enabling the task. It can clarify the appropriate target of benchmarking (e.g., facilitate the curation of key out-of-scope settings, along the lines discussed in footnote 5). Thus, such abstraction plays a central role in teasing out what details are essential to take into account and what specific design challenges arise from what specific adverse conditions.

### 3.2. Foundations of Research and Teaching

A logistical reason to establish the discipline of robotics is that contemporary civilization enshrines disciplines in universities that commit substantial resources to departments for their propagation and advancement in the broadest interests of human knowledge. This may change—many prophesy, and some already find evidence of, great disruption. But for present purposes, it is convenient to envision the prospects for a discipline of robotics through the lens of its departmental manifestation in a research university. Here, pedagogical imperatives confer the greatest intellectual benefit. Delivery to a novice provides the best motivation for and evidence of a deep understanding. The conceptual barriers between the sorts of benchmark monographs just discussed underscore the huge advantage their impressive authors and indebted readers would all accrue from the obligation to explain to a sophomore robotics major how they fit together.

The next most important role of the robotics department is to hire its replacements. While the key criteria for wise faculty appointments remain creative talent and intellectual ambition, the accompanying arguments about what direction to push in and why play a crucial role in the maturing of a discipline. The droning on about how the department needs not just one but actually three more scholars in the area of one's five most recent publications ultimately confers significant intellectual value in the aggregate, however near unendurable in the moment. For the vision of what should come ahead must be contested not merely in dropfuls by the grant but in bushels—or, with extreme luck, in tons—by the career. As robotic technology's impacts deepen, these vital arguments about where the fundamental questions lie are increasingly camouflaged or distorted when cramped into the mold of neighboring disciplines. Correspondingly, their potential benefit is lost to the field whose present coherence and future invention depend upon them.

If the history of computer science holds the lessons for robotics urged in this article, then not the least important role of the discipline's creation will be to referee the tussle over theory and practice. In one convincing reading (115), the discipline of computing emerged from specialized corporate training programs (in the 1950s) to educate a practice of software engineering promulgated by universities (in the 1960s), shortly encapsulated within an academically focused canon of theory (in the 1970s and 1980s), the escape beyond which was engineered by a creed of disciplinary problem solving that persists to the present day. The cycles of tension, expansion of purview, and reemergence of consensus regarding curricula and foundational agendas that characterize departmental incarnations of disciplines seem to provide essential ballast for any technology that boasts the accelerating social impact to which the argument now turns.

### 3.3. Imperatives of Social Impact

Ready or not, robots seem finally to be on the way. They have already transformed factories. Bold announcements and acquisitions by large corporate actors herald their appearance with greater autonomy and in less structured settings throughout the commercial sphere. But roboticists understand that such pronouncements are the mark of irrational exuberance (116) and dangerously misleading product advertising (117). Following nearly a decade of promised disruption, automated vehicles at level 4 of the SAE J3016 classification (118) seem unlikely to operate securely in the face of general highway hazard scenarios for another decade, while full autonomous operation at level 5 is many decades away (119). Of course, the very notion of levels is suspect given the huge importance of the local cityscape, an environmental context whose vital characterization is in its very infant stages (120). Meanwhile, robots in still less structured settings do not deploy with much of any repeatable pattern: Successful applications result from elaborate, one-off, multihuman team exertion and still do not function in any predictable manner, failing regularly—or, worse, succeeding unexpectedly—from setting to setting (121).

Notwithstanding the accumulating multiple fatal accidents (117), physically embodied agents endowed with poorly understood, sloppily conceived, demonstrably dangerous partial autonomy are already being released into the human and natural environment by both commercial and state actors. While there is no dearth of similarly sloppy practice in the software industry, growing evidence suggests that the increasing power and practicability of formal methods are beginning to play an important role in at least life-critical applications (122). Robots, to the extent that they are useful at all, must be presumed to fall into this same category of life-critical applications. The demand-side pressure for such technology is rapidly growing, and many would-be customers will not want to delay the benefits of—much less impose a moratorium on—apparently useful machines no matter how imperfectly characterized. But there is presently no available formal methodology of correct robotics even to offer in case industry seeks it—or society comes to demand it.

## 4. HOW CAN WE GET IT?

The synthetic sciences are so young that adapting the right model for robotics will require both deliberate introspection and historical insight. Kinematics, the discipline focused on design of mechanical motion, was the first aspirant to synthetic science (123). It plays a role in robotics nearly as critical as computing but seems less instructive because it concerns behaviors whose architectures do not require an internal state, and seems never to have been concerned with physical limits. Cybernetics burst into mid-twentieth-century science with the proposal that emerging theories of information flow and its regulation could unify the study of animal nervous systems, computing machines, economies, electromechanical circuits, languages, psychopathologies, and social organizations (124). Under the weight of these breathtaking burdens—unsupportable even by the dazzling brilliance (and weakly armored emotions) of its proponent (125)—the field shattered into new engineering disciplines whose narrowed foci reflected different aspects of its origins in dynamical control theory (126). Embracing mathematical synthesis, these offspring—modern-day control, communications, and signal processing—fled the domain of synthetic science by rejecting the commitment to any specific physical setting. Synthetic biology is a fascinating younger cousin of robotics that is seemingly even less settled in its foundations. There remains the example of computer science.

From that perspective, the problems that robots face (and that humans must help them overcome) can be formulated in terms of a space of agent–environment states within which agent-initiated actions instigate transitions to new states toward some task-designated goal subspace (127). Uncertainty is rife: Models are wrong by intent of abstraction, and information is incomplete and noisy. Thus, inevitably, solutions—policies for progressing to goals

### THE CHALLENGE OF PROBLEM 3

Addressing Problem 3 is bound up in solving its antecedents: Progress in building and programming robots invites a speculative view of how to reason about their design.

#### Problem 1

Progress is being made using task–environment templates to reason about resource requirements (Section 4.1.1.1) and the transparency–dexterity trade-offs in deploying them (Section 4.1.1.2). Outstanding challenges include seeking fundamental limits to power, bit rate, and grip (Section 4.1.1.1) and formalizing principles of codesign for distributing the available resources across the body's interior and surface interface to the environment (Section 4.1.1.2).

#### Problem 2

Progress is being made by advances in the qualitative theory of robot hybrid dynamical systems (Section 4.1.2.1) arising from the consensus work-interface model (Section B.1) and the consequent prospects for grounded symbols (Section 4.1.2.2). Outstanding challenges include advancing and deepening the categorical account of template compositions (Section 4.1.3.1) and reworking them to ensure the empirical utility of the associated functional programming languages (Section 4.1.3.2).

#### Problem 3

A well-defined notion of task domain might be conceived as arising from the application of available grounded compositional operators (Section 4.1.3.1) to the available lexicon of grounded symbols (Section 4.1.2.2) via expressions allowed in the resulting programming language (Section 4.1.3.2). The prospects for defining and reasoning about types of environments seem to rest on further progress in modeling a robot's information interface (Section B.2.3).

(Section 4.1.2.2)—must be applied iteratively, requiring the agent to check its progress and replan according to the mismatch between anticipated and actual accomplishment, notwithstanding its noisy perception (Section 4.1.2.1). This iterated replanning view defines a closed loop that locks problem solving into the setting of dynamical systems theory (Section A). Section 4.1 brings this dynamical systems point of view to bear on the problems of robot body and program design, as summarized in the sidebar titled The Challenge of Problem 3. Section 4.2 takes an analogous view of the problems facing the birth of a new discipline.

## 4.1. Building and Testing Robots, Theories, and Programs

Speculating about approaches to Problem 3 necessitates some account of its antecedents. Section 4.1.1 addresses the problem of bodies (Problem 1), Section 4.1.2 proposes a theoretical framework that might yield programming languages of work (Problem 2), and Section 4.1.3 imagines what it might be like to actually use them in putting robots to task (Problem 3).

**4.1.1. Robots.** As commercially available robot technologies improve their capabilities to operate in more diverse environments, the disciplinary project of robot design and building comes into finer resolution. Researchers[18] must articulate what fundamental resources (e.g., power,

[18]Once again, time and space constraints restrict the scope of this article to the consideration of general-purpose robot architectures for general-purpose environments. For example, that restrictive scope entirely

information rate, or grip)—or perhaps some more carefully refined or newly identified fundamental limit not considered in Section 2.2.1—are better recruited or coordinated in their new designs, or else their contributions are more suited to evaluation by markets than by peer review.[19]

In that context, Problem 1 might be more carefully articulated as follows. Suppose a designer is given some material budget affording a sensorimotor and grip endowment with known and reliable scaling and distribution properties. Now, how do required tasks situated in discovered environments dictate a morphology and the distribution of power, bit rate, and gripping resources assigned it across space and time? Section 4.1.1.1 assesses the prospects for developing well-characterized appropriate robotic material resources, and Section 4.1.1.2 reviews what is known about how to distribute them.[20]

***4.1.1.1. Resources: power, information, grip, and their trade-offs.*** For the foreseeable future, roboticists must closely study the requirements of specific problem triples (Section 2.1.3) in order to design their robot's actuarium. This article makes a case in Section 4.1.3.1 for using templates (Section C.2.3.1) as the modeling framework for so doing. Recent work on projectile launch against gravity (131) presents an archetypal example of how to pursue such challenging analysis for task domains—here, single-shot leaps or hammering. Modeling the complex interplay between power constraints (speed–torque curve), compliance nonidealities (spring inertia), and grip limitations (latch geometry) to achieve launch energy over a range of environments (load mass) yields fundamental insight into what power train may be necessary or sufficient for what regions of this task domain (131). Analogously, the first vertical dynamical climbing robot (132) was achieved by insights from a bioinspired template (133), revealing the necessity for parallel springs to assist the available motor-specific power in supporting the machine's three-orders-of-magnitude increase in load (albeit with the simplification of assuming perfect grip) (134).

The need for grip seems even more complicated to characterize and trade off in the context of such task–environment pairings.[21] Considered as a metamaterial property, grip seems most carefully studied within the two-decade-old literature on synthetic dry adhesives, arguably initiated by the discovery of the van der Waals force mechanism underlying Gecko toe attachment and detachment (66). Principal figures of merit entail strength (typically measured in terms of surface energies) and reversibility (relevant quantities entailing time rate and magnitudes of load-to-preload and reversal ratios) (70). But numerous additional criteria, such as durability, propensity for fouling versus self-cleaning, and the difficulty of performing many of the relevant measurements, greatly complicate its physical characterization (70). Engineers' growing insight stemming from carefully informed bioinspiration (136–138) and improving materials and fabrication methods have spurred notable advances in rapidly reversible high-strength bonds through the ingenious arrangement of

---

ignores such crucial application areas as design for physical (128) or psychological (129) human–robot interaction.

[19]By the same token, robot companies that care about advancing the discipline underlying their technologies must be willing to expose to the broad research community (with suitable nondisclosure protections) hardware interfaces that permit testing of new theory.

[20]Another important, fascinating topic that lies beyond the scope of this article is the development of new materials for sensing. For example, given the huge role played by olfaction in the evolution of animal cognition (130), it is remarkable that robotics has not yet found a way to widely integrate some corresponding technology.

[21]Discussion of grip is a particularly illuminating setting for understanding that resource requirements can only be characterized with respect to a pairing of task and environment. For example, locomotion constrains animals' use of their remarkable grippers, while, simultaneously, animal gaits are known to vary dramatically depending upon the friction and adhesive properties of the substrate (135).

hierarchical mechanisms (139) exploiting anisotropic compliance (140). Correspondingly, the environment's affordance (141) of grip has commanded at least a comparable degree of attention regarding its native composition (142, 143) or design (144) as well as assessment, whether by remote anticipation (145) or direct proprioception (146, 147). It seems plausible that the urgently needed characterization of grip may be emerging with advances departing from these two opposite poles of the task–environment axis.

By contrast, equally central, yet almost entirely ignored within robotics, is the question of whether the trade-offs among grip, power, and information rate are fundamental or merely artifacts of presently available (or perhaps even poorly deployed) technology. Apparently, all three resources are bound up together and simultaneously coordinated as well as colimited in a robot's exploitation of any one. Greater power implies an ability to more rapidly and securely grasp and release using whatever adhesion a given object's surface affords. Reciprocally, substrates with a higher friction coefficient afford broader ground reaction force cones that increase the stance travel distance along which a given actuation power budget can add kinetic energy to a running body. Meanwhile, more timely contact information is required to increase the impulse that can be lent the body by the same power train and traction condition. Or, again reciprocally, the more secure the grip, the faster the proprioceptively gleaned information about an object's mass distribution. An urgent agenda for the discipline of robotics is to uncover the relationships among these resources and their ultimate physical colimitations (Section C.1.1).

Computer engineers' ability to trade energy against information is well established (19). Preliminary explorations of the nature and implications of analogous information–energy rate limits relevant to robotics are beginning to appear in the literature. Empirical observation suggests that mass-specific force (bonding strength) rather than power fundamentally limits actuator work rates (148). However, the constrained interaction between mechanical power and rapidly received, computed, and transmitted information has begun to be established at the mesoscale as affecting both output ports (76) and input channels (149).

***4.1.1.2. Distribution: compliance, modularity, and codesign.*** Compliance ideally is characterized by a memoryless force-extension function or, equivalently, a scalar-valued potential energy function. The term memoryless means that there is no internal state; hence, in the absence of any intrinsic time constants, ideal compliant elements incur no power limits and can support arbitrarily fast energy flows with no losses. Springs are good (150): Given the inevitable limits on actuation power, a designer is strongly motivated to introduce compliant elements in the body that can intermediate between the actuators' slow extraction of joules from the energy supply and the loads' fast time constants associated with the kinetic energy shifts required by the target mass states. However, the introduction of compliance typically incurs lowered information rates on the proprioceptive interfaces: diminished transparency (in their channels) and dexterity (through their ports).[22]

Channels suffer since, once coupled to an inertial load, the compliant elements dramatically alter the overall system time constants. The desire for transparency (the ability of an actuator to quickly and accurately read the loads' states) motivates the introduction of direct-drive (neither gearing nor compliance) robot technology (73), dramatically increasing the specific power, albeit

[22] Once again, limitations of space preclude a proper treatment in this article of so-called soft robotics—a popularizing term for the systematic introduction of tunable compliance afforded by recent advances in materials science and engineering.

at the cost of tricky trade-offs in specific force for locomotory systems (151). At the ports, series compliance has long been understood as offering enhanced output force accuracy (152) at the cost of severe bandwidth loss (153). Alternatively, parallel compliance can be used to amplify force magnitude, but only at specifically designated phases in the work loop (132), limiting the generality of tasks that can be performed.

The loss of transparency (accuracy–bandwidth trade-offs in proprioceptive channels) and dexterity (accuracy–bandwidth or magnitude–timing trade-offs through proprioceptive ports) motivates the consideration of sensorimotor specialization, whereby, for example, compliance can be associated with high-power actuation at the body core (i.e., proximal to the mass center) (154), where neither dexterity nor state information is crucial, whereas highly dexterous, sensitive, lowered power actuation can be placed at the periphery (i.e., distal to the mass center), where the body meets the environment (155). By contrast, the appeal of composable resources has motivated proposals for the programmability (156) and reconfigurability (157) of modules, some originating in the familiar traditions of mechatronics (158–160) and others from microelectromechanical systems (161) or materials science (71). The hierarchical, multiscale nature of biological morphology confers a decisive advantage on engineering designs clever enough to achieve it (138), heightening the challenge of finding simple recipes for composition of either form or function.

The general question of how to rationally distribute a robot's physical resources to perform a set of tasks in a class of environments has come to be called the problem of codesign (Section C.1.2). A more constrained version of this problem inspired by the common recourse to preflexes (162, 163) (Section B.2.1) observed in animals seeks to design what has come to be called morphological computation (Section C.1.2). Accelerating progress in posing and solving such design problems gives the promise of advancing the field's insight into how to pose Problem 3.

**4.1.2. Theories.** The mechanics of work expressed by the first equation of Section B.1.1 imbues any reasoning about robotics with the study of dynamics. Embracing the iterated version of Newell & Simon's (127) general problem-solving formulation further implies closed-loop dynamics that result from feedback.[23] Absent a general information interface model (Section B.2.1), it is now convenient to assume that this feedback takes the form of assigning to actuators some function of the entire history of sensor readings.

These assumptions bring to bear the theory of dynamical systems (Section A) and, thereby, the tools of topology that hold a relationship to robotics roughly analogous to that exhibited by logic relative to computer science. Originating in Poincaré's investigation of celestial mechanics (164), topology engages robotics through dynamics (Section B.1) to present an intrinsic, robust account of uncertainty and cost (Section 4.1.2.1)[24] and offer a formal characterization of grounded symbols (Section 4.1.2.2).

[23]For example, the distinction between deliberative and reactive planning that is so firmly entrenched in the idioms of contemporary robotics blurs when one considers that reactive policies must be constructed in advance of their execution, while deliberative policies will inevitably be iterated in some fashion. At best, one might expect that these categories describe a relationship wherein policies seem deliberative to other policies at a finer spatiotemporal scale that they call out and reactive to spatiotemporally broader policies that call them in turn.

[24]Limitations of space and insight restrict this article's focus to the work exchanged between a robot and its environment, as specified in Section B.1. This comes at the expense of more elaborated uncertainty and cost models whose careful representation awaits the development of an information interface specification (Section B.2). Meanwhile, of course, methods of both stochastic filtering and control (110) and optimization (165) are, justifiably, deeply ingrained in robotics practice. Their formal integration will unquestionably be essential to an eventual discipline. Brief mention of the optimality point of view is made in Section 4.1.2.1.

*4.1.2.1. Intrinsic models of uncertainty and cost.* Uncertainty in the models and measurements underlying a robot's interface to the environment motivates the consideration of chains (166)—solutions of dynamical systems with arbitrary, small, but arbitrarily persistent errors (Section A.1). From the foundational view, this completely unstructured model of uncertainty is very attractive: It is intrinsic in the sense that there is no requirement for further models of information (Section B.2.1). Happily, the formal notion of a chain can be extended to the hybrid setting (167), and with it a version of steady-state behavior (168) for at least a large subclass of physically practicable but well-behaved versions of the work-interface model (Section B.1) derived from Reference 169. However, usefully more structured representations of uncertainty, such as parameterized families of models, lead to differential inclusions that can also be incorporated in useful hybrid systems models (170). It seems urgent to establish which empirically effective robotics models do or do not admit what version of chains with accompanying guarantees of well-behaved steady state.[25]

Those guarantees include the existence of an intrinsic scalar-valued Lyapunov function down which flows must decrease along the way to their steady-state attracting sets (Section A.2). While there is no canonical choice of such functions in general, when carried over to the setting of hybrid robot dynamics (Section B.1), it is appropriate to assume that they will be closely related to the physical total energy (171).[c] This article takes the speculative step of simply presuming that such a natural energy-like function is available to play the role of whatever version of a Lyapunov function the particular dynamics will afford (Section A.2). Then, in the smooth case (Section A.3), its derivative along the system's motions will yield an expression of mechanical power (instantaneous energy expenditure). Henceforth, it is convenient to simply refer to the various versions of these scalar-valued functions as energy landscapes and their scalar-valued descent rate functions as power landscapes.

*4.1.2.2. From goal primitives to tasks via grounded symbols.* Models of mechanical behavior (Section B.1.1) describe the interaction of spatiotemporally continuous, and hence uncountable, quantities (energy and information flows), but languages are defined over countable alphabets. With the tools and concepts of the previous discussion (Section 4.1.2.1) in place, it becomes possible to address the essential barrier to the agenda of Problem 3: the gulf between signals and symbols (Section C.2.1). One concludes that the valleys of the energy landscape—the basins of attracting invariant manifolds perceptually marked by their energy sublevel sets[n]—are viable candidates for physically grounded symbols. Happily, these symbolic goal primitives also come ready-equipped with actual task specifications, as follows.

Historically, robotic tasks were confined to motion planning in presumed known workspaces—initially Euclidean spaces, necessitating only robot kinematics models (172), and subsequently rigid placements in spaces punctured by fixed obstacles (173). Thus, path planning was the first algorithmically posed problem of robotics and remains a central focus of the field (174). The computational complexity of a deterministic solution must grow exponentially in the degrees of freedom (175), motivating a shift toward sampling-based formulations (176) that can only be probabilistically complete. Yet finding a free path between a given initial–final pair of configurations is essentially a topological problem whose complexity can be alternatively quantified

[25]Here and throughout the article, steady state is a convenient but potentially misleading term for the robust, long-term behaviors exhibited by dynamical systems. Section A.1 offers a brief sketch of the powerful theory establishing the emergence and persistence of such attracting sets, whose complicated spatiotemporal structure lies far beyond the tidy equilibrium or oscillatory behaviors that the colloquial use of the term might intuitively connote.

via the cardinality of a lexicon over the space of pairs each of whose symbols is a continuous path planning function (177). The unavoidable need for repeated replanning now enforces a reformulation in dynamical terms: iterated maps or vector field flows that bring (almost) any initial choice to a desired final configuration (178).

This reformulation was first expressed in terms of artificial potential energy (179) and subsequently shown in principle to afford almost global dynamical solutions to any motion planning problem (180).[26] The sequential composition of artificial energy basins (184) (Section C.2.3.1) was anticipated by the notion of preimage back-chaining (185) inherited from AI (186). The proposal to use artificial potential functions as a specification of the effective impedance that a robot should present to the environment (187) represents an important parallel starting point in the agenda to encode goal primitives via reference dynamics rather than reference motions (Section C.2.2).[27] These ideas can be extended to algorithmically generate artificial potential functions for a large variety of robot dynamics models of the sort modeled by the first equation of Section B.1.1 (189).

Because the combinatorial complexity of symbol manipulation can be dramatically reduced by appeal to the algebra of basin adjacency revealed by a vector field planner (190),[g] this approach to task planning seems worth embracing even from the perspective of computational efficiency alone. In this view, it seems urgent to work out the topological perplexity (182) (see footnote 26) of task domains along lines parallel to the ongoing progress in characterizing topological complexity (191). More broadly, the adjacency or disconnection of basins reveals intrinsic "spelling" rules for transitional tasks arising from specific environmental affordances (94, 192, 193) or task relationships (194) in a manner touched upon in Section B.2.2.

The question now arises of how to ground such planning vector fields in the dynamics of a working robot given by the first equation of Section B.1.1. Conceived literally in the manner of their origins as artificial potential fields (179), the goals and antigoals[m] of the associated quasi-static gradient field can be immediately rendered as dynamically grounded symbols by recourse to Lord Kelvin's insight that dissipative second-order systems asymptotically reach the minima of their potential energy (171, 195). However, the motion planning literature has appropriately emphasized the importance of kinodynamic plans (196) entailing trajectories with specifically tailored transient properties. Analogously tailored properties can be imposed upon planning fields, and those desirable transient properties can be closely approximated by the appropriately compensated working robot dynamics (197–199), as illustrated in Section C.2.2.3.

Supplemental Material >

The **Supplemental Appendix** also briefly addresses the alternative of generating reference dynamics indirectly from a cost function (200) using methods of optimal control (Section C.2.2.2). A preference for task specification via direct reference dynamics emerges from a number of considerations,[o] but the driving motivation stems from the conclusion that reference dynamics

[26]Here, the term almost global means that the defect (i.e., the complement of the basin of the attracting final destination) has an empty interior in the configuration space. In typical application settings (indeed, barring pathologically wild cases; M.D. Kvalheim, manuscript in preparation), basins have the same homotopy type as their attracting sets (181). Hence, single-point goals typically have contractible basins (topological disks) that cannot cover the entirety of noncontractible configuration spaces. This motivates the problem of determining how few basins must be required to do so, as exemplified in Reference 182, coined as topological perplexity in Reference 183.

[27]This idea has grown to be influential enough in animal motor science to have achieved clinical application in human rehabilitation therapies (188). See footnote 30 for references to the animal motor science literature that further extend the concept into the composition of motor primitives in a manner analogous to that traced for robotics in Section 4.1.3.1.

better serve the purposes of composition to be developed below in Section 4.1.3. Broadly speaking, the danger of appeal to optimal methods at any one level of a deep specification hierarchy is their parochial nature.[28]

### 4.1.3. Programming.

This section considers the prospects for a functional programming language at the most basic level of a robot's interface to the environment. Section 4.1.3.1 discusses the availability of grounded compositions for the grounded task specification primitives just introduced (Section 4.1.2.2). Section 4.1.3.2 explores the availability of a type theory capable of treating such compositions as formal combinators (29) whose evaluation has an operational semantics in the formal properties of the task specifications so composed.[29]

#### *4.1.3.1. Composition of task specifications.*

A far wider realm of tasks than the motion planning setting of Section 4.1.2.2 seems to require specification not merely in terms of set-theoretic goals but also with respect to reference dynamics. Many examples from biology suggest that animals solve the degrees-of-freedom problem (202) (Section C.1.1) by using low-dimensional abstract templates (203) (Section C.2.2). More broadly still, chain-grounded goal symbols (Section 4.1.2.2) represent intrinsically steady-state behavior, whereas much of a robot's work can be expected to entail transitional maneuvers. Thus, at the very least, the agenda of programming work seems to require a syntax for building up specifications of behaviors entailing compositions of dynamical primitives. To require that these compositions be grounded is to specify which formal properties of constituent task primitives are inherited by the results of the composition.

Construing the construction and embedding of reference dynamics (Section C.2.2.3) as a hierarchical composition (Section C.2.3.1) instantiates the concept of an anchored template (203)—a module of behavior that can be represented and composed via a symbol grounded in a physically embodied sensorimotor behavior (Section 4.1.2.2). Here, the guaranteed property of the composition is that the behavior in the resulting higher-dimensional anchoring space converges toward a lower-dimensional subspace whose dynamics is a change of coordinates away from that of the template (204, 205).

From the perspective of the agenda pursued by this article, Raibert's hoppers (113) contributed the most important of any advances to twentieth-century robotics. On the one hand, their reliance on dynamical equilibria (stable limit cycles rather than mere point attractors) underscored the primary role of energy in robot tasks. On the other hand, they pioneered empirical candidates for parallel composition (Section C.2.3.1). From a purely formal point of view, there is nothing simpler: The parallel composition of two functions simply takes as domain and codomain the Cartesian products of the individuals' corresponding sets and evaluates them independently. A suitably elaborated version of this idea can be used to define products over the hybrid dynamics category (167). However, in robotics, as in any setting of mechanical circuitry (39), conjoining two physical systems inevitably entails cross talk. An urgent problem is to express more relaxed (206)

[28]Consider the succinct assessment "premature optimization is the root of all evil (or at least most of it) in programming" (3, p. 671). The problem is similarly eloquently addressed from a historical perspective in the context of the opposition to compositional methods evinced by layer-level optimality specialists on the eve of the VLSI revolution (201).

[29]The reader will observe that this is a far cry from the deep layers of specification languages urged in Section 2.2.2. In response, the author will again plead the immaturity of a young field against the considerable challenges outlined there. Potentially, the language emerging at this level of direct environmental interface will come to be seen as analogous to the assembly languages of computers, upon which much more abstract and useful—but still grounded—languages can be built (Section C.2.3.2).

parallel compositions that can distinguish safe from inimical cross talk (207, 208) in a categorical setting (Section C.2.3.1). Here, the grounding requirement is that the steady-state dynamics of the product system are guaranteed to be included in the product of the constituents' steady-state dynamics.

A version of sequential composition using Lyapunov sublevel sets (184) has been expressed categorically as well (167). However, the present working version is defined only for basins associated with the attractors arising from chain analysis rather than basins associated with hyperbolic attracting sets established via smooth analysis, exemplified in **Figure 1** in the **Supplemental Appendix** (Section C.2.3.1) and the powerful, exquisitely detailed applications (209) of bifurcation theory (210) that follow. This is a major motivation to understand how and under what conditions the hybrid dynamics category of Reference 167 can be refined to work with higher-resolution (but more rigid) hyperbolic attracting manifolds (Section A.3). Here, the grounding requirement is that the steady-state dynamics of the paired sequence are included in the steady-state dynamics of the second system of the pair. Enforcing this property using the sublevel sets of the associated energy landscapes yields an effective means of stringing together highly energetic transitional maneuvers (211).

Supplemental Material >

***4.1.3.2. Grounded type theory.*** The compositional framework just discussed now invites a reprise of the conceptual passage from modules to functional programs briefly highlighted in Section 2.2.2, suggesting a transplant from computer science into an analogous development for robotics (Section C.2.3).[30] Crucially, the purpose of modules is to be reused in varied compositions, promoting the construction of more complex behaviors from simpler constituents with known, reliable properties. As hinted just above, this amounts to the requirement for a category-theoretic (216) treatment of the primitives and their compositions. That requirement arises in the form of an unavoidable link whereby types emerge intrinsically from and at once intrinsically define categories (217). In other words, the physical grounding of the type theory is established through its category-theoretic model.

Type theories for robotics have long been proposed, but their associated categories typically remain unspoken and hence ungrounded. For example, the signals of functional reactive programming (218) take their domain in the reals with an unspecified (apparently arbitrary) codomain; therefore, the combinators applied to compose their user-accessible signal functions have no physically specified operational semantics. In consequence, the applicability of their appealingly broad, elegant type theories (219) to any particular robot architecture in any class of environments is indeterminate.

Conversely, a long-developing literature on physically grounded symbolic specification of robot motion control compositions (220–222) yields a version that presents as a context-free grammar (223). Locating the complexity of a specification language in the Chomsky hierarchy is surely important for managing trade-offs between ease of expression and cost of evaluation. However, it is not a substitute for characterizing the behaviors specified. Absent some categorical analysis of the mathematical objects constructed by such grammars, one guesses that the operational semantics are roughly comparable to those of arbitrary hybrid systems, whose most interesting qualitative properties are typically undecidable (224).

A contrasting approach to behavioral specification obtains from appeal to purely linguistic representation. For the price of translating the description of coupled hybrid robot–environment

[30]A similarly analogous conceptual passage may be detected in animal motor science, which has progressed from identifying muscle synergies (212) and then multiple modules (213) to forming compositional hypotheses (214) and, ultimately, to the prospect of their clinical application (215).

dynamics given by the first equation of Section B.1.1 into modal logic, recourse to model checking yields computationally effective behavioral verification (i.e., correctness guarantees or failure diagnostics) of plans and controllers written out in the same syntax (225). Moving up the Chomsky hierarchy, an analogous verification tool has been developed for models, tasks, and policies written out in a more expressive but computationally costlier context-free grammar (226). Formal interfaces to structured natural language (227) and probabilistic interfaces to human natural language (228) further enrich the ease of expression. A substantial advance in such approaches to robot programming is an explicit representation of the entire problem triple: architecture, task, and environment. A central challenge is the grounding gap—the computational complexity (229) and rigidity (230)—facing any linguistic representation of the physical world (231).

Eminently practicable functional programming environments have been developed for popular contemporary robot operating systems (232), and increasingly powerful autonomous high-level task planners have been built with more formal versions of such tools (233). It remains to exercise them with physically grounded categorically generated type theories (Section C.2.3.2). Here arises a new challenge to ensure that the formalism serves the purposes of safe and expressive task specification rather than generating sterile, impracticable nostrums. This can be achieved only by a constant interplay between the hypotheses of the designs and the refutations of their empirical examination.

## 4.2. Building and Deepening Collaborations

This article has focused on the need to reach more deeply into the physical, mathematical, and teleological foundations that underlie robotics. It seems fitting to end with a glimpse of the need to reach out. Accordingly, Section 4.2.1 outlines the importance to robotics of its interdisciplinary relationships (including the importance of establishing an explicit departmental identity) within the university, and Section 4.2.2 urges using that platform to lead the charge for greater equity and diversity in science and society.

**4.2.1. Institutions.** The foregoing discussions have amply rehearsed the crucial role that other disciplines have to play in robotics. Indeed, today's prevailing conception of the field seems to be that of a technological breeding ground where electrical, mechanical, and materials engineers can all collaborate with computer scientists to spawn wonderful contraptions. Surely, it is clear to these other disciplines how valuable a role robotics can play in motivating their favorite specialties. For reasons touched upon in Section 3.2, absent an explicit departmental identity in the universities, the affection of other units does not necessarily advance the field.

It is striking to revisit Gorn's (234) proposal from more than half a century ago for a discipline called "the computer and information sciences." That tripartite argument prescribed the intellectual focus (mechanical languages), characterized the boundaries by listing its adjacent disciplines (electrical engineering, linguistics, mathematics, philosophy, and psychology), and strategized the politics of departmental startup (escape from a neighboring department fueled by demand for its service courses).[31] As just remarked, this article focuses nearly exclusively on the first concern. Gorn's second line of argument seems deserving of its own independent treatment at a time closer to the widespread establishment of robotics departments. There are many other adjacent and more distal disciplines that both contribute to and benefit from robotics. Beyond the obvious adjacent

[31]It appears that this third argument, Gorn's wise appeal to service courses as the basis for a declaration of independence from colonializing neighbors, remains a tantalizing fantasy, awaiting greater penetration of robot technology into society.

disciplines of computer science, engineering, and mathematics, throughout this article there has been constant mention of the vital interplay with biology.[32] There are similarly deep connections to be established with history and philosophy. The reach of robotics clearly extends far deeper into the humanities, encountering the arts as well. All of these connections will be needed to establish the discipline of robotics.

**4.2.2. People.** Attracting the best and brightest young minds is of course the most essential driver of any field. But contemporary robotics, drawing its researchers largely from those established disciplines most notoriously homogeneous with respect to gender, race, and ethnicity, will continue to suffer a consequently diminished pace of innovation (235) until more strenuous effort achieves greater diversity. The enduring, paradoxical (236), and pernicious (237) nature of inequitable access[33] has been documented at every step (239) along the leaky STEM pipeline, including disparities in mentorship even at the doctoral level (240), reaching increasingly disproportionate numbers at successively higher ranks (241). Accumulating evidence that neither specific public school redesign (242) nor even recourse to competition in eliciting it (243) has substantially broadened access to achievement prompts a growing chorus seeking newly vigorous, intentional recruitment of existing social (244) and legal (245) structures to break apart the many interlocking obstacles lying beyond the reach of mere educational reform.

But robotics portends a unique social impact that the discipline must embrace and deliver for human benefit. On the one hand, urgent moral issues (246) compounding dramatic 50-year declines in manufacturing jobs and the [seemingly consequent (247)] decoupling of income from GDP growth (248) oblige the creators of automation technologies to recognize their implicit policy concomitants (249) and become more thoughtful about their social impacts (250).[34] On the other hand, the technology is irresistibly fascinating. Thus, of all disciplines, robotics is charmed in its visceral appeal, as evidenced by millennia of human dreams and nightmares about the inanimate made vital. Robotics researchers owe it to their field as well as their society to leverage that popular fascination in recruitment, retention, and promotion to secure the participation of presently underrepresented groups. Beyond the promise of intellectual advance, we will need the broadest possible diversity of disciplinary experts to collectively balance the benefits of automation against the future of human work while ensuring that our burgeoning technological prowess works to counter social injustice (253) and spread those benefits across cultures and classes (254).

## 5. CONCLUSION

The Future Issues list below summarizes the overarching next steps that this article proposes as particularly urgent to advance the foundations of robotics. But the article is first of all offered

[32]A great deal of confusion concerns the relationship of animals to robots and the role their juxtaposition has to play in advancing both sciences. A proper commentary would require at least a dedicated essay of comparable length to this one.

[33]These inequities have been carefully studied largely with respect to the US domestic population. The powerful intellectual advantages conferred by [now imperiled (238)] historical US access to the elite of international STEM populations and the innovation-pumping cultural diversity they bring does not mitigate the damage incurred by missing out on the underserved local talent.

[34]Other observers emphasize the centuries-long experience of technology-driven job disruption leading to greater productivity and accelerated creation of new jobs (251). Nevertheless, social rejection might still be triggered not by a wholesale retreat of employment but rather by an inability to educate the workforce to track the rising high-end demand (252).

in the hope of helping new researchers hone their proposals by exploiting alignments, exploring connections, or exposing fallacies introduced by the observations and arguments made along the way. Even better would be if some pronouncement here influences the introduction—and best of all, imaginably, the conclusion—of someone's next scientific paper to support or refute one of those many opinions. For such readers, the **Supplemental Appendix** may hopefully offer much more usefully specific targets of attack, and the Summary Points list that leads it off is intended to outline the connections between the conceptual questions raised here and the technical machinery involved in addressing them. More broadly, this article will have achieved a large part of its purpose if it spurs other researchers, young and old, to articulate coherent foundations of robotics that improve upon, contest, or outright reject and replace this account in a manner that better promotes actionable fundamental research. In the long run, a synthetic science of robotics will emerge anyway, and it may be at least historically interesting to look back at the concerns that attracted and the confusions that bedeviled one toiling denizen of its predisciplinary era.

Supplemental Material >

## FUTURE ISSUES

1. Scientific foundations necessary to advance the capabilities and safety of emerging robot technologies will require the identification of fundamental physical limits as well as models that characterize types of environments in which types of machines can be expected to operate appropriately.
2. Intellectual advances necessary to undergird those technologies with these foundations require a constant interplay between hypothesized type theories for assigning tasks to architectures in environments and empirical study of physical machines in uncurated settings that can support such theories or refute them.
3. Disciplinary developments necessary to foster such intellectual advances include the creation of robotics departments capable of systematic collaboration with adjacent disciplines and populated by the greatest diversity of thinkers that the age-old human fascination with robots promises to achieve.

## DISCLOSURE STATEMENT

As the text indicates, this review reflects the author's personal view of the entire field of robotics—not necessarily as it exists today but as it might come into clearer form as an intellectual discipline. The author asserts that the discussion of all articles cited is accurate, fair, and unbiased and is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of their treatment in this review.

## ACKNOWLEDGMENTS

The time and thought necessary to complete this article were underwritten by Office of Naval Research grant N00014-16-1-2817, awarded as a Vannevar Bush Faculty Fellowship sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering. The material was originally conceived and delivered as the 2011 Henry Paynter Lecture to the faculty of the Department of Mechanical Engineering at the Massachusetts Institute of Technology. I am extremely grateful to Professor Hari Asada and his MIT colleagues, who invited the address and provided lively feedback throughout the day and evening (and, indeed, over the course of many

discussions before and since) that significantly advanced my thinking about these ideas. A next chance to develop the themes was occasioned by my delivering the 2012 David Goldman Lecture to the faculty of the Technion Autonomous Systems Program at the Israel Institute of Technology. I thank Professor Elon Rimon and his Technion colleagues for their kind hosting throughout another very stimulating week (and, again, many years prior and subsequent) of exchanges related to these ideas.

Those having even minimal familiarity with my work will already know that the source of any good ideas to be found here or elsewhere in my writing lies in the individual lessons taught to me and the collective culture created around me by the amazing group of PhD students and postdoctoral fellows I have had the wonderful good fortune to grow up with over the course of a nearly four-decade career. Many of my phenomenal collaborators (students, peers, and elders alike) are responsible for specific technical insights discussed here, and I have tried to acknowledge this by appropriate citation.[35] My teacher, Professor K.S. Narendra, taught me how to think.

## LITERATURE CITED


1. Simon HA. 2019. *The Sciences of the Artificial*. Cambridge, MA: MIT Press
2. Schmidt M. 2010. Do I understand what I can create? In *Synthetic Biology: The Technoscience and Its Societal Consequences*, ed. M Schmidt, A Kelle, A Ganguli-Mitra, H Vriend, pp. 81–100. Dordrecht, Neth.: Springer
3. Knuth DE. 1974. Computer programming as an art. *Commun. ACM* 17(12):667–73
4. Milner R. 1993. Elements of interaction: Turing Award lecture. *Commun. ACM* 36(1):78–89
5. Harper R. 2013. *Practical Foundations for Programming Languages*. Cambridge, UK: Cambridge Univ. Press
6. Puliafito C, Mingozzi E, Longo F, Puliafito A, Rana O. 2019. Fog computing for the Internet of Things: a survey. *ACM Trans. Internet Technol.* 19:18
7. Shalf J. 2020. The future of computing beyond Moore's law. *Philos. Trans. R. Soc. A* 378:20190061
8. Hopcroft JE, Ullman JD. 1979. *Introduction to Automata Theory, Languages and Computation*. Reading, MA: Addison-Wesley
9. Fein L. 1959. The role of the university in computers, data processing, and related fields. *Commun. ACM* 2(9):7–14
10. Denning PJ. 2007. Computing is a natural science. *Commun. ACM* 50(7):13–18
11. Landauer R. 1961. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* 5:183–91
12. Markov IL. 2014. Limits on fundamental limits to computation. *Nature* 512:147–54
13. Falkoff AD, Iverson KE, Sussenguth EH. 1964. A formal description of SYSTEM/360. *IBM Syst. J.* 3:198–261
14. Moore GE. 2006. Cramming more components onto integrated circuits, reprinted from Electronics, vol. 38, number 8, April 19, 1965, pp. 114 ff. *IEEE Solid-State Circuits Soc. Newsl.* 11(3):33–35
15. Dennard RH, Rideout VL, Bassous E, Leblanc AR. 1974. Design of ion-implanted MOSFET's with very small physical dimensions. *IEEE J. Solid-State Circuits* 9:256–68
16. Bohr M. 2007. A 30 year retrospective on Dennard's MOSFET scaling paper. *IEEE Solid-State Circuits Soc. Newsl.* 12(1):11–13
17. Krishnan S, Garimella SV, Chrysler G, Mahajan R. 2007. Towards a thermal Moore's law. *IEEE Trans. Adv. Packag.* 30:462–74
18. Pop E. 2010. Energy dissipation and transport in nanoscale devices. *Nano Res.* 3:147–69


[35] I am particularly grateful to Sam Burden, Jean-Luc Cambier, Paul Gustafson, Matt Kvalheim, Matt Mason, Lisa Miracchi, and Vassilis Vasilopoulos for detailed comments and advice that substantially improved the quality of this article.

19. DeHon A. 2015. Fundamental underpinnings of reconfigurable computing architectures. *Proc. IEEE* 103:355–78
20. Theis TN, Wong HSP. 2017. The end of Moore's law: a new beginning for information technology. *Comput. Sci. Eng.* 19:41–50
21. Bennett C. 1973. Logical reversibility of computation. *IBM J. Res. Dev.* 17:525–32
22. Carlton DB, Lambson B, Scholl A, Young AT, Dhuey SD, et al. 2011. Computing in thermal equilibrium with dipole-coupled nanomagnets. *IEEE Trans. Nanotechnol.* 10:1401–4
23. Mead C, Conway L. 1980. *Introduction to VLSI Systems*. Reading, MA: Addison-Wesley
24. Bell CG, Newell A. 1971. *Computer Structures: Readings and Examples*. New York: McGraw-Hill
25. McCarthy J. 1960. Recursive functions of symbolic expressions and their computation by machine, part I. *Commun. ACM* 3(4):184–95
26. McCarthy J. 1962. *A basis for a mathematical theory of computation*. Memo 31, Artif. Intell. Proj., Mass. Inst. Technol., Cambridge, MA
27. Martin-Lof P. 1984. Constructive mathematics and computer programming. *Philos. Trans. R. Soc. A* 312:501–18
28. Wood AK. 2012. *Warships of the Ancient World*. Oxford, UK: Osprey
29. Pierce BC. 2002. *Types and Programming Languages*. Cambridge, MA: MIT Press
30. Appel AW, Beringer L, Chlipala A, Pierce BC, Shao Z, et al. 2017. Position paper: the science of deep specification. *Philos. Trans. R. Soc. A* 375:20160331
31. McCarthy J. 1981. History of Lisp. In *History of Programming Languages*, ed. RL Wexelblat, pp. 173–85. New York: Academic
32. Cohen PR. 1995. *Empirical Methods for Artificial Intelligence*. Cambridge, MA: MIT Press
33. Russell SJ, Norvig P. 2009. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice Hall
34. Reynolds M, Cortese A, Liu Q, Wang W, Cao M, et al. 2020. Surface electrochemical actuators for micron-scale fluid pumping and autonomous swimming. *Bull. Am. Phys. Soc.* In press
35. Oishi K, Klavins E. 2014. Framework for engineering finite state machines in gene regulatory networks. *ACS Synthet. Biol.* 3:652–65
36. Reverdy PB, Vasilopoulos V, Koditschek DE. 2020. Motivation dynamics for autonomous composition of navigation tasks. *IEEE Trans. Robot.* In press
37. Miracchi L. 2019. A competence framework for artificial intelligence research. *Philos. Psychol.* 32:588–633
38. Robot. Ind. Assoc. 2020. Joseph Engelberger: the Father of Robotics. *Robotic Industries Association*. **https://www.robotics.org/joseph-engelberger/about.cfm**
39. Whitney DE. 1996. Why mechanical design cannot be like VLSI design. *Res. Eng. Des.* 8:125–38
40. Mead C. 1989. *Analog and Neural Systems*. Boston: Addison-Wesley
41. Schuman CD, Potok TE, Patton RM, Birdwell JD, Dean ME, et al. 2017. A survey of neuromorphic computing and neural networks in hardware. arXiv:1705.06963 [cs.NE]
42. Achour S, Sarpeshkar R, Rinard M. 2016. Configuration synthesis for programmable analog devices with Arco. In *PLDI '16: Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 177–93. New York: ACM
43. Mirvakili SM, Hunter IW. 2018. Artificial muscles: mechanisms, applications, and challenges. *Adv. Mater.* 30:1704407
44. Meijer K, Bar-Cohen Y, Full RJ. 2003. Biological inspiration for musclelike actuators of robots. In *Biologically Inspired Intelligent Robots*, ed. Y Bar-Cohen, CL Breazeal, pp. 26–45. Bellingham, WA: Soc. Photo-Opt. Instrum. Eng.
45. Full RJ. 1989. Mechanics and energetics of terrestrial locomotion: bipeds to polypeds. In *Energy Transformations in Cells and Animals*, ed. W Wieser, E Gnaiger, pp. 175–82. New York: Thieme
46. Hunter IW, Hollerbach JM, Ballantyne J. 1992. A comparative analysis of actuator technologies for robotics. *Robot. Rev.* 2:299–342
47. Zhang J, Sheng J, O'Neill CT, Walsh CJ, Wood RJ, et al. 2019. Robotic artificial muscles: current progress and future perspectives. *IEEE Trans. Robot.* 35:761–81
48. Spenko MJ, Saunders JA, Haynes GC, Cutkosky MR, Rizzi AA, et al. 2008. Biologically inspired climbing with a hexapedal robot. *J. Field Robot.* 25:223–42

49. Ruina A. 1998. Nonholonomic stability aspects of piecewise holonomic systems. *Rep. Math. Phys.* 42:91–100
50. Kubow TM, Full RJ. 1999. The role of the mechanical system in control: a hypothesis of self-stabilization in hexapedal runners. *Philos. Trans. R. Soc. B* 354:849–61
51. Holmes P, Full RJ, Koditschek DE, Guckenheimer J. 2006. The dynamics of legged locomotion: models, analyses, and challenges. *SIAM Rev.* 48:207–304
52. McGeer T. 1990. Passive dynamic walking. *Int. J. Robot. Res.* 9:62–82
53. Collins SH, Wisse M, Ruina A. 2001. A three-dimensional passive-dynamic walking robot with two legs and knees. *Int. J. Robot. Res.* 20:607–15
54. Schmitt J, Holmes P. 2000. Mechanical models for insect locomotion: dynamics and stability in the horizontal plane I. *Theory Biol. Cybernet.* 83:501–15
55. Parrondo JMR, Horowitz JM, Sagawa T. 2015. Thermodynamics of information. *Nat. Phys.* 11:131–39
56. Debiossac M, Grass D, Alonso JJ, Lutz E, Kiesel N. 2020. Thermodynamics of continuous non-Markovian feedback control. *Nat. Commun.* 11:1360
57. Messler RW. 1993. *Joining of Advanced Materials*. Boston: Butterworth-Heinemann
58. Kadic M, Milton GW, van Hecke M, Wegener M. 2019. 3D metamaterials. *Nat. Rev. Phys.* 1:198–210
59. Sussman DM, Cho Y, Castle T, Gong X, Jung E, et al. 2015. Algorithmic lattice kirigami: a route to pluripotent materials. *PNAS* 112:7449–53
60. Yuan H, Pikul J, Sung C. 2018. Programmable 3-D surfaces using origami tessellations. In *Origami$^7$: Proceedings of the 7th International Meeting on Origami in Science, Mathematics, and Education*, Vol. 3: *Engineering One*, ed. RJ Lang, M Bolitho, Z You, pp. 893–906. St. Albans, UK: Tarquin
61. Guseinov R, McMahan C, Pérez J, Daraio C, Bickel B. 2020. Programming temporal morphing of self-actuated shells. *Nat. Commun.* 11:237
62. Chan TS, Carlson A. 2019. Physics of adhesive organs in animals. *Eur. Phys. J. Spec. Top.* 227:2501–12
63. Mason MT. 2018. Toward robotic manipulation. *Annu. Rev. Control Robot. Auton. Syst.* 1:1–28
64. Cutkosky MR, Wright PK. 1986. Friction, stability and the design of robotic fingers. *Int. J. Robot. Res.* 5:20–37
65. Tian Y, Pesika N, Zeng H, Rosenberg K, Zhao B, et al. 2006. Adhesion and friction in gecko toe attachment and detachment. *PNAS* 103:19320–25
66. Autumn K, Sitti M, Liang YA, Peattie AM, Hansen WR, et al. 2002. Evidence for van der Waals adhesion in gecko setae. *PNAS* 99:12252–56
67. Murray RM, Li Z, Sastry SS. 1994. *A Mathematical Introduction to Robotic Manipulation*. Boca Raton, FL: CRC
68. Johnson A, Koditschek D. 2013. Legged self-manipulation. *IEEE Access* 1:310–34
69. Milner R. 1999. *Communicating and Mobile Systems: The Pi Calculus*. Cambridge, UK: Cambridge Univ. Press
70. Eisenhaure J, Kim S. 2017. A review of the state of dry adhesives: biomimetic structures and the alternative designs they inspire. *Micromachines* 8:125
71. McEvoy MA, Correll N. 2015. Materials that couple sensing, actuation, computation, and communication. *Science* 347:1261689
72. Gholipour B, Bastock P, Craig C, Khan K, Hewak D, Soci C. 2015. Amorphous metal-sulphide microfibers enable photonic synapses for brain-like computing. *Adv. Opt. Mater.* 3:635–41
73. Asada H, Youcef-Toumi K. 1987. *Direct-Drive Robots: Theory and Practice*. Cambridge, MA: MIT Press
74. Kim S, Laschi C, Trimmer B. 2013. Soft robotics: a bioinspired evolution in robotics. *Trends Biotechnol.* 31:287–94
75. Gregorio P, Ahmadi M, Buehler M. 1997. Design, control, and energetics of an electrically actuated legged robot. *IEEE Trans. Syst. Man Cybernet. B* 27:626–34
76. Piccoli M, Yim M. 2015. Anticogging: torque ripple suppression, modeling, and parameter selection. *Int. J. Robot. Res.* 35:148–60
77. De A, Stewart-Height A, Koditschek D. 2019. Task-based control and design of a BLDC actuator for robotics. *IEEE Robot. Autom. Lett.* 4:2393–400

78. Ordonez C, Gupta N, Collins EG, Clark JE, Johnson AM. 2012. Power modeling of the XRL hexapedal robot and its application to energy efficient motion planning. In *Adaptive Mobile Robotics*, ed. AKM Azad, NJ Cowan, MO Tokhi, GS Virk, RD Eastman, pp. 689–96. Singapore: World Sci.
79. Makadia A, Patterson A, Daniilidis K. 2006. Fully automatic registration of 3D point clouds. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 1297–304. Piscataway, NJ: IEEE
80. Arslan O. 2019. Statistical coverage control of mobile sensor networks. *IEEE Trans. Robot*. 35:889–908
81. Bajcsy R. 1988. Active perception. *Proc. IEEE* 76:966–1005
82. Bajcsy R, Aloimonos Y, Tsotsos JK. 2018. Revisiting active perception. *Auton. Robots* 42:177–96
83. Cowan NJ, Lee J, Full RJ. 2006. Task-level control of rapid wall following in the American cockroach. *J. Exp. Biol*. 209:1617–29
84. Cowan NJ, Fortune ES. 2007. The critical role of locomotion mechanics in decoding sensory systems. *J. Neurosci*. 27:1123–28
85. Cowan NJ, Ankarali MM, Dyhr JP, Madhav MS, Roth E, et al. 2014. Feedback control as a framework for understanding tradeoffs in biology. *Integr. Comp. Biol*. 54:223–37
86. Miracchi L. 2017. Generative explanation in cognitive science and the hard problem of consciousness. *Philos. Perspect*. 31:267–91
87. Brooks RA. 1991. Intelligence without representation. *Artif. Intell*. 47:139–59
88. Sontag ED. 1998. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. New York: Springer
89. Francis BA, Wonham WM. 1976. The internal model principle of control theory. *Automatica* 12:457–65
90. Brown GS, Campbell DP. 1948. *Principles of Servomechanisms: Dynamics and Synthesis of Closed-Loop Control Systems*. New York: Wiley
91. Mason MT. 1993. Kicking the sensing habit. *AI Mag*. 14(1):58–59
92. Erdmann M, Mason MT. 1988. An exploration of sensorless manipulation. *IEEE J. Robot. Autom*. 4:369–79
93. Brooks R. 1986. A robust layered control system for a mobile robot. *IEEE J. Robot. Autom*. 2:14–23
94. Roberts SF, Koditschek DE, Miracchi LJ. 2020. Examples of Gibsonian affordances in legged robotics research using an empirical, generative framework. *Front. Neurorobot*. 14:12
95. Metz C. 2019. Turing Award won by 3 pioneers in artificial intelligence. *New York Times*, Mar. 27. **https://www.nytimes.com/2019/03/27/technology/turing-award-ai.html**
96. Narendra KS, Parthasarathy K. 1990. Identification and control of dynamical systems using neural networks. *IEEE Trans. Neural Netw*. 1:4–27
97. Esteves C, Allen-Blanchette C, Makadia A, Daniilidis K. 2018. Learning SO(3) equivariant representations with spherical CNNs. In *Computer Vision – ECCV 2018*, ed. V Ferrari, M Hebert, C Sminchisescu, Y Weiss, pp. 54–70. Cham, Switz.: Springer
98. Buehler M, Iagnemma K, Singh S, eds. 2009. *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*. Berlin: Springer
99. Sofge E. 2015. The DARPA Robotics Challenge was a bust. *Popular Science*, July 6. **https://www.popsci.com/darpa-robotics-challenge-was-bust-why-darpa-needs-try-again**
100. Marcus G. 2012. Why making robots is so darn hard. *New Yorker*, Dec. 13. **https://www.newyorker.com/news/news-desk/why-making-robots-is-so-darn-hard**
101. Guizzo E, Ackerman E. 2015. The hard lessons of DARPA's Robotics Challenge. *IEEE Spect*. 52(8):11–13
102. Schmelzer R. 2018. Why are robotics companies dying? *Forbes*, Oct. 29. **https://www.forbes.com/sites/cognitiveworld/2018/10/29/why-are-robotics-companies-dying**
103. Vanderborght B. 2019. Robotic dreams, robotic realities. *IEEE Robot. Autom. Mag*. 26(1):4–5
104. Mervis J. 2020. U.S. Lawmakers unveil bold $100 billion plan to remake NSF. *Science*, May 26. **https://www.sciencemag.org/news/2020/05/us-lawmakers-unveil-bold-100-billion-plan-remake-nsf**
105. Kerry CF, Karsten J. 2017. *Gauging investment in self-driving cars*. Rep., Brookings Inst., Washington, DC
106. Bhana Y. 2015. Drone technology: a tricky take-off or the future of ecommerce? *TranslateMedia*, Feb. 9. **https://www.translatemedia.com/us/blog-us/drone-technology-tricky-take-off-future-ecommerce**

107. Rosenberg N. 1976. *Perspectives on Technology*. Cambridge, UK: Cambridge Univ. Press
108. Rizzo U, Barbieri N, Ramaciotti L, Iannantuono D. 2020. The division of labour between academia and industry for the generation of radical inventions. *J. Technol. Transf.* 45:393–413
109. Latombe JC. 1991. *Robot Motion Planning*. New York: Springer
110. Thrun S, Burgard W, Fox D. 2006. *Probabilistic Robotics*. Cambridge, MA: MIT Press
111. Mason MT, Salisbury JK Jr. 1985. *Robot Hands and the Mechanics of Manipulation*. Cambridge, MA: MIT Press
112. Mason MT. 2001. *Mechanics of Robotic Manipulation*. Cambridge, MA: MIT Press
113. Raibert MH. 1986. *Legged Robots That Balance*. Cambridge, MA: MIT Press
114. Westervelt ER, Grizzle JW, Chevallereau C, Choi JH, Morris B. 2007. *Feedback Control of Dynamic Bipedal Robot Locomotion*. Boca Raton, FL: CRC
115. Tedre M, Simon, Malmi L. 2018. Changing aims of computing education: a historical survey. *Comput. Sci. Educ.* 28:158–86
116. Hewitt C, Kumar V. 2018. The gap in CS, mulling irrational exuberance. *Commun. ACM* 61(11):8–9
117. Dixon L. 2020. Autonowashing: the greenwashing of vehicle automation. *Transp. Res. Interdiscip. Perspect.* 5:100113
118. SAE Int. 2014. *Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems*. Stand. J3016, SAE Int., Warrendale, PA
119. Shladover SE. 2018. Connected and automated vehicle systems: introduction and overview. *J. Intel. Transp. Syst.* 22:190–200
120. Soteropoulos A, Mitteregger M, Berger M, Zwirchmayr J. 2020. Automated drivability: toward an assessment of the spatial deployment of level 4 automated vehicles. *Transp. Res. A* 136:64–84
121. Miller ID, Cladera F, Cowley A, Shivakumar SS, Lee ES, et al. 2020. Mine tunnel exploration using multiple quadrupedal robots. *IEEE Robot. Autom. Lett.* 5:2840–47
122. Zhang F, Niu W. 2019. A survey on formal specification and verification of system-level achievements in industrial circles. *Acad. J. Comput. Inf. Sci.* 2:22–34
123. Moon FC. 2003. Franz Reuleaux: contributions to 19th century kinematics and theory of machines. *Appl. Mech. Rev.* 56:261–85
124. Wiener N. 1948. *Cybernetics: Or Control and Communication in the Animal and the Machine*. New York: Wiley & Sons
125. Conway F, Siegelman J. 2006. *Dark Hero of the Information Age: In Search of Norbert Wiener, the Father of Cybernetics*. New York: Basic Books
126. Mayr O. 1971. Maxwell and the origins of cybernetics. *Isis* 62:425–44
127. Newell A, Simon HA. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall
128. Krieger A, Susil RC, Menard C, Coleman JA, Fichtinger G, et al. 2005. Design of a novel MRI compatible manipulator for image guided prostate interventions. *IEEE Trans. Biomed. Eng.* 52:306–13
129. Fasola J, Mataric MJ. 2012. Using socially assistive human–robot interaction to motivate physical exercise for older adults. *Proc. IEEE* 100:2512–26
130. Jacobs LF. 2012. From chemotaxis to the cognitive map: the function of olfaction. *PNAS* 109:10693–700
131. Ilton M, Bhamla MS, Ma X, Cox SM, Fitchett LL, et al. 2018. The principles of cascading power limits in small, fast biological and engineered systems. *Science* 360:eaao1082
132. Clark J, Goldman DI, Lin PC, Lynch G, Chen TS, et al. 2008. Design of a bio-inspired dynamical vertical climbing robot. In *Robotics: Science and Systems III*, ed. W Burgard, O Brock, C Stachniss, pp. 9–16. Cambridge, MA: MIT Press
133. Goldman DI, Chen TS, Dudek DM, Full RJ. 2006. Dynamics of rapid vertical climbing in cockroaches reveals a template. *J. Exp. Biol.* 209:2990–3000
134. Lynch GA, Clark JE, Lin PC, Koditschek DE. 2012. A bioinspired dynamical vertical climbing robot. *Int. J. Robot. Res.* 31:974–96
135. Federle W, Labonte D. 2019. Dynamic biological adhesion: mechanisms for controlling attachment during locomotion. *Philos. Trans. R. Soc. B* 374:20190199
136. Autumn K, Dittmore A, Santos D, Spenko M, Cutkosky M. 2006. Frictional adhesion: a new angle on gecko attachment. *J. Exp. Biol.* 209:3569–79

137. Santos D, Spenko M, Parness A, Kim S, Cutkosky M. 2007. Directional adhesion for climbing: theoretical and practical considerations. *J. Adhes. Sci. Technol.* 21:1317–41
138. Cutkosky MR. 2015. Climbing with adhesion: from bioinspiration to biounderstanding. *Interface Focus* 5:20150015
139. Asbeck AT, Kim S, Cutkosky MR, Provancher WR, Lanzetta M. 2006. Scaling hard vertical surfaces with compliant microspine arrays. *Int. J. Robot. Res.* 25:1165–79
140. Kim S, Spenko M, Trujillo S, Heyneman B, Santos D, Cutkosky MR. 2008. Smooth vertical surface climbing with directional adhesion. *IEEE Trans. Syst. Man Cybernet. C* 24:65–74
141. Jamone L, Ugur E, Cangelosi A, Fadiga L, Bernardino A, et al. 2018. Affordances in psychology, neuroscience, and robotics: a survey. *IEEE Trans. Cogn. Dev. Syst.* 10:4–25
142. Li C, Zhang T, Goldman DI. 2013. A terradynamics of legged locomotion on granular media. *Science* 339:1408–12
143. Othayoth R, Thoms G, Li C. 2020. An energy landscape approach to locomotor transitions in complex 3D terrain. *PNAS* 117:14987–95
144. Boothroyd G. 1994. Product design for manufacture and assembly. *Comput.-Aided Des.* 26:505–20
145. Papadakis P. 2013. Terrain traversability analysis methods for unmanned ground vehicles: a survey. *Eng. Appl. Artif. Intell.* 26:1373–85
146. Shill JJ, Collins EG Jr., Coyle E, Clark J. 2015. Tactile surface classification for limbed robots using a pressure sensitive robot skin. *Bioinspir. Biomimet.* 10:016012
147. Wu XA, Huh TM, Sabin A, Suresh SA, Cutkosky MR. 2019. Tactile sensing and terrain-based gait control for small legged robots. *IEEE Trans. Robot.* 36:15–27
148. Marden JH, Allen LR. 2002. Molecules, muscles, and machines: universal performance characteristics of motors. *PNAS* 99:4161–66
149. Kenneally G, Chen WH, Koditschek DE. 2018. Actuator transparency and the energetic cost of proprioception. In *Proceedings of the 2018 International Symposium on Experimental Robotics*, ed. J Xiao, T Kröger, O Khatib, pp. 485–95 Cham, Switz.: Springer
150. Alexander R. 1990. Three uses for springs in legged locomotion. *Int. J. Robot. Res.* 9:53–61
151. Kenneally G, De A, Koditschek DE. 2016. Design principles for a family of direct-drive legged robots. *IEEE Robot. Autom. Lett.* 1:900–7
152. Pratt GA, Williamson MM. 1995. Series elastic actuators. In *Proceedings of the 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 1, pp. 399–406. Piscataway, NJ: IEEE
153. Loughlin C, Albu-Schäffer A, Haddadin S, Ott C, Stemmer A, et al. 2007. The DLR lightweight robot: design and control concepts for robots in human environments. *Ind. Robot* 34:376–85
154. Duperret J, Kramer B, Koditschek DE. 2016. Core actuation promotes self-manipulability on a direct-drive quadrupedal robot. In *2016 International Symposium on Experimental Robotics*, ed. D Kulić, Y Nakamura, O Khatib, G Venture, pp. 147–59. Cham, Switz.: Springer
155. Topping TT, Kenneally G, Koditschek D. 2017. Quasi-static and dynamic mismatch for door opening and stair climbing with a legged robot. In *2017 IEEE International Conference on Robotics and Automation*, pp. 1080–87. Piscataway, NJ: IEEE
156. Gilpin K, Rus D. 2010. Modular robot systems. *IEEE Robot. Autom. Mag.* 17(3):38–55
157. Rubenstein M, Cornejo A, Nagpal R. 2014. Programmable self-assembly in a thousand-robot swarm. *Science* 345:795–99
158. Yim M, Shen WM, Salemi B, Rus D, Moll M, et al. 2007. Modular self-reconfigurable robot systems. *Robot. Autom. Mag. IEEE* 14(1):43–52
159. Daudelin J, Jing G, Tosun T, Yim M, Kress-Gazit H, Campbell M. 2018. An integrated system for perception-driven autonomy with modular robots. *Sci. Robot.* 3:eaat4983
160. Tosun T, Sung C, McCloskey C, Yim M. 2019. Optimal structure synthesis for environment augmenting robots. *IEEE Robot. Autom. Lett.* 4:1069–76
161. Bourgeois J, Goldstein SC. 2015. Distributed intelligent MEMS: progresses and perspectives. *IEEE Syst. J.* 9:1057–68
162. Jindrich DL, Full RJ. 2002. Dynamic stabilization of rapid hexapedal locomotion. *J. Exp. Biol.* 205:2803–23

163. Daley MA, Voloshina A, Biewener AA. 2009. The role of intrinsic muscle mechanics in the neuromuscular control of stable running in the guinea fowl. *J. Physiol.* 587:2693–707
164. Diacu F, Holmes P. 1999. *Celestial Encounters: The Origins of Chaos and Stability*. Princeton, NJ: Princeton Univ. Press
165. Bemporad A, Morari M. 1999. Control of systems integrating logic, dynamics, and constraints. *Automatica* 35:407–27
166. Conley CC. 1978. *Isolated Invariant Sets and the Morse Index*. Providence, RI: Am Math. Soc.
167. Culbertson J, Gustafson P, Koditschek DE, Stiller PF. 2020. Formal composition of hybrid systems. *Theory Appl. Categ.* 35:1634–82
168. Kvalheim MD, Gustafson P, Koditschek DE. 2020. Conley's fundamental theorem for a class of hybrid systems. arXiv:2005.03217 [math.DS] (in review for *SIAM J. Appl. Dyn. Syst.*)
169. Johnson AM, Burden SA, Koditschek DE. 2016. A hybrid systems model for simple manipulation and self-manipulation systems. *Int. J. Robot. Res.* 35:1354–92
170. Goebel R, Sanfelice RG, Teel A. 2009. Hybrid dynamical systems. *IEEE Control Syst. Mag.* 29(2):28–93
171. Koditschek DE. 1989. The application of total energy as a Lyapunov function for mechanical control systems. In *Dynamics and Control of Multibody Systems*, ed. JE Marsden, PS Krishnaprasad, JC Simo, pp. 131–57. Providence, RI: Am. Math. Soc.
172. Whitney DE. 1969. Resolved motion rate control of manipulators and human prostheses. *IEEE Trans. Man-Mach. Syst.* 10:47–53
173. Reif JH. 1979. Complexity of the mover's problem and generalizations. In *20th Annual Symposium on Foundations of Computer Science*, pp. 421–27. Piscataway, NJ: IEEE
174. LaValle SM. 2006. *Planning Algorithms*. Cambridge, UK: Cambridge Univ. Press
175. Canny JF. 1988. *Complexity of Robot Motion Planning*. Cambridge, MA: MIT Press
176. LaValle SM, Kuffner JJ. 2001. Randomized kinodynamic planning. *Int. J. Robot. Res.* 20:378–400
177. Farber M. 2003. Topological complexity of motion planning. *Discrete Comput. Geom.* 29:211–21
178. Koditschek DE. 1992. Task encoding: toward a scientific paradigm for robot planning and control. *Robot. Auton. Syst.* 9:5–39
179. Khatib O. 1986. Real-time obstacle avoidance for manipulators and mobile robots. *Int. J. Robot. Res.* 5:90–98
180. Koditschek DE, Rimon E. 1990. Robot navigation functions on manifolds with boundary. *Adv. Appl. Math.* 11:412–42
181. Bhatia NP, Szegö GP. 2002. *Stability Theory of Dynamical Systems*. Berlin: Springer
182. Baryshnikov Y, Shapiro B. 2014. How to run a centipede: a topological perspective. In *Geometric Control Theory and Sub-Riemannian Geometry*, ed. G Stefani, U Boscain, J-P Gauthier, A Sarychev, M Sigalotti, pp. 37–51. Cham, Switz.: Springer
183. Baryshnikov Y. 2015. *Topological perplexity in feedback stabilization*. Paper presented at the Summer Graduate Program in Mathematics on Topological Methods in Complex Systems, Institute for Mathematics and Its Applications, Minneapolis, MN, July 25. **https://faculty.math.illinois.edu/~ymb/talks/perp_talk/perp.html**
184. Burridge RR, Rizzi AA, Koditschek DE. 1999. Sequential composition of dynamically dexterous robot behaviors. *Int. J. Robot. Res.* 18:534–55
185. Lozano-Perez T, Mason MT, Taylor RH. 1984. Automatic synthesis of fine-motion strategies for robots. *Int. J. Robot. Res.* 3:3–24
186. Fikes RE, Nilsson NJ. 1971. STRIPS: a new approach to the application of theorem proving to problem solving. *Artif. Intell.* 2:189–208
187. Hogan N. 1985. Impedance control: an approach to manipulation: part I—theory. *J. Dyn. Syst. Meas. Control* 107:1–7
188. Guastello S, Nathan D, Johnson MJ. 2009. Attractor and Lyapunov models for reach and grasp movements with application to robot-assisted therapy. *Nonlinear Dyn. Psychol. Life Sci.* 13:99–121
189. Bloch A, Chang DE, Leonard N, Marsden J. 2001. Controlled Lagrangians and the stabilization of mechanical systems. II. Potential shaping. *IEEE Trans. Autom. Control* 46:1556–71
190. Arslan O, Guralnik DP, Koditschek DE. 2016. Coordinated robot navigation via hierarchical clustering. *IEEE Trans. Robot.* 32:352–71

191. Farber M, Grant M, Lupton G, Oprea J. 2019. An upper bound for topological complexity. *Topol. Appl.* 255:109–25
192. Johnson AM, Koditschek DE. 2013. Toward a vocabulary of legged leaping. In *2013 IEEE International Conference on Robotics and Automation*, pp. 2568–75. Piscataway, NJ: IEEE
193. Brill A, De A, Johnson AM, Koditschek DE. 2015. Tail-assisted rigid and compliant legged leaping. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6304–11. Piscataway, NJ: IEEE
194. Haynes GC, Cohen FR, Koditschek DE. 2011. Gait transitions for quasi-static hexapedal locomotion on level ground. In *Robotics Research: The 14th International Symposium ISRR*, pp. 105–21. Berlin: Springer
195. Koditschek DE. 1989. Robot planning and control via potential functions. In *The Robotics Review 1*, ed. JJ Craig, O Khatib, T Lozano-Pérez, pp. 349–67. Cambridge, MA: MIT Press
196. Donald B, Xavier P, Canny J, Reif J. 1993. Kinodynamic motion planning. *J. ACM* 40:1048–66
197. Koditschek DE. 1987. Adaptive techniques for mechanical systems. In *Proceedings of the Fifth Yale Workshop on Applications of Adaptive Systems Theory*, pp. 259–65. New Haven, CT: Yale Univ.
198. Rizzi A. 1998. Hybrid control as a method for robot motion programming. In *1998 IEEE International Conference on Robotics and Automation*, Vol. 1, pp. 832–37. Piscataway, NJ: IEEE
199. Ayanian N, Kallem V, Kumar V. 2011. Synthesis of feedback controllers for multiple aerial robots with geometric constraints. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3126–31. Piscataway, NJ: IEEE
200. Bemporad A, Morari M. 1999. Robust model predictive control: a survey. In *Robustness in Identification and Control*, ed. A Garulli, A Tesi, pp. 207–26. London: Springer
201. Conway L. 2012. Reminiscences of the VLSI revolution: how a series of failures triggered a paradigm shift in digital design. *IEEE Solid-State Circuits Mag*. 4(4):8–31
202. Bernstein N. 1967. *The Coordination and Regulation of Movements*. Oxford, UK: Pergamon
203. Full R, Koditschek D. 1999. Templates and anchors: neuromechanical hypotheses of legged locomotion on land. *J. Exp. Biol*. 202:3325–32
204. Libby T, Johnson AM, Chang-Siu E, Full RJ, Koditschek DE. 2015. Comparative design, scaling, and control of appendages for inertial reorientation. *IEEE Trans. Robot*. 32:1380–98
205. Eldering J, Kvalheim M, Revzen S. 2018. Global linearization and fiber bundle structure of invariant manifolds. *Nonlinearity* 31:4202–45
206. De A, Burden SA, Koditschek DE. 2018. A hybrid dynamical extension of averaging and its application to the analysis of legged gait stability. *Int. J. Robot. Res*. 37:266–86
207. De A, Koditschek DE. 2018. Averaged anchoring of decoupled templates in a tail-energized monoped. In *Robotics Research: Volume 2*, ed. A Bicchi, W Burgard, pp. 269–85. Cham, Switz.: Springer
208. De A, Koditschek DE. 2018. Vertical hopper compositions for preflexive and feedback-stabilized quadrupedal bounding, pacing, pronking, and trotting. *Int. J. Robot. Res*. 37:743–78
209. Franci A, Golubitsky M, Bizyaeva A, Leonard NE. 2020. A model-independent theory of consensus and dissensus decision making. arXiv:1909.05765 [math.OC]
210. Guckenheimer J, Holmes P. 1983. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. New York: Springer
211. Topping TT, Vasilopoulos V, De A, Koditschek DE. 2019. *Composition of templates for transitional pedipulation behaviors*. Paper presented at the International Symposium on Robotics Research, Hanoi, Vietnam, Oct. 6–10. **https://repository.upenn.edu/ese_papers/860**
212. Bizzi E, Mussa-Ivaldi FA, Giszter S. 1991. Computations underlying the execution of movement: a biological perspective. *Science* 253:287–91
213. Ting LH, McKay JL. 2007. Neuromechanics of muscle synergies for posture and movement. *Curr. Opin. Neurobiol*. 17:622–28
214. Ting L, Chiel H, Trumbower R, Allen J, McKay JL, et al. 2015. Neuromechanical principles underlying movement modularity and their implications for rehabilitation. *Neuron* 86:38–54
215. Taborri J, Agostini V, Artemiadis PK, Ghislieri M, Jacobs DA, et al. 2018. Feasibility of muscle synergy outcomes in clinics, robotics, and sports: a systematic review. *Appl. Bionics Biomech*. 2018:3934698
216. Ghrist RW. 2014. *Elementary Applied Topology*. Scotts Valley, CA: CreateSpace

217. Barr M, Wells C. 1990. *Category Theory for Computing Science*. New York: Prentice Hall
218. Hudak P, Courtney A, Nilsson H, Peterson J. 2002. Arrows, robots, and functional reactive programming. In *Advanced Functional Programming: Revised Lectures*, ed. J Jeuring, SLP Jones, pp. 159–87. Berlin: Springer
219. Perez I, Bärenz M, Nilsson H. 2016. Functional reactive programming, refactored. *ACM SIGPLAN Not*. 51(12):33–44
220. Brockett RW. 1988. On the computer control of movement. In *1988 IEEE International Conference on Robotics and Automation*, pp. 534–40. Piscataway, NJ: IEEE
221. Murray R, Deno D, Pister K, Sastry S. 1992. Control primitives for robot systems. *IEEE Trans. Syst. Man Cybernet*. 22:183–193
222. Manikonda V, Krishnaprasad PS, Hendler J. 1999. Languages, behaviors, hybrid architectures, and motion control. In *Mathematical Control Theory*, ed. J Baillieul, JC Willems, pp. 199–226. New York: Springer
223. Hristu-Varsakelis D, Egerstedt M, Krishnaprasad PS. 2003. On the structural complexity of the motion description language MDLe. In *42nd IEEE International Conference on Decision and Control*, Vol. 4, pp. 3360–65. Piscataway, NJ: IEEE
224. Doyen L, Frehse G, Pappas GJ, Platzer A. 2018. Verification of hybrid systems. In *Handbook of Model Checking*, ed. EM Clarke, TA Henzinger, H Veith, R Bloem, pp. 1047–110. Cham, Switz.: Springer
225. Kress-Gazit H, Fainekos G, Pappas G. 2009. Temporal-logic-based reactive mission and motion planning. *IEEE Trans. Robot*. 25:1370–81
226. Dantam N, Stilman M. 2013. The motion grammar: analysis of a linguistic method for robot control. *IEEE Trans. Robot*. 29:704–18
227. Kress-Gazit H, Fainekos GE, Pappas GJ. 2008. Translating structured English to robot controllers. *Adv. Robot*. 22:1343–59
228. Tellex S, Kollar T, Dickerson S, Walter MR, Banerjee AG, et al. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pp. 1507–14. Palo Alto, CA: AAAI Press
229. Saha I, Ramaithitima R, Kumar V, Pappas GJ, Seshia SA. 2014. Automated composition of motion primitives for multi-robot systems from safe LTL specifications. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1525–32. Piscataway, NJ: IEEE
230. Wong KW, Ehlers R, Kress-Gazit H. 2018. Resilient, provably-correct, and high-level robot behaviors. *IEEE Trans. Robot*. 34:936–52
231. Tellex S, Gopalan N, Kress-Gazit H, Matuszek C. 2020. Robots that use language. *Annu. Rev. Control Robot. Auton. Syst*. 3:25–55
232. Cowley A, Taylor CJ. 2011. Stream-oriented robotics programming: the design of roshask. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1048–54. Piscataway, NJ: IEEE
233. Kortik S, Saranli U. 2019. Robotic task planning using a backchaining theorem prover for multiplicative exponential first-order linear logic. *J. Intell. Robot. Syst*. 96:179–91
234. Gorn S. 1963. The computer and information sciences: a new basic discipline. *SIAM Rev*. 5:150–55
235. Hofstra B, Kulkarni VV, Galvez SMN, He B, Jurafsky D, McFarland DA. 2020. The diversity–innovation paradox in science. *PNAS* 117:9284–91
236. Stoet G, Geary DC. 2018. The gender-equality paradox in science, technology, engineering, and mathematics education. *Psychol. Sci*. 29:581–93
237. Merolla DM, Jackson O. 2019. Structural racism as the fundamental cause of the academic achievement gap. *Sociol. Compass* 13:e12696
238. NAFSA. 2020. *Losing talent: an economic and foreign policy risk America can't ignore*. Policy Resour., NAFSA, Washington, DC
239. Quinn DM, Cooc N. 2015. Science achievement gaps by gender and race/ethnicity in elementary and middle school: trends and predictors. *Educ. Res*. 44:336–46
240. Mendoza-Denton R, Patt C, Fisher A, Eppig A, Young I, et al. 2017. Differences in stem doctoral publication by ethnicity, gender and academic field at a large public research university. *PLOS ONE* 12:e0174296
241. Whittaker JA, Montgomery BL, Martinez Acosta VG. 2015. Retention of underrepresented minority faculty: strategic initiatives for institutional value proposition based on perspectives from a range of academic institutions. *J. Undergrad. Neurosci. Educ*. 13:A136–45

242. Berends M. 2015. Sociology and school choice: what we know after two decades of charter schools. *Annu. Rev. Sociol.* 41:159–80
243. Jabbar H, Fong CJ, Germain E, Li D, Sanchez J, et al. 2019. The competitive effects of school choice on student achievement: a systematic review. *Educ. Policy*. **https://doi.org/10.1177/0895904819874756**
244. Mishra S. 2020. Social networks, social capital, social support and academic success in higher education: a systematic review with a special focus on 'underrepresented' students. *Educ. Res. Rev.* 29:100307
245. Jenkins C. 2020. Before we put $100 billion into AI. . . . *VentureBeat*, Aug. 8. **https://venturebeat.com/2020/08/08/before-we-put-100-billion-into-ai**
246. Johnson AM, Axinn S. 2013. The morality of autonomous robots. *J. Mil. Ethics* 12:129–41
247. Vardi MY. 2016. Are robots taking our jobs? *The Conversation*, Apr. 6. **https://theconversation.com/are-robots-taking-our-jobs-56537**
248. Bernstein A, Raman A. 2015. The great decoupling: an interview with Erik Brynjolfsson and Andrew McAfee. *Harvard Business Review*, June. **https://hbr.org/2015/06/the-great-decoupling**
249. Zysman J, Kenney M. 2018. The next phase in the digital revolution: intelligent tools, platforms, growth, employment. *Commun. ACM* 61(2):54–63
250. Hecht B, Wilcox L, Bigham JP, Schöning JP, Hoque E, et al. 2018. It's time to do something: mitigating the negative impacts of computing through a change to the peer review process. *ACM Future of Computing Blog*, Mar. 29. **https://acm-fca.org/2018/03/29/negativeimpacts**
251. World Econ. Forum. 2018. *The future of jobs report 2018*. Rep., World Econ. Forum, Geneva
252. Autor DH. 2015. Why are there still so many jobs? The history and future of workplace automation. *J. Econ. Perspect.* 29(3):3–30
253. Howard A, Kennedy M III. 2020. Robots are not immune to bias and injustice. *Sci. Robot.* 5:eabf1364
254. Johnson AM, Admoni H, Berry CA. 2020. Supporting black scholars in robotics. *IEEE Spectrum*, Sept. 10. **https://spectrum.ieee.org/automaton/at-work/education/supporting-black-scholars-in-robotics**