

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 04-05-2017		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 14-Jan-2016 - 13-Jan-2017	
4. TITLE AND SUBTITLE Final Report: Study Proposal:Transforming Terrestrial Agility At All Scales			5a. CONTRACT NUMBER W911NF-16-1-0035		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Todd Murphey, Eva Kanso, Yasamin Mostofi, Evangelos Theodorou			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Northwestern University Evanston Campus 1801 Maple Avenue Evanston, IL 60201 -3149			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 68154-EG.3		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This study is the result of a year long effort to assess how to meet the Army's requirements for off-road agility over the next several decades. The resulting recommendations are in the areas of control, sensing and perception, and energy. These recommendations include a wide range of topics, including information-based control, collaborative active sensing, dedicated computing elements designed for perception and control, model-based control in model-free and data-driven settings, and energy harvesting.					
15. SUBJECT TERMS terrestrial agility, control, perception					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Todd Murphey
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 847-467-1041

Report Title

Final Report: Study Proposal:Transforming Terrestrial Agility At All Scales

ABSTRACT

This study is the result of a year long effort to assess how to meet the Army's requirements for off-road agility over the next several decades. The resulting recommendations are in the areas of control, sensing and perception, and energy. These recommendations include a wide range of topics, including information-based control, collaborative active sensing, dedicated computing elements designed for perception and control, model-based control in model-free and data-driven settings, and energy harvesting.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received

Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Prof. Murphey was a member of the National Academies / National Research Council Committee on Counter-Unmanned Aircraft System (CUAS) Capability for Battalion-and-Below Operations (2016)

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Todd Murphey	0.04	
Yasamin Mostifi	0.04	
Eva Kanso	0.04	
Evangelos Theodorou	0.04	
FTE Equivalent:	0.16	
Total Number:	4	

Names of Under Graduate students supported

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

1 a. University of Southern California

1 b. Contracts & Grants

3720 S. Flower St.

Los Angeles CA 900890701

Sub Contractor Numbers (c): \$15,000.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Eva Kanso will participate in the Army Research Office (ARO) study on terrestrial :

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. University of Southern California

1 b. 837 W. 36th Place, STO 330

Los Angeles CA 900890001

Sub Contractor Numbers (c): \$15,000.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Eva Kanso will participate in the Army Research Office (ARO) study on terrestrial :

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. University of Southern California

1 b. Dept of Contracts and Grants Marina

University Of Southern California

Marina Del Rey CA 900890001

Sub Contractor Numbers (c): \$15,000.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Eva Kanso will participate in the Army Research Office (ARO) study on terrestrial :

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. University of Southern California

1 b. 4676 Admiralty Way

Suite 1001

Marina Del Rey CA 902926611

Sub Contractor Numbers (c): \$15,000.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Eva Kanso will participate in the Army Research Office (ARO) study on terrestrial :

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. University of Southern California

1 b. Dept. Of Contracts And Grants

University Of Southern California

Marina Del Rey CA 902926611

Sub Contractor Numbers (c): \$15,000.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Eva Kanso will participate in the Army Research Office (ARO) study on terrestrial :

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. University of Southern California

1 b. 3720 S. Flower Street

3rd Floor

Los Angeles CA 900890701

Sub Contractor Numbers (c): \$15,000.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Eva Kanso will participate in the Army Research Office (ARO) study on terrestrial :

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. University of California - Santa Barbara

1 b. 3227 Cheadle Hall

3rd floor, MC 2050

Santa Barbara CA 931062050

Sub Contractor Numbers (c): \$14,798.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Yasamin Mostofi will participate in the ARO study on terrestrial agility between sun

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. University of California - Santa Barbara

1 b. Cheadle Hall, Room 3227

Santa Barbara CA 931062050

Sub Contractor Numbers (c): \$14,798.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Prof. Yasamin Mostofi will participate in the ARO study on terrestrial agility between sun

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. Georgia Institute of Technology

1 b. 225 NORTH AVENUE

ATLANTA GA 303320357

Sub Contractor Numbers (c): \$14,798.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Pof. Evangelos Theodorou will participate in the ARO study on terrestrial agility between

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

1 a. Georgia Institute of Technology

1 b. 225 NORTH AVENUE

ATLANTA GA 303320357

Sub Contractor Numbers (c): \$14,798.00

Patent Clause Number (d-1): 52.227-11

Patent Date (d-2): 5/1/14 12:00AM

Work Description (e): Pof. Evangelos Theodorou will participate in the ARO study on terrestrial agility between

Sub Contract Award Date (f-1): 1/14/16 12:00AM

Sub Contract Est Completion Date(f-2): 1/13/17 12:00AM

Inventions (DD882)

Scientific Progress

The scientific progress is entirely contained in the report document, and summarizes the study group's view of areas of research that can improve terrestrial agility in the next two decades. The areas emphasized are in control, sensing, and energy management.

Technology Transfer

Study: Transforming Terrestrial Agility At All Scales

Todd Murphey (Lead), Northwestern University, Evanston, IL 60208

Eva Kanso, University of Southern California, Los Angeles, CA 90089

Yasamin Mostofi, University of California, Santa Barbara, Santa Barbara, CA 93106

Evangelos Theodorou, Georgia Institute of Technology, Atlanta, GA 30332

Abstract

This study is the result of a year long effort to assess how to meet the Army’s requirements for off-road agility over the next several decades. The resulting recommendations are in the areas of control, sensing and perception, and energy. These recommendations include a wide range of topics, including information-based control, collaborative active sensing, dedicated computing elements designed for perception and control, model-based control in model-free and data-driven settings, and energy harvesting.

1 Introduction

The U.S. Army has a need for high levels of mobility on the ground, away from roads and other support infrastructure. The goal of this study is to assess opportunities that could transform ground-based platforms in terms of both speed and agility, in scenarios ranging from teleoperation to fully autonomous execution. But what defines agility? Certainly some examples of agility bring to mind fast motion. Efficiency also would seem to play a role, since motion with limitless power is rare. But other factors play a role as well, including correct perception for motion, where sensing is integrated into decision making sufficiently quickly to make time critical decisions. Lastly, agility would sometimes seem to depend on correct long-term reasoning about physical cause and effect; for instance, parkour athletes have to make many interdependent decisions before a reasonable plan can be settled upon, and then enact that plan in time constrained settings. This study aims to identify research activities that will bring agility to advanced systems that need to be deployed in unstructured environments.

Our recommendations are in four categories of research, though these categories should be interpreted as deeply intertwined—they are separable now, but in many regards one of our main conclusions is that these areas need to become *inseparable* over time. These delineations are Control, Sensing and Perception, and Energy.

We additionally attempt to discuss how agility might change at different scales (e.g., vehicle scale, human scale, and “small” scale—by which we mean a variety of scales smaller than those on the order of mammals). What sort of principles of motion, decision-making, and design can be expected to translate well between these different scales? Moreover, at these scales, what examples can be exploited for design principles?

Terrestrial animals can swerve, jump, climb, turn and stop abruptly. Their remarkable stability and maneuverability provide an attractive paradigm for engineered robotic vehicles. However, robotic vehicles remain no match to their biological counterparts in agility and ability to deal with variable environmental conditions. What will it take to develop robotic vehicles as stealthy as a fox? Or a robotic flyer as agile as a hummingbird? Or a swarm of robots that collaborate together as a colony of ants, a flock of birds or a school of fish? And in these cases, is agility an attribute of individuals or of the group as a whole? The outstanding issues preventing transition of these capabilities to robotic applications are at once technological and scientific. Technological advances in smart materials and mechanical actuators and sensors are imperative for the development of future generations of autonomous and agile robots. Concurrently, there is a need for basic scientific research to accompany these technological advances and allow them to be successfully incorporated in the design and development of agile robotic vehicles. And although we may use biological observations as evidence that agility is possible, we also recognize that many agile technologies will not mimic the biology that motivates them.

Before moving on to the main section of this document, we note that there are two important ideas missing from the discussion. First, agility for very large vehicles—those on the order of tens of meters, such as long-haul trucks and larger—is ignored. Although one might argue that the energetic requirements for rapidly accelerating such large systems may make “agility” for such systems impractical, it is primarily the case that we simply did not make any progress thinking about that scale. Secondly, we do not consider human-in-the-loop issues, nor do we consider any of the associated psychological factors that may come into play when a person is interacting with an agile system. These considerations are undoubtedly extremely important, but for the purposes of this study we focus on autonomous agility, understanding that many of the recommendations here may play a role in human-machine agility.



Figure 1: Examples of agile and scalable behavior in animals at various length scales. Top row: many-legged locomotion, flying, swimming and galloping modes of locomotion in biological organisms. Bottom row: behavior of same organisms at the collective level.

Our recommendations may be summarized as follows. First, in the area of control, we recommend that new theoretical foundations for control be developed that enable flexible response to a continually-evolving environment. We consider several versions of this idea, including information-theoretic control techniques, goal-oriented principles of motion and decision-making, and the impact that changes in computing hardware may have. The overarching theme, however, is more general—that controlled response to state needs to be reformulated from scratch to deal with variable-morphology, variable-environment systems, rather than simply “adding” onto the optimization-based, linearization-based understanding of control that has dominated the past half century. Second, in the area of sensing and perception, we recommend that cooperative active sensing using ambient environmental signals should be pursued. The specific example we discuss is the use of bluetooth and wifi signals to map the interior of buildings, but the broader picture is that relatively incapable, isolated systems may radically improve their situational understanding by cooperatively exploiting ambient signals. We also argue that robotics and autonomy have matured to a level where it should be possible to create dedicated computation for sensing and actuation closed-loop computation. In this section, we additionally advocate for automating the data-to-control-response pipeline, specifically using recently developed techniques in model-based control for model-free systems. Third, we end with a discussion of energy in the context of energy harvesting and micron-scale systems. In many regards, profound changes in energy storage and management are simultaneously the most important issues for terrestrial agility and the area for which our recommendations are least concrete. Nevertheless, we argue that identifying situations where energy harvesting can be used will have important consequences, specifically in the case of long time horizons or very small scales.

2 Actuation and motor control

Locomotion is one of the most basic behaviors of organisms, yet even the simplest motions often require the coordination of nervous, sensory and motor systems with the external environment. A central problem in biolocomotion, and motor control in general, is understanding how these systems are coordinated to achieve a locomotory task. Which aspects of coordination are “passive” and which aspects require “active” control? A large body of evidence shows that animals accomplish a variety of motor tasks reliably and repeatedly but with large variability on the level of the individual degrees of freedom. Trial-to-trial variability in individual degrees of freedom is on average larger than variability at the task level. These observations are fundamentally incompatible with the mainstream control paradigm that enforces precise control of each variable and degree-of-freedom to achieve a “desired trajectory.” (There are exceptions to control goals being specified through the use of a desired trajectory—e.g., energetically-specified goals [1, 2, 3, 4, 5, 6, 7] and thermodynamically-specified [8, 9, 10, 11]—but these represent a tiny minority of approaches in control.) Supporting and allowing trial-to-trial variability suggests that one should focus on using only goal-directed corrections and explicitly allowing variability that does not interfere with the goal. Therefore, there is a need to build a mathematical framework that allows the development of design rules and control strategies that use actuation only when and where it is necessary. Below we will discuss several examples of what such a framework might look like. But we first point out differing principles of control—focusing on geometric, optimality-based, and information-theoretic

formulations—that provide much of the history relevant to control synthesis.

Geometric Mechanics as a Description of Physics-Based Control and Actuation. Bio-inspired locomotion is based on the interaction of internal or “shape”-related degrees-of-freedom with the surrounding environment. One of the successes in control has been in the development of geometry-based methods for control; these have provided a unified framework for studying locomotory systems at a wide range of length scales. Roughly speaking, the configuration space can be decomposed, using tools from reduction and symmetry, into shape variables and variables associated with the net locomotion. Locomotion gaits can be produced by prescribing different paths in the shape space [12, 13, 14, 15, 16, 17]. However, ignoring the shape dynamics comes at a price. It assumes control of all shape variables and does not provide freedom or flexibility in actuation.

To be able to explain, evaluate and emulate the remarkable agility of biological systems under variable environmental conditions, this framework needs to be expanded to allow the shape space to be decomposed into actively-controlled variables and passively-responding flexible elements. The choice of the actuated variables depends on the specific task (e.g., rapid maneuvers versus periodic locomotion) and environmental conditions. This view allows for “redundancy” in actuation, similar to the redundancy in biological systems, in the sense that, out of the available shape variables, multiple combinations of active and passive elements might produce the same desired output, albeit at different energies, speeds, and stability properties; thus the need for actuation strategies that optimize the system’s performance in terms of, not only efficiency, but also stability and agility.

Optimality and the Minimal-Intervention Principle. Optimality-based control synthesis [18, 19, 20, 21, 22, 23, 24] are among the most used control synthesis techniques, partially because of their tractability in the case of linear systems [25]. However, as pointed out above, sometimes an optimality-based principle can over-specify the control approach. The minimal intervention principle—of which there are many variants—provides an alternative. Qualitatively, the minimal intervention principle can be stated as follows: deviations in individual variables are corrected only when they interfere with the task performance [26]. This approach relies on the intrinsic physics of the system and requires active control only if and when needed. The research question is to decide when and where control is needed. Such decisions obviously depend on the system itself and its environment. *The challenge here is to develop strategies that allow the “control” to follow the physics.* Some versions of the minimal intervention principle have extremely helpful computational properties; for instance, in [27], framing the minimal intervention problem as finding the next action that best improves the long-time horizon behavior *relative* to a nominal choice of control (often the zero control) leads to an analytic solution for arbitrary nonlinear piecewise continuous systems. Moreover, the resulting control coincides with the optimal solutions, but at much less computational cost. The key point is that reframing control questions could lead to dramatic computational savings, scalability, and other factors contributing to real-time capabilities; the minimal intervention principle appears to have many desirable characteristics as a principle of motion.

Moreover, minimal intervention is sometimes a more “natural” problem specification. For instance, minimal intervention is particularly relevant in light of the recent development in soft robotics [28, 29] which provide a new paradigm for robotics and actuation where complete control over all degrees of freedom is not feasible. Alternatively, rehabilitation devices may only want to intervene when safety is at stake, to ensure that a subject is as responsible as possible for his or her own movement [30].

Applications should suggest even more dramatic reframing of the control problem—for instance, what should the control problem be when there is no model, but there is plenty of data? In this case, it is important to explore data-driven versus model-based control and finite versus infinite dimensional representations of a robotic system. Learning algorithms based on artificial neural nets could provide an efficient, model-independent framework for versatile motion design and planning [31, 32]. We will discuss this particular issue more in Section 3.

Variable Geometry Systems. An application that seems both natural and challenging for current formulations of control is that of variable geometry systems. We define variable geometry systems to be those that can change their morphological structure as part of their locomotion strategy. For instance, robots that can transition between legged




Figure 2: Geometric picture underlying undulatory locomotion.

and wheeled locomotion have variable geometry, with two possible geometries to choose from. Other systems could, of course, have more choices of morphology, all the way to a continuum of choices. The question, then, is how should one control such a system, both in terms of the physical inputs and in terms of the morphological choices?

Variable geometry systems have two distinct classes of decision variables, and as a result we identify three scenarios or modes of operation. In the first operational mode, there is a change in the geometry while the control (such as steering and throttle) remain constant. In the second operational mode there is the dual situation in which the geometry remains fixed but control commands have to be optimized. The third operational mode is the scenario where geometry and control commands have to be simultaneously optimized. There are already control-oriented tools for all three of these operational modes: methods in hybrid control [33, 34] approach the problem as a scheduling problem while minimizing an objective function, methods in information theoretic control capture the changes in geometry and structure through information measures [35], and direct optimization methods can approach the problem through integer-constrained optimization [36]. All these methods suffer from scalability issues, slow computation time, and sensitivity to the fidelity of the model.

Terrestrial agility, enabled by variable geometry designs, will need to have these problems solved, with concrete algorithms that can act on both the internal geometry of the vehicle as well as control actions. Stochastic control methods that rely on information theoretic measures could provide one way to derive this controllers, so here we include discussion of this area specifically as well as some of its near-term goals. Then, at the end of this section, we provide our perspective on a longer-term research agenda.

Information Theoretic Formulations of Control. The information theoretic formulation of control is very general, applicable to both differentiable and nondifferentiable system, both in their dynamics and their objectives—and as a result can model systems with variable geometry. Stochastic control is a representative special case and can be used as a first step for computationally practical approaches. Numerical methods for stochastic control using sampling have been extensively studied [37, 38, 39, 40]; however, most of these approaches suffer from scalability issues. Recently new methods have been developed that overcome the issue of scalability by computing optimal control policies using forward sampling and adaptive importance sampling. These methods rely on the connection between information theoretic dualities between free energy and relative entropy, path integrals and the dynamics programming principle [41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52]. Resulting algorithms such as the Model Predictive Path integral control are able to steer terrestrial vehicles such as the GT AutoRally car in an aggressive fashion [53, 54].

However, there are open challenges that emerge in the area of terrestrial agility and would have to be addressed before these methods could be applied to real vehicles in real operating environments. These near-term challenges are not only related to the information theoretic formulation of stochastic control, but are relevant to alternative formulations as well.

- *High Dimensionality and Scalability:* The first challenge is related to the dimensionality of the decision variables especially the ones related to the internal geometry of the terrestrial vehicle. Optimization tools for high dimensional optimization problems and tools used for training neural networks for deep learning are of potential interest as well.
- *Optimization in Multiple Time/Spatial Scales and Hierarchies:* The second challenge is related to the time scale of optimization. It is possible that the two types of decision variable have to be optimized in different time scales. This means that iterative optimization algorithms have to be developed for optimization that alternates between different time scales. On this front, elements on stochastic calculus for importance sampling for multi-scale diffusion processes will be important in cases where optimization is performed via sampling [55, 56]. Multi-scale diffusion processes can be used to create hierarchies for decision making under uncertainty and planning. Hierarchical decision making using an information theoretic formulation has the potential to create a computational framework that can unify high level reasoning and voluntary decision making with low level reactive reflexive control.
- *Multiple Models and Hybrid Stochastic Control:* Depending on the properties of the terrain the state space dynamics of the vehicle may rapidly change, potentially discontinuously so. This rapid change may not be captured by the stochastic representation of the problem statement. In this case multiple state space representations will be required to capture the entire spectrum of the dynamics of the terrestrial vehicle. Hybrid representations and multiple models [57] should be used in cases where terrestrial vehicles have *variable geometry* or *variable environment* (e.g., short-flight or amphibious capabilities). The aforementioned algorithms should be reactive,

adaptive, stochastic and able to handle different possibilities and scenarios in forward planning. It almost goes without saying that this is an ambitious goal, but it is nevertheless one that is likely solvable within the scope of current state-of-the-art techniques.

More ambitious goals in control techniques include the following.

Policy-Based Stochastic Control Using Tensor Train Decompositions. A classical approach to solving stochastic control problems is to use Markov Chain Model Carlo approximation methods [37]. The key ingredient of this method is that it relies on the generation of a grid in the state space and solution of the Bellman equation on this grid by forcing the so-called consistency condition to hold. One of the major limitations of this class of methods has been the lack of scalability and its sensitivity to regularity of the state space grid. Recent developments in the area of tensor train decomposition and compressed sensing [58, 59] seem to alleviate the issue of scalability to some degree by solving the Bellman equation in a lower dimensional representation of the initial problem formulation. The framework is system-specific and requires the representation of dynamics as products of functions of each dimension of the state space. These approaches are used for tasks that involve stabilization or tracking. Application of the aforementioned computational methodologies to terrestrial vehicles will require work on incorporating hybrid stochastic dynamics, multiple model-based stochastic control and multi-scale dynamics.

Stochastic Control Using Linearization. In the third class of stochastic control methods algorithms, such as differential dynamic programming (DDP), approximate the HJB PDE locally along nominal trajectories. Specifically, the DDP algorithm back propagates the value function along a trajectory using local quadratic approximations of the cost, value function, and system dynamics evaluated on the nominal trajectory. Locally optimal open- and closed-loop control policies are then computed based on the quadratically approximated value function [60]. The same basic principles were used to develop iterative linear quadratic regulators [61, 62]. The stochastic differential dynamic programming (S-DDP) algorithm considers stochastic system dynamics with additive control- and state-dependent noise and finds optimal open- and closed-loop control policies to minimize the expectation of a given cost [63, 61]. Unlike sampling based methods, the S-DDP algorithm's performance is significantly affected by its ability to propagate the considered system's configuration and accuracy of the linearization of the resultant stochastic discrete dynamics. Therefore future research on stochastic control methods using linearization should include the development of integrators that encode mechanical information. For instance, variational integrators constrain (or strongly enforce) the conservation of fundamental mechanical quantities such as momentum, energy, the symplectic form (related to volume preservation for stochastic processes), and geometric constraints [64, 65, 66]. These numerical methods can also be linearized, providing a mechanically-relevant linearization for purpose of control [67]. The key point is that advances in numerical can directly impact practical control systems, particular ones that depend on forward simulation of the mechanics and backwards simulation of the sensitivity.

Moreover, uncertainty potentially can be represented explicitly in the form of stochastic differential equations or by using deterministic representations such generalized Polynomial Chaos (gPC) theory and Gaussian Process (GP) regression. (It is important to emphasize that deterministic representation of uncertainty come with the drawback of increasing the dimensionality of the initial stochastic representation.) Therefore these representations expand the state space dimensions and make the application of trajectory optimization methods challenging.

There exist important questions related to uncertainty representation and the way that this representation can be incorporated within linearization-based stochastic control methods. The choice of uncertainty representation will result in different mathematical tools for useful linearization and variational integration. For example, in the case of stochastic differential equations, investment should be made in the development of variational integrators for generalized models of stochasticity that go beyond diffusion processes, motivated by the fact that real mechanical uncertainty is almost never of the form of a diffusion process. In the case of polynomial chaos representation, research should be steered towards the development of scalable variational integrators. More generally, uncertainty quantification that is both computable and anticipates the synthesis of control is essential in the near term.

Infinite Dimensional Representations: Many vehicle designs could include continuum-based designs, rather than relying on rigid parts. In the case of infinite dimensional representations, there are challenges related to whether there exist computational algorithms for control and how scalable or real time these algorithms could be. Infinite dimensional representations of dynamical systems such as partial differential equations could be used in case of high fidelity morphing/variable geometry representations. In the literature of control theory, control of PDE is also known under the name of *distributed control*. Work in the area of distributed control involves the generalization of the two principles of optimality namely Pontryagin's Maximum Principle and Dynamic Programming principles [68, 69].

There is a great potential for distributed control algorithms within the area of terrestrial agility for vehicles that have morphing capabilities which go beyond multi-body physics. On the theoretical side, there is existing work on

optimal control of PDEs however most of this work is at the level of existence and uniqueness of solutions and less on the computational and algorithmic side. Motivated by the need to create highly agile terrestrial vehicles there are many opportunities related to the development of computational algorithms for real time shaping and morphing. The major difficulty for this research will be on the scalability and real time requirements of the computation. One possible way towards making progress on that front is to ask the question which types of algorithms scale for simpler finite dimensional representations of systems in robotics and autonomy and then proceed with work towards generalizing or *lifting* these algorithms so that they are applicable to infinite dimensional representations.

In the robotics and learning communities there is a consensus that certain classes of optimal control methods are scalable. These classes include Policy Gradient or Policy Search methods [70, 71, 72, 73, 74, 75, 76, 51], model-based trajectory optimization methods such as Differential Dynamic Programming [60] and its variations [77, 78, 79, 63, 80, 81, 82], and stochastic control methods using sampling [53, 54, 48, 83]. Future efforts towards the development of scalable, real time algorithms for morphing in terrestrial vehicles should include the *lifting* of these classes of methods to the infinite dimensional spaces. Some prior work on the control of stochastic PDEs using the exponential transformation of the value function already exists however the algorithmic and computational aspects of this work are completely unexplored [84, 85, 86].

2.1 Hardware for Control Computation

Recent advances in the area of stochastic control and machine learning have shown that parallelism is a key ingredient for state of art algorithms to meet the real time requirements of autonomous systems operating in dynamic environments. From the decision making point of view, parallelism supports the computation of reactive stochastic controllers. From the machine learning and perception point of view, parallelism allowed researchers to train multiple layers of deep neural networks and help them to investigate the best architectures for a given task. There are different solutions that support parallelization. Parallelization can be achieved in software using high end-low power GPUs but it can also be performed in hardware using different technologies such as FPGAs and neuromorphic computation.

Within the context of terrestrial agility and autonomy, it is essential to have hardware that supports parallelization for real time decision making and learning while maintaining low power consumption. The FPGA technology could be seen as an intermediate step towards the development of completely neuromorphic hardware that promises computation with very low power consumption. FPGA technology has been used in many application within autonomy but mostly for subtasks such as image and signal processing, vision and filtering. FPGAs have been used in high speed trading and Monte Carlo simulation in the financial and banking sector however the trading models used are low dimensional-typically scalar [87] and rather simple in terms of nonlinearities. One of the questions here is whether this technology can be used to perform real time decision making for high speed navigation of terrestrial vehicles with variable geometry.

Neuromorphic technology relies on a new micro-device (called memristors) and promises high speed computation under low power consumption [88, 89, 90, 91, 92]. Memristor devices are inspired by spiking neurons and the energetic efficiency of biological neural networks. One the fundamental questions is whether this technology can be used for forward sampling in hardware of thousand of trajectories in stochastic control or for representing in hardware deep neural networks architectures and performing online real time training. Future terrestrial vehicles should be equipped with neuromorphic chips [93] for implementing decision making, learning and sensing in an energetically efficient way. Moreover, control architectures will need to take these computational capabilities into account.

2.2 Long Term Recommendations

We recommend investment in control advances that generalize geometric methods to situations where variability of approach is built into the mathematical representation of the control problem. Flexibility of the control synthesis with respect to strategy is essential, and will require relaxing some of the assumptions associated with classical methods in control while adding new axioms of motion. Integrating information-driven metrics into control synthesis at an intrinsic level—rather than estimation or linear process models—should be prioritized. This new theory would include advantages of geometric, stochastic, minimal-intervention methods, but would presumably not look like any one of them; it would be more than simply modifying previously existing theories.

A key point is that control theory likely needs to be reformulated based on a new set of requirements, requirements that no longer treat the “trajectory” as the most primitive description of the goal, requirements that build in assumptions about computation, sensing and perception, and requirements that reflect needs for flexibility and adapt-

ability. Rather than simply modifying existing theoretical foundations, these foundations should look quite different from the ones we see now.¹

Most importantly, developing theoretical tools that lead to a simplification of computation and implementation for control is an important goal to pursue [94]. Currently, each new requirement (e.g., the inclusion of stochastic terms in the statement of a model) leads to increasing computational complexity and sophistication of implementation (e.g., stochastic control requires either solving partial differential equations or evaluation of large numbers of sample paths). However, it seems unlikely that each new requirement (e.g., incorporating uncertainty into decisions) should lead to dramatic increases in complexity. Instead, what is needed is something like the original Kalman filter insight—that filtering problems with slightly more structure than the Wiener process assumptions can be solved with dramatically simpler computations and implementations.

3 Sensing, Data, and Perception

In this section we discuss sensing and perception, and specifically how these might change dramatically as a result of, and contribute to, advanced forms of agility. Moreover, how should agility be impacted by the ubiquity of data? Again, we look to biology for some guidance.

Organisms leverage a wide variety of sensory modalities, such as vision, olfaction or flow sensing, in order to respond to changes in their environment. To deploy robotic vehicles in unknown terrains, they must be able to rely on online sensory systems to identify and decipher the important features of their environment. Sensing of environmental cues consists of four complex and interdependent components: (i) the signal or environmental cue to be sensed or detected, (ii) the type and layout of the sensory system used to measure the signal, (iii) the decoding algorithms that analyze the sensory measurements and decipher the signal, and (iv) the decisions that are made to obtain favorable signals. Lastly, in the case of multiple sensory modalities, there is a need for strategies to integrate the sensory information from various sources. Each of these topics is the subject of intensive research in specific model systems, for example, vision in fruit flies [95, 96, 97], hydrodynamic sensing in marine animals [98, 99, 100, 101, 102, 103], chemical sensing in insects and zooplankton [104, 105, 106, 107, 108], et cetera. However, so far, there is no framework that enables the seamless integration of the knowledge developed in these various disciplines, nor is there a systematic way to predict how they affect each other. Nevertheless, some of these observations point to new opportunities in sensing that would have been impossible—or at least impractical—even a few years ago. We start this discussion with an example of sensing using ambient communication signals.

Sensing and Perception with on-board communication signals. Can agile platforms use on-board transceivers and the associated generated signals for sensing and perception in uncertain environments? This, for instance, could enable through-wall sensing, which is not possible with traditional on-board sensors such as laser scanners. More importantly, for small micron-scale platforms, carrying advanced imaging/vision systems may not be possible. Can they then use MEMS-based RF radios for sensing and perception on very small platforms? MEMS-based RF radios creates the potential for micro-miniaturization of radios, which can provide a viable option for micron-scale platforms.

Traditionally, sensing and imaging with electromagnetic waves is achieved at frequencies such as 10^{16} Hz (e.g., x-ray), while sensing and perception with everyday RF signals, such as WiFi and Bluetooth, is much more rare, despite the fact that these signals also permeate walls and other obstacles. Let X represent the vector of the stacked up voxel values of the 3D unknown space that the platforms are interested in reconstructing by using RF transmissions. X can be a binary vector, with each element representing the presence or absence—the occupancy—of an object in the corresponding voxel. Equivalently, $X(i)$ can represent the material property of the corresponding voxel, or any other feature of interest, so long as the RF signals in question are sensitive to that feature—that RF signals contain information about the feature. Let $y = f(X, T) + n$ represent the vector of the RF-based measurements collected by the platforms, with T denoting a matrix whose i^{th} column is the way-points of the trajectory of the i^{th} robot, and n representing measurement noise. Function f depends on the specifics of the RF-based measurements (e.g., signal strength, transmission delay, phase, et cetera). The platforms are then interested in estimating the value of X , given

¹For instance, imagine if control theory were, from its infancy, developed assuming that a deep learning network was the source of prediction, rather than an ordinary differential equation. Such a control approach would have very different characteristics than those currently in use—most obviously, *time* would not necessarily play a fundamental role in the description of a system. Indeed, even the notion of a vector space would be only distally relevant to such a theory. Instead, the entire theory would hinge on features (potentially spatiotemporally defined) and their descriptions. We are not advocating for such a theory specifically, but are advocating that major advances in agility will require such a level of theoretical reorganization.

the measurement y . Furthermore, they are interested in designing T , i.e., their trajectories, in an online and adaptive manner, based on an online estimation of X , and under resource constraints. Fortunately, there is a large literature on search, information maximization, and volume coverage (e.g., [109, 110, 111, 112, 113, 10, 11]) that can be adapted to this scenario. Nevertheless, applications such as cooperative robotic imaging using low frequency RF signals have substantial challenges.

Underlying Challenges: There are key fundamental challenges preventing us from robotic imaging at these lower frequencies. First, several propagation phenomena, such as scattering off of objects not on the direct line of sight, are largely negligible at x-ray frequencies, while they cannot be neglected at WiFi frequencies. This then requires a more advanced wave interaction model to represent f , which can make solving the imaging problem computationally infeasible, especially on very small platforms. Second, the problem can quickly become considerably under-determined. For instance, consider the simple case that a linear wave approximation suffices to represent f . We then have $y = A(T)X + n$. However, the number of the measurements collected by the agile platforms is typically much smaller than the size of the unknown space due to resource constraints. This results in a considerably under-determined system and severe ambiguity in estimating X .

We then ask: Can utilizing mobile platforms, co-optimizing the path planning and RF sensing, and novel signal processing have the potential to get around some of the aforementioned challenges? Consider mobility as an example. With mobile platforms, we have the possibility of optimizing the positioning of transmitting/receiving antennas (i.e., designing the matrix T). This can create unprecedented possibilities for sensing (and perception) with RF signals. For instance, the platforms can potentially create a very large virtual antenna, with a desired beam pattern of interest, which can best serve their particular sensing and perception task. In other words, there is a rich space to explore at the intersection of mobility and RF-based sensor design for cooperative perception on small platforms.

On the other hand, recent advances in signal processing such as sparse signal processing [114, 115] allow us to get around the under-determined nature the occupancy sensing problem by utilizing sparse representations of the signal of interest. Suppose that X has a sparse representation in another domain, i.e., it can be represented as a linear combination of a small set of vectors: $X = \Phi x$, where x is sparse. Let $\Psi(T) = A(T) \times \Phi$. Sparse signal processing allows us to sample a signal of interest in a compressed manner, thus getting around the seemingly large-size of the unknown space. The underlying principle of it is the fact that “most” signals are compressible to a smaller space, which is the basis for the success of compression methods. And while sparse processing uses sparsity after data has been collected, sparse sensing attempts to exploit sparsity beforehand by reconstructing a signal with very sparse samples. Recent work has shown that this is possible under certain conditions, e.g., when matrix Ψ satisfies the Restricted Isometry Condition (RIC) [116]. Most systems do not automatically satisfy the RIC condition, which leads to the question: Can agile platforms then take advantage of their mobility to successively design T such that Ψ satisfies the RIC condition?

In summary, there is a rich unexplored space at the intersection of agility, RF sensing, and signal processing, which has the potential to enable agile platforms, especially micron-scale platforms, with perception capabilities, under resource constraints and through walls. Fig. 3 shows our preliminary result along these lines where two vehicles produced this 3D image of the area shown in the right figure, based on an approach that builds on the aforementioned principles [117].




Figure 3: 3D robotic imaging through walls with on-board WiFi signals.

Near-Term Important Questions:

1. What are the characteristics of the robotic paths that would satisfy the RIC condition for perception on small platforms?
2. Can agile platforms utilize the underlying localized properties that govern the design of most spaces for perception with limited resources? For instance, we can model the 3D space of interest as a Markov Random Field (MRF) in order to model the spatial dependencies among local neighbors. Using the MRF model, we can then use the Hammersley-Clifford Theorem, for instance, to express the probability distribution of the voxel values in terms of local dependencies, which creates possibilities for solving the perception problem through graph-based propagation approaches.
3. How can MEMS RF sensing and perception enable advanced sensing in otherwise inaccessible areas? Can MEMS RF sensors communicate with each other and form a chain of reliable communications and situational understanding inside an inaccessible building. Can combining such sensors with external mobile vehicles provide higher resolution estimates of the interior of a building?

Throughout this discussion of cooperative, RF-based sensing, agility of mobile systems plays an important role—how can the mobile systems be automated so that they are *efficiently responsive* to both the question that needs answering—e.g., estimate the geometry of the interior of a building—and to data as it is acquired—e.g., the RF signals themselves and the probability distribution over the voxel states describing the building’s interior geometry. Hence, agility for such a cooperative system is not necessarily intrinsically connected to the agility of individual platforms. Instead, agility is the responsiveness of the entire group to the task. This characterization of agility in the context of cooperative systems is applicable to other cooperative tasks as well. As we learn how to integrate sensory and actuation systems into coherent wholes, subsystems will need to actively cooperate, explicitly exploiting each other in pursuit of a task.

Integrated Sensory and Actuation Systems. The development of versatile and robust bio-inspired robotic vehicles, whether terrestrial, aerial or underwater, requires integrated sensory and actuation systems that work collaboratively to achieve the operational demands on locomotion and manipulation. Such collaboration between sensing and actuation requires the two or more systems to actively and passively modulate their elements in both space and time. Both live organisms and man-made autonomous vehicles such as robots, self-driving cars and deep-ocean AUVs, have finite computational resources; they must therefore multitask and develop schemes to dynamically allocate resources to different tasks, such as multi-sensory acquisition and integration, locomotion planning and motor control. Moreover, organisms often employ strategies for navigation in complex uncertain environments that involve switching between egocentric and geocentric schemes.

Egocentric versus geocentric navigation strategies: Egocentric navigation refers to when an organism or robot uses innate or pre-acquired information to move along a path without any realtime feedback. Basically, this scheme executes an actuation strategy with no realtime sensory information or with a strategy that anticipates environmental variability. This strategy requires memory and information storage and can be efficient, but can succeed only in a constant environment. In geocentric navigation, the organism or robot constantly probes its location relative to environmental cues and adjusts its strategy in real time. This strategy requires continuous course corrections and is often slow, but can cope with fluctuating environments. Accurate yet efficient navigation in a variable environment requires switching between these two strategies or blending the two strategies continuously. A key ingredient lies in understanding and optimizing the switching parameters and their dependence on the natural environment.

So how should blending the balance between egocentric and geocentric—or, alternately, between model-based and data-based, or even intrinsic and extrinsic—reasoning be determined? Ultimately, this comes down to managing complexity—how computationally intensive an approach is—while maintaining robustness. But complexity is relative and one implementation of a given idea may be very computationally intensive while another is not. As an example, under what conditions can a sophisticated control policy and its dependence on sensors be simplified all the way down so that no computation is required at all? This idea, which we reference as Cyber-free Autonomous Systems, provides a kind of bound on how much can be achieved through model-based design, while also making clear the role that data, and reconfigurable computing, might play.

Cyber-free Autonomous Systems: An interesting question related to sparsity arises in model-based control: if a control policy—thought of as a set of voxels in a high dimensional state space—happens to be sparse, how can one exploit its

sparsity? And if a control policy is very sparse, is computation required at all? Specifically, do robotic systems need computing in the form of reconfigurable signal processing? What level of complexity is really needed for reliable control, and can that control be pushed away from the CPU and towards the “nervous system” of an autonomous system? For instance, can flight control in a complex fluid be pushed entirely into hardware accelerators embedded into a machine?

As an example, we took the spring-loaded inverted pendulum (SLIP) model (studied extensively in the locomotion community [118, 119, 120, 121]) and automatically extracted a finite-state machine (FSM) representation of the control approach. We did this using the techniques in [122], where an optimal control technique is used to compute an approximation of the policy at a particular state and compression techniques are used to avoid excessive evaluation of the control. The resulting FSM for SLIP locomotion involves only a five state FSM that could be easily embedded in dedicated hardware (e.g., an FPGA or a hardware accelerator). Moreover, the logic of the finite state machine could potentially be directly encoded in the mechanism of locomotion itself (e.g., mechanical ratchets), though this would require yet another computational process to map potential finite state machines to potential physical designs.

The only failures for the example occur when the state is outside of the states of the FSM; these can be corrected by offloading control computation to a CPU. This latter point is critical; the key role that reconfigurable computing plays in this example is as a backstop to the dedicate, simple processing. So computation itself can be thought of as a resource to be dynamically allocated to a control process—something similar to “attention” in animals. Automating this characteristic so that deeply embedded processes can be generally responsible for system behavior—while computationally-intensive capabilities can be recruited as needed—would provide one way to investigate the trade-offs between intrinsic, model-based capabilities and extrinsic, data-based capabilities.

Moreover, there is a variation of cyber-free systems that should be considered that further blurs the distinction between computational control and physically-embedded control. In the event that an FPGA or hardware accelerator is being used, it is entirely possible that the occasional recruitment of computation for a task would naturally lead to rewriting the FPGA code, on comparatively slow time intervals. Self-coding systems that adapt their control policies, while managing complexity, in response to environmental changes, mechanical changes, et cetera, would lead to continuous communication between the slowly-varying intrinsic dynamics and the sparse (in time) use of explicit computation. As an example of a theoretical setting in which this idea is well posed, we next discuss some recent advances in the use of Koopman operators [123] for control.

Active Learning and Spectral Methods. Data-driven methods are good at describing dynamics in an information-rich setting. Hence, we can expect they could be applied to vehicles running when there is substantial data acquisition opportunity. The same methods are less effective for control synthesis and for human-in-the-loop (HITL) systems, partially because a system must be pushed away from equilibrium before data about it can be expected to be meaningful. (The reasons for this essentially come down to the fact that at equilibria information is low—for instance as measured via Fisher information—so long as the equilibria is stable.) As vehicles are in environments that are increasingly poorly understood, we can expect that controlled response to data-driven representations of highly dynamic processes will be important.

Spectral methods, and Koopman operators [123] in particular, have recently become of substantial interest again in the dynamics community. Roughly speaking, spectral methods replace a finite dimensional ordinary differential equation (typically nonlinear) that depends on and maps the state at one time to the state at a later time with an infinite dimensional linear operator $\mathcal{K}$ (called a Koopman operator) that maps functions of the state at one time to functions of the state at a later time. If an ordinary differential equation is in the (x, y) plane, the Koopman operator $\mathcal{K}$ will map functions of (x, y) to functions of (x, y) , and these functions can come from any function space over (x, y) (e.g., polynomials, Fourier, et cetera). For example, functions could include $x, y, xy, x + y, x^7y^3$, and the linear operator $\mathcal{K}$ takes one set of values of these functions and returns a new set of values of these functions, thereby implicitly updating the state itself. Although moving from a finite dimensional setting to an infinite dimensional setting does create computational challenges, the linear representation of the system evolution enables one to generate data-driven models without prior assumptions about the dynamics by examining the spectrum of the data. Specifically, the linear operator can be approximated by a finite set of eigenvalues and eigenfunctions, making the approximation of $\mathcal{K}$ reasonably easy to compute for a finite choice of basis functions.

The advantages to terrestrial agility in this context are substantial. If a model of a mechanical system interacting with its environment can be constructed directly in terms of data, many of the subtle complexities associated with terramechanics could be avoided. The fine-scale differences between different types of granular media would cease to be a major problem for the control, because the modeling would be based on what the vehicle actually does in response

to control inputs rather than focusing on low-level physics-based models. In this context, one could focus entirely on whether or not the model was actionable—that is, whether or not changes in the model lead to changes in control decisions. The challenges in using a data-driven model are many, as discussed below, but the foundation of using data to directly obtain actionable models would enable online systems to depend on the data they are accumulating instead of models developed in environmentally benign circumstances.

Operationalizing Spectral Methods for Online Control of Unmodeled Systems: Little work has been done using Koopman operators for control. One option is to compute linear quadratic controllers in the infinite dimensional space, though that does not necessarily make the problem more tractable. Computing LQ controllers in high dimensions is complicated by computing Riccati equations in high dimensions. Moreover, the interpretation of the basis functions becomes nonobvious, not least because linear systems with linear quadratic cost *always* have global optima, and the state-based nonlinear problem does not. Another possibility is to use the Koopman representation to reconstruct a nonlinear state-dependent ordinary differential equation, and synthesize control for that using nonlinear optimal control. A third possibility is to reconstruct analytic physics-based models—for example, models that treat mechanical symmetries as constraints on the data-driven process—directly from the spectrum. For instance, if the underlying system generating data has Hamiltonian structure, one can directly compute the Hamiltonian from the Koopman representation. From the Hamiltonian, one gets a deterministic model in the form of a discrete time deterministic system (through the discrete-time Legendre transform); this model satisfies all the symmetries that the Hamiltonian imposes (e.g., conservation laws, constraints). In each of these approaches, the quality of the data will play an important role in the quality of the resulting control.

Exploration For Data-Driven Representations: The Koopman formalism (and other related formalisms for data-driven processes) rely on collecting measurements that are informative but not too destructive to the system. For example, controlled exploration that destabilizes a system would be unacceptable for most engineered systems. Measurements all taken at dynamic equilibrium are unlikely to contain much, if any, information about dynamic structure. Measurements taken too far outside of normal operating conditions can create spurious effects and create artificial computation burdens. Balancing these, through formally integrating spectral information needs into routine control, is an important part of maintaining a model during operation. How should one do this?

To investigate, we did an example of Koopman-based control, using a Sphero robot—a sphere with a kinematic vehicle inside it to drive the sphere forward—to drive on the floor [124]. Our approach followed the second approach listed above—we first compute $\mathcal{K}$ and then map that back to a nonlinear state-space representation where classical model-predictive control techniques can be employed. The robot has an outer sphere and inner kinematic device, so the geometry is very sophisticated. Nevertheless, we started by modeling the Sphero as a double integrator system ($\ddot{x} = u_1, \ddot{y} = u_2$) in the (x, y) plane, which at low velocity is reasonable. The reference trajectory was parameterized to be sufficiently aggressive to excite the systems internal nonlinearities that cannot be captured completely by the minimal state representation. Then, with a human in the loop, the Sphero was driven at higher speeds to collect data to compute $\mathcal{K}$. The resulting controller, based on $\mathcal{K}$, demonstrates that at low speeds, the extra information in $\mathcal{K}$ does not impact the control very much—the system is too close to the equilibrium where the control is already performing well. But as the trials involve driving at faster speeds, the control starts to exploit the extra information about nonlinearity in the state-space reconstruction based on $\mathcal{K}$, resulting in several factors of error improvement. A downside of this approach is that we cannot *interpret* the resulting nonlinear system—its nonlinearities arise partially due to the geometry of the robot, partially due to the physics of the robot, and partially due to nonlinearities in the contact physics between the robot and the floor. (This is where the mechanics-based approach listed above could be pivotal.) But the upside is that we see real improvement in control performance with minimal intervention and no explicit modeling effort, while being constrained to real-time operation.

With this example in mind, several questions that one can ask (and hopefully answer) include the following. How should a data-driven system *acquire* data, ignoring at least momentarily optimizing the details of how that data is processed? At what point does the data become *actionable*—that is, under what conditions can a data-driven model be trusted sufficiently to start incorporating model updates into decisions? And how can these models be related to meaningful physical models, exploiting what we do know (e.g., continuity in time, entropy increases, symmetries of motion, energetic passivity of frictional surface interactions, et cetera) while allowing the data to illuminate what we do not know (e.g., details of nonlinear models, dimensions of underlying state space, manifold structure)? These questions are not specific to the Koopman formalism—any data-driven process will eventually have to come to terms with them in one way or another.

3.1 Long Term Recommendations

Based on the above discussions, we have several recommendations for long-term investment.

1. Identify cooperation-dependent tasks—such as those discussed above using RF sensing—and develop cooperative, distributed systems for those tasks. Integrate multiple scales of devices (e.g., MEMS-RF devices and mobile units) into cooperative sensing tasks.
2. Develop a mechanically constrained, data-driven approach to control, using mechanical information to constrain data-driven processes and information measures appropriate for parameter-free processes to drive exploration based on those measures. Use these approaches to develop *active* self-learning systems, that constantly probe their internal models while executing missions.
3. Develop both a theory and practical implementation for cyber-free systems. Part of this could be developing the capacity for computing extremely sparse control policies (e.g., [122]), where only a handful of states in a finite state machine approximate the policy.

4 Energy

The previous two sections implicitly assume that terrestrial systems have access to unlimited energy, or at least great quantities of it. We consider here two opportunities for enhancing agility while managing energy, beyond the typical approaches of trying to design efficient systems with efficient control systems. Specifically, we consider here how we can take advantage of spatial and temporal scales to harvest environmental energy.

4.1 Battery-Free Agile Systems:

Energy drainage is one of the major challenges for the autonomous operation of unmanned vehicles. A battery of a typical ground vehicle, such as a Pioneer robot, lasts around 30 minutes, which can be disruptive to many, perhaps most, practical robotic operations. We ask whether a vehicle can re-charge on the field using existing signals in its environment? While the idea of solar-powered vehicles has been explored, the feasibility of recharging with the existing RF signals in the environment is unclear. Examples of such signals would include those emitted by nearby TV or radio stations, those emitted by other unmanned systems, or those emitted by remote human operators. RF signals are ubiquitous these days and can provide a promising additional source for robotic field charging. Therefore, the main question is if, when, and at what scale it is possible to do such a charging?

New results have recently demonstrated the possibility of wireless charging in very short distances. For instance, [125] has designed a high-efficiency circuit to convert the ambient RF signals to DC current, which is then used to power temperature and humidity meters. Despite relatively little work in this area, however, the idea of using existing RF signals for charging is fairly unexplored at the time being. Here, we are further interested in unmanned vehicles, which presents more challenges than charging a non-mobile platform. Nevertheless, unmanned vehicles have several advantages. First, they can potentially be very small. Second, depending on their mission, they can potentially execute a mission over a long time horizon.

Questions:

1. Can an unmanned vehicle utilize RF signals for powering up its onboard units? How much is fundamentally possible? How far and how strong should the transmitted signal be? What can such harvested energy power up? For instance, can this generated energy be used for communication? Can this generated energy be used for sensing?
2. What does this enable for small-size agile systems, such as cell-size agile platforms? These platforms need little energy to sense or communicate as each has limited capabilities. With the advances in MEMS technology, micro-miniaturization of radio transceivers is becoming a possibility [126]. What possibilities would this then present for the small-size agile platforms in terms of energy harvesting?
3. While the progress on wireless energy harvesting relies heavily on the advances in circuit design, agility can play a key role in how mobile platforms best take advantage of this harvesting technology. Specifically, the

performance of wireless harvesting relies on two key factors: 1) the strength of the available electric field (induced by a nearby transmission), and 2) harvesting potential of the receiver on the platform, which is directly related to the state-of-the-art in circuit design. Consider the first factor. Then, how can a number of platforms, each with a small transmission power, path plan to optimum positions and optimally beamform to generate a strong enough field for a receiving platform to harvest energy? Beamforming is a cooperative communication technique where a number of antennas optimally design their transmission phase to arrive in phase at the receiver, thereby increasing the received signal power considerably. Then, in the context of a number of mobile agile units, how can they best position themselves to optimize the generated beam pattern cooperatively, for the purpose of energy harvesting?

4.2 Micron-Scale Agile Systems

As the RF energy harvesting discussion suggests, there may be energetic advantages possible for cell-sized autonomous systems; small systems may simply be in a better position to exploit energy harvesting techniques. However, to be useful, one would need meaningful, mission-relevant autonomy at such a scale, and likely need to achieve that autonomy without substantial on-board computation. Imagine cell-sized mechanical devices with minimal computation (expressed in terms of actions on bits), minimal actuation (e.g., electrostatic adhesion), and minimal sensing (e.g., gross compass readings). What sensing ranges could such a system have (e.g., inches to thousands of miles)? What could it sense/collect (e.g., chemical agent samples, radioactivity)? How could such a device be powered (e.g., thermal resonators and diodes [127], RF signals as discussed above)? Moreover, what would it mean for a micron-scale system to be agile? And what purpose would such a system serve? For instance, could a collection of motes gain access to a facility, take a distributed picture of its interior, and exit again?

There is no power-based locomotion that can be expected to enable distances multiple order of magnitude larger than the length scale of the locomotor; exploitation of environmental energy sources will be necessary. For instance, chemical adhesion provides relatively low energy control authority. Very small devices could therefore use macro-scale hosts—people, animals, vehicles—for transportation. Moreover, the control goal at these small scales would not be for a particular autonomous device to meet mission criteria; instead, statistical notions of success across all the collection of systems would suffice—we do not need *all* of the devices to achieve a goal if a sufficient number do.

As an example, consider Fig. 4. This simulation has macro-scale agents that are moving randomly in the environment, moving in and out of an enclosed space through a doorway. Moreover, there are autonomous systems with sensing in the form of detection of nearby passing agents and an internal compass. Moreover, they are equipped with the ability to adhere to the agents. With a very simple rule applied to the sensory data computed in a manner similar to the SLIP model mentioned earlier in Section 3, the simulated agents are able to enter the enclosed space, using only the environmental energy (in the form of macro-scale agents moving) for transport.




Figure 4: Three snapshots of simple autonomous systems (red dots) using passerby agents (circles) to locomote to a location (blue $\times$) in a confined space by turning on and off their adhesion control authority.

The example in Fig. 4 is a cartoon of a bigger idea—how can we design extremely simple systems, with simple input-output relationships encodable directly in the physics of the device, to achieve a mission? This will require “algorithmic materials”—materials that implicitly compute and implement responses to state—and the computational

ability to design them. This latter point, that design challenges become increasingly central to many of the questions considered in this study, is a primary point in the long term recommendations we make here.

4.3 Long Term Recommendations

Energy harvesting would enable persistent presence in adversarial conditions, and investing in energy harvesting techniques that target ambient sources such as ubiquitous RF signals and naturally-occurring thermal fluctuations could provide consistent power sources for autonomous systems. Particularly at small scales, this capability would require the ability to match mission needs to physical capabilities (e.g., the molecular machinery available enable cell-sized systems). Developing computational methods that enable the design of such cyber-free systems would radically reshape what autonomy looks like at these scales.

5 Conclusion

This study considers three main areas of investment for improving agility: control technology, sensing and perception, and energy. Some of the recommendations are reasonably obvious if technologically challenging—developing new control techniques to manage uncertainty, cooperative strategies for sensing, and energy harvesting techniques are all seen in current, mainstream research. However, some of the recommendations are of a more radical nature. We argue that control approaches may need to be fundamentally reformulated from scratch to address the multi-faceted requirements that terrestrial agility may generate. We argue that computational synthesis techniques for “cyber-free” approaches to control may be critical to offloading most agility-related control to dedicated processing. Indeed, extremely small systems may be the only systems capable of exploiting energy harvesting, but would require such a cyber-free approach to their control. Lastly, both macro- and micro-scale systems are unlikely to have high quality physics-based models, and therefore the ability to seamlessly integrate data-based techniques into control and perception; data-based control will then require active exploration, subject to safety and mission constraints. The main message of this study is that there are multiple areas of research that could have dramatic impact on practically deployed systems. Many of these are not currently well represented in the literature, suggesting future areas of investment and work.

References

- [1] Isabelle Fantoni, Rogelio Lozano, and Mark W Spong. Energy based control of the Pendubot. *IEEE Trans. Autom. Control*, 45(4):725–729, Apr. 2000.
- [2] Thamer Albahkali, Ranjan Mukherjee, and Tuhin Das. Swing-up control of the Pendubot: an impulse–momentum approach. *IEEE Transactions on Robotics*, 25(4):975–982, Aug. 2009.
- [3] Yury Orlov, Luis T Aguilar, Leonardo Acho, and Adain Ortiz. Swing up and balancing control of Pendubot via model orbit stabilization: Algorithm synthesis and experimental verification. In *Proc. IEEE Conf. Decision and Control*, pages 6138–6143, 2006.
- [4] Xin Xin and Taiga Yamasaki. Energy-based swing-up control for a remotely driven Acrobot: Theoretical and experimental results. *IEEE Trans. Control Syst. Technol.*, 20(4):1048–1056, Jul. 2012.
- [5] X Xin and M Kaneda. Analysis of the energy-based swing-up control of the Acrobot. *Int. J. of Robust and Nonlinear Control*, 17(16):1503–1524, Nov. 2007.
- [6] Anthony M Bloch, Melvin Leok, Jerrold E Marsden, and Dmitry V Zenkov. Controlled Lagrangians and stabilization of the discrete cart-pendulum system. In *Proc. IEEE Conf. Decision and Control and European Control Conf.*, pages 6579–6584, 2005.
- [7] Karl Johan Åström and Katsuhisa Furuta. Swinging up a pendulum by energy control. *Automatica*, 36(2):287–295, Feb. 2000.
- [8] George Mathew and Igor Mezic. Metrics for ergodicity and design of ergodic dynamics for multi-agent system. *Physica D-nonlinear Phenomena*, 240(4-5):432–442, Feb 2011.
- [9] S.E. Scott, T.C. Redd, L. Kuznetsov, I. Mezic, and C.K. Jones. Capturing deviation from ergodicity at different scales. *Physica D*, 238(26):1668–1679, 2009.
- [10] L. Miller, Y. Silverman, M. A. MacIver, and T.D. Murphey. Ergodic exploration of distributed information. *IEEE Transactions on Robotics*, 32(1):36–52, 2016.
- [11] A Wilson, J. Schultz, and T. D. Murphey. Trajectory synthesis for Fisher information maximization. *IEEE Transactions on Robotics*, 30(6):1358–1370, 2014.
- [12] Joseph E Avron and Oren Raz. A geometric theory of swimming: Purcell’s swimmer and its symmetrized cousin. *New Journal of Physics*, 10(6):063016, 2008.
- [13] Joel W Burdick, Jim Radford, and Gregory S Chirikjian. A ‘sidewinding’ locomotion gait for hyper-redundant robots. In *IEEE International Conference on Robotics and Automation*, pages 101–106, 1993.
- [14] Howie M Choset. *Principles of robot motion: theory, algorithms, and implementation*. MIT press, 2005.
- [15] Alessandro De Luca, Giuseppe Oriolo, and Claude Samson. Feedback control of a nonholonomic car-like robot. *Robot Motion Planning and Control*, pages 171–253, 1998.
- [16] Ross L Hatton and Howie Choset. Geometric motion planning: The local connection, stokes theorem, and the importance of coordinate choice. *The International Journal of Robotics Research*, 30(8):988–1014, 2011.
- [17] E. Kanso, J.E. Marsden, C.W. Rowley, and J.B. Melli-Huber. Locomotion of articulated bodies in a perfect fluid. *Journal of Nonlinear Science*, 15:255–289, 2005.
- [18] L.S. Pontryagin, V.G. Boltyanskii, R.V. Gamkrelidze, and E.F. Mischenko. *The Mathematical Theory of Optimal Processes*, K.N. Trirogoff (transl.), L.W. Neustadt (ed.). Interscience Publishers, New York, NY, USA, 1962.
- [19] R. Bellman. On the theory of dynamic programming. In *Proceedings of the National Academy*, volume 38.8, page 716, 1952.

- [20] Daniel Liberzon. *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, Princeton, New Jersey, USA, 2012.
- [21] Frank Allgower, Rolf Findeisen, and Zoltan K Nagy. Nonlinear model predictive control: From theory to application. *J. Chinese Inst. of Chemical Engineers*, 35(3):299–316, May 2004.
- [22] Eduardo F Camacho and Carlos Bordons Alba. *Model Predictive Control*. Springer-Verlag, London, UK, 2 edition, 2013.
- [23] Lars Grüne and Jürgen Pannek. *Nonlinear Model Predictive Control*. Springer, London, UK, 2011.
- [24] Frank Allgöwer and Alex Zheng. *Nonlinear Model Predictive Control*, volume 26. Birkhäuser Basel, Basel, Switzerland, 2000.
- [25] Brian D. O. Anderson and John B. Moore. *Optimal Control: Linear Quadratic Methods*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1990.
- [26] Emanuel Todorov and Michael I Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235, 2002.
- [27] A. Ansari and T. D. Murphey. Sequential Action Control: Closed-form optimal control for nonlinear and hybrid systems. *IEEE Transactions on Robotics*, 32(5):1196–1214, 2016.
- [28] Daniela Rus and Michael T Tolley. Design, fabrication and control of soft robots. *Nature*, 521(7553):467–475, 2015.
- [29] Robert F Shepherd, Filip Ilievski, Wonjae Choi, Stephen A Morin, Adam A Stokes, Aaron D Mazzeo, Xin Chen, Michael Wang, and George M Whitesides. Multigait soft robot. *Proceedings of the National Academy of Sciences*, 108(51):20400–20403, 2011.
- [30] A. Mavrommati, A. Ansari, and T. D. Murphey. *Trends in Control and Decision-Making for Human-Robot Collaboration Systems*, chapter Assistive Optimal Control-on-Request with Application in Standing Balance Therapy and Reinforcement, pages 131–156. Springer, 2017. Eds. Y. Wang and F. Zhang.
- [31] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press Cambridge, 1998.
- [32] Mattia Gazzola, Andrew A Tchieu, Dmitry Alexeev, Alexia de Brauer, and Petros Koumoutsakos. Learning to school in the presence of hydrodynamic interactions. *Journal of Fluid Mechanics*, 789:726–749, 2016.
- [33] T. Caldwell and T. D. Murphey. Projection-based iterative mode scheduling for switched systems. *Nonlinear Analysis: Hybrid Systems*, 21:59–83, 2016.
- [34] Henrik Axelsson, Y Wardi, Magnus Egerstedt, and El Verriest. Gradient descent approach to optimal mode scheduling in hybrid dynamical systems. *Journal of Optimization Theory and Applications*, 136(2):167–186, 2008.
- [35] Yorai Wardi, Magnus Egerstedt, and Muhammad Umer Qureshi. Hamiltonian-based algorithm for relaxed optimal control. In *IEEE Conference on Decision and Control (CDC)*, pages 7222–7227, 2016.
- [36] Sorin C Bengea and Raymond A DeCarlo. Optimal control of switching systems. *Automatica*, 41(1):11–27, 2005.
- [37] H. J. Kushner and P. G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, London, UK, UK, 1992.
- [38] V. A. Huynh, S. Karaman, and E. Frazzoli. An incremental sampling-based algorithm for stochastic optimal control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2865–2872, May 2012.
- [39] P. Chaudhari, S. Karaman, D. Hsu, and E. Frazzoli. Sampling-based algorithms for continuous-time POMDPs. In *American Control Conference*, pages 4604–4610, June 2013.

- [40] V. A. Huynh, S. Karaman, and E. Frazzoli. An incremental sampling-based algorithm for stochastic optimal control. *International Journal of Robotic Research*, 35(4):305–333, 2016.
- [41] W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Springer, New York, 1993.
- [42] W.H. Fleming. Exit probabilities and optimal stochastic control. *Applied Math. Optim*, 9:329–346, 1971.
- [43] W. H. Fleming and W. M. McEneaney. Risk-sensitive control on an infinite time horizon. *SIAM J. Control Optim.*, 33:1881–1915, November 1995.
- [44] P. Dai Pra, L. Meneghini, and W. Runggaldier. Connections between stochastic control and dynamic games. *Mathematics of Control, Signals, and Systems (MCSS)*, 9(4):303–326, 1996-12-08.
- [45] S. K. Mitter and N. J. Newton. A variational approach to nonlinear estimation. *SIAM J. Control Optim.*, 42(5):1813–1833, May 2003.
- [46] E.A Theodorou and E. Todorov. Relative entropy and free energy dualities: Connections to path integral and KL control. In *the Proceedings of IEEE Conference on Decision and Control*, pages 1466–1473, Dec 2012.
- [47] D.A Braun, P.A Ortega, E. A. Theodorou, and S. Schaal. Path integral control and bounded rationality. In *IEEE Symposium on Adaptive Dynamic Programming And Reinforcement Learning (ADPRL)*, pages 202–209, April 2011.
- [48] E. A. Theodorou. Nonlinear stochastic control and information theoretic dualities: Connections, interdependencies and thermodynamic interpretations. *Entropy*, 17(5):3352, 2015.
- [49] O. Arslan, E. A. Theodorou, and P. Tsiotras. Information-theoretic stochastic optimal control via incremental sampling-based algorithms. In *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, pages 1–8, Dec 2014.
- [50] E. Todorov. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*, 106(28):11478–11483, 2009.
- [51] E.A. Theodorou, J. Buchli, and S. Schaal. A generalized path integral approach to reinforcement learning. *Journal of Machine Learning Research*, (11):3137–3181, 2010.
- [52] H. J. Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment*, 11:P11011, 2005.
- [53] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Aggressive driving with model predictive path integral control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1433–1440, May 2016.
- [54] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. Rehg, B. Boots, and E.A. Theodorou. Information theoretic MPC for model-based reinforcement learning. In *IEEE Conference on Robotics and Automation (ICRA)*, 2017.
- [55] Paul Dupuis, Konstantinos Spiliopoulos, and Hui Wang. Importance sampling for multiscale diffusions. *Multi-scale Modeling & Simulation*, 10(1):1–27, 2012.
- [56] Carsten Hartmann, Christof Schütte, Marcus Weber, and Wei Zhang. Importance sampling in path space for diffusion processes with slow-fast variables. *Probability Theory and Related Fields*, pages 1–52, 2017.
- [57] Daniel Liberzon. *Switching in Systems and Control*. Springer Science & Business Media, 2012.
- [58] J. I. Alora, A. Gorodetsky, S. Karaman, Y. Marzouk, and N. Lowry. Automated synthesis of low-rank control systems from sc-LTL specifications using tensor-train decompositions. In *IEEE Conference on Decision and Control (CDC)*, pages 1131–1138, 2016.

- [59] Alex Gorodetsky, Sertac Karaman, and Youssef Marzouk. Efficient high-dimensional stochastic optimal motion control using tensor-train decomposition. In *Proceedings of Robotics: Science and Systems*, Rome, Italy, July 2015.
- [60] D. H. Jacobson and D. Q. Mayne. *Differential Dynamic Programming*. American Elsevier Pub. Co., New York, 1970.
- [61] E. Todorov and Weiwei Li. A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *American Control Conference*, pages 300–306 vol. 1, June 2005.
- [62] Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *International Conference on Informatics in Control, Automation, and Robotics*, pages 222–229, 2004.
- [63] E.A. Theodorou, Y. Tassa, and E. Todorov. Stochastic differential dynamic programming. In *American Control Conference*, pages 1125–1132, 2010.
- [64] J. E. Marsden and M. West. Discrete mechanics and variational integrators. *Acta Numerica*, 10:357–514, 2001.
- [65] E. Hairer, C. Lubich, and G. Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer-Verlag, 2006.
- [66] Nawaf Bou-Rabee and Houman Owhadi. Stochastic variational integrators. *arXiv preprint arXiv:0708.2187*, 2007.
- [67] E. Johnson, J. Schultz, and T. Murphey. Structured linearization of discrete mechanical systems for analysis and optimal control. *IEEE Transactions on Automation Science and Engineering*, 12(1):140–152, 2015.
- [68] H.O. Fattorini. *Infinite Dimensional Optimization and Control Theory*. Cambridge studies in advanced mathematics. Cambridge University Press, 1999.
- [69] A. Bensoussan, G. Da Prato, M.C. Delfour, and S. Mitter. *Representation and Control of Infinite Dimensional Systems*. Systems & Control: Foundations & Applications. Birkhäuser Boston, 2007.
- [70] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In S. A. Solla, T. K. Leen, and K.-R. Müller, editors, *Advances in Neural Processing Systems 12*. MIT Press, 2000.
- [71] J. Peters and S. Schaal. Policy gradient methods for robotics. In *Proceedings of the IEEE International Conference on Intelligent Robotics Systems (IROS 2006)*, 2006.
- [72] J. Peters and S. Schaal. Natural actor critic. *Neurocomputing*, 71(7-9):1180–1190, 2008.
- [73] J. Kober and J. Peters. Policy search for motor primitives. In D. Schuurmans, J. Benigio, and D. Koller, editors, *Advances in Neural Information Processing Systems 21 (NIPS 2008)*. Cambridge, MA: MIT Press, 2008.
- [74] F. Stulp, E.A. Theodorou, and S. Schaal. Reinforcement learning with sequences of motion primitives for robust manipulation. *IEEE Transactions on Robotics*, 28(6):1360–1370, 2012.
- [75] J. Buchli, F. Stulp, E. A. Theodorou, and S. Schaal. Learning variable impedance control. *International Journal of Robotics Research*, pages 820–833, 2011.
- [76] E. Todorov. Policy gradients in linearly-solvable MDPs. In John D. Lafferty, Christopher K. I. Williams, John Shawe-Taylor, Richard S. Zemel, and Aron Culotta, editors, *NIPS*, pages 2298–2306. Curran Associates, Inc., 2010.
- [77] Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO (1)*, pages 222–229, 2004.
- [78] J. Morimoto and C. Atkeson. Minimax differential dynamic programming: An application to robust biped walking. In *Advances in Neural Information Processing Systems 15*. MIT Press, Cambridge, MA, 2002.

- [79] E. Todorov and W. Li. A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *Proceedings of the American Control Conference (ACC)*, pages 300–306, 2005.
- [80] Y. Pan and E.A. Theodorou. Probabilistic differential dynamic programming. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems* 27, pages 1907–1915. Curran Associates, Inc., 2014.
- [81] Y. Tassa, N. Mansard, and E. Todorov. Control-limited differential dynamic programming. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1168–1175, May 2014.
- [82] Sergey Levine and Pieter Abbeel. Learning neural network policies with guided policy search under unknown dynamics. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems* 27, pages 1071–1079. Curran Associates, Inc., 2014.
- [83] Grady Williams, Andrew Aldrich, and Evangelos A Theodorou. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40:2:344–357, 2017.
- [84] Jin Feng. Large deviation for diffusions and Hamilton-Jacobi equation in Hilbert spaces. *The Annals of Probability*, 34(1):321–385, 2006.
- [85] Amarjit Budhiraja, Paul Dupuis, and Vasileios Maroulas. Large deviations for infinite dimensional stochastic dynamical systems. *The Annals of Probability*, 36(4):1390–1420, 2008.
- [86] Giuseppe Da Prato and Arnaud Debussche. Control of the stochastic Burgers model of turbulence. *SIAM Journal on Control and Optimization*, 37(4):1123–1149, 1999.
- [87] C. Leber, B. Geib, and H. Litz. High frequency trading acceleration using FPGAs. In *International Conference on Field Programmable Logic and Applications*, pages 317–322. IEEE Computer Society, 2011.
- [88] D. Soudry, D. Di Castro, A. Gal, A. Kolodny, and S. Kvatinsky. Memristor-based multilayer neural networks with online gradient descent training. *IEEE Transactions on Neural Networks and Learning Systems*, 26(10):2408–2421, Oct 2015.
- [89] M. Al-Shedivat, R. Naous, G. Cauwenberghs, and K. N. Salama. Memristors empower spiking neurons with stochasticity. *IEEE J. Emerg. Sel. Topics Circuits Syst.*, 5(2):242–253, 2015.
- [90] R. Naous, M. Al-Shedivat, and K. N. Salama. Stochasticity modeling in memristors. *IEEE Transactions on Nanotechnology*, 15(1):15–28, Jan 2016.
- [91] D. Kudithipudi, Q. Saleh, C. Merkel, J. Thesing, and B. Wysocki. Design and analysis of a neuromemristive reservoir computing architecture for biosignal processing. *Frontiers in Neuroscience*, 9:502, 2016.
- [92] G. S. Snider. Spike-timing-dependent learning in memristive nanodevices. In *2008 IEEE International Symposium on Nanoscale Architectures*, pages 85–92, June 2008.
- [93] S. Choudhary, S. Sloan, S. Fok, A. Neckar, E. Trautmann, P. Gao, T. Stewart, C. Eliasmith, and K. Boahen. Silicon Neurons That Compute. In *International Conference on Artificial Neural Networks*, Lecture Notes in Computer Science, pages 121–128. Springer Berlin Heidelberg, 2012.
- [94] W Brockett. Minimum attention control. In *IEEE Conference on Decision and Control*, pages 2628–2632, 1997.
- [95] Alexander Borst, Juergen Haag, and Dierk F Reiff. Fly motion vision. *Annual Review of Neuroscience*, 33:49–70, 2010.
- [96] Michael H Dickinson, Claire T Farley, Robert J Full, MAR Koehl, Rodger Kram, and Steven Lehman. How animals move: an integrative view. *Science*, 288(5463):100–106, 2000.
- [97] Gaby Maimon, Andrew D Straw, and Michael H Dickinson. Active flight increases the gain of visual motion processing in drosophila. *Nature neuroscience*, 13(3):393–399, 2010.

- [98] Brendan Colvert and Eva Kanso. Fishlike rheotaxis. *Journal of Fluid Mechanics*, 793:656–666, 2016.
- [99] Brendan Colvert, Kevin Chen, and Eva Kanso. Local flow characterization using bioinspired sensory information. *Journal of Fluid Mechanics*, 2017.
- [100] Brendan Colvert, Kevin Chen, and Eva Kanso. Bioinspired sensory systems for shear flow detection. *Journal of Nonlinear Science*, 2017.
- [101] Guido Dehnhardt, Björn Mauck, Wolf Hanke, and Horst Bleckmann. Hydrodynamic trail-following in harbor seals (*phoca vitulina*). *Science*, 293(5527):102–104, 2001.
- [102] Geoffrey R Spedding. Wake signature detection. *Annual Review of Fluid Mechanics*, 46:273–302, 2014.
- [103] Leif Ristroph, James C Liao, and Jun Zhang. Lateral line layout correlates with the differential hydrodynamic pressure on swimming fish. *Physical Review Letters*, 114(1):018102, 2015.
- [104] Jeannette Yen, Norman T Nicoll, et al. Setal array on the first antennae of a carnivorous marine copepod, *euchaeta norvegica*. *J Crustac Biol*, 10:218–224, 1990.
- [105] Jeannette Yen and J Rudi Strickler. Advertisement and concealment in the plankton: what makes a copepod hydrodynamically conspicuous? *Invertebrate Biology*, pages 191–205, 1996.
- [106] Jeannette Yen, Marc J Weissburg, and Michael H Doall. The fluid physics of signal perception by mate-tracking copepods. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 353(1369):787, 1998.
- [107] Jeannette Yen, Anne C Prusak, Michael Caun, Michael Doall, Jason Brown, and J Rudi Strickler. Signaling during mating in the pelagic copepod, *temora longicornis*. In *Handbook of Scaling Methods in Aquatic Ecology: Measurement, Analysis, Simulation*, pages 149–159. CRC Press, 2003.
- [108] Jeannette Yen, Jennifer K Sehn, Kimberly Catton, Andrew Kramer, and Orlando Sarnelle. Pheromone trail following in three dimensions by the freshwater copepod *hesperodiaptomus shoshone*. *Journal of Plankton Research*, 33(6):907–916, 2011.
- [109] P. Yang, R. A. Freeman, and K. M. Lynch. Multi-agent coordination by decentralized estimation and control. *IEEE Transactions on Automatic Control*, 53(11):2480–2496, 2008.
- [110] T. H. Chung, V. Gupta, J. W. Burdick, and R. M. Murray. On a decentralized active sensing strategy using mobile sensor platforms in a network. In *IEEE Int. Conf. on Decision and Control (CDC)*, pages 1914–1919, 2004.
- [111] A. J. Davison and D. W. Murray. Simultaneous localization and map-building using active vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):865–880, 2002.
- [112] H. Zhou and S. Sakane. Mobile robot localization using active sensing based on Bayesian network inference. *Robotics and Autonomous Systems*, 55(4):292–305, 2007.
- [113] Enric Galceran and Marc Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12):1258 – 1276, 2013.
- [114] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information theory*, 52(2):489–509, 2006.
- [115] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information theory*, 52(4):1289–1306, 2006.
- [116] D. Needell and R. Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. *Foundations of Computational Mathematics*, 9(3):317–334, 2009.
- [117] C.R. Karanam and Y. Mostofi. 3D through-wall imaging with unmanned aerial vehicles using WiFi. In *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2017.

- [118] M. Hutter, C.D. Remy, M.A. Hoepflinger, and R. Siegwart. ScarLETH: design and control of a planar running robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 562–567, Sept 2011.
- [119] B. Dadashzadeh, H.R. Vejdani, and J. Hurst. From template to anchor: A novel control strategy for spring-mass running of bipedal robots. In *Proc. IEEE Conf. Intelligent Robots and Syst.*, pages 2566–2571, Sept. 2014.
- [120] Devin Koepl and Jonathan Hurst. Impulse control for planar spring-mass running. *J. of Intelligent & Robotic Syst.*, 74(3-4):589–603, Jun. 2014.
- [121] Ioannis Poulakakis and Jessie W Grizzle. The spring loaded inverted pendulum as the hybrid zero dynamics of an asymmetric hopper. *IEEE Trans. Autom. Control*, 54(8):1779–1793, Aug. 2009.
- [122] A. Mavrommati and T. D. Murphey. Automatic synthesis of control alphabet policies. In *IEEE Conference on Automation Science and Engineering*, pages 313 – 320, 2016.
- [123] Marko Budišić, Ryan Mohr, and Igor Mezić. Applied Koopmanism. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(4):047510, 2012.
- [124] G. de la Torre, I. Abraham, and T. D. Murphey. Model-based control using data-driven models. In *Robotics: Science and Systems*, 2017.
- [125] U. Olgun, C. Chen, and J. Volakis. Design of an efficient ambient wifi energy harvesting system. *IET Microwaves, Antennas & Propagation*, 6(11):1200–1206, 2012.
- [126] Toward the Creation of the World’s Smallest Radio.
https://www.nrl.navy.mil/content_images/05Acoustics_Houston.pdf.
- [127] Anton L Cottrill and Michael S Strano. Analysis of thermal diodes enabled by junctions of phase change materials. *Advanced Energy Materials*, 5(23), 2015.