Future Directions of Intelligent Physical Systems

A Workshop on the Foundations of Intelligent Sensing, Action and Learning (FISAL)

October 19–20, 2015

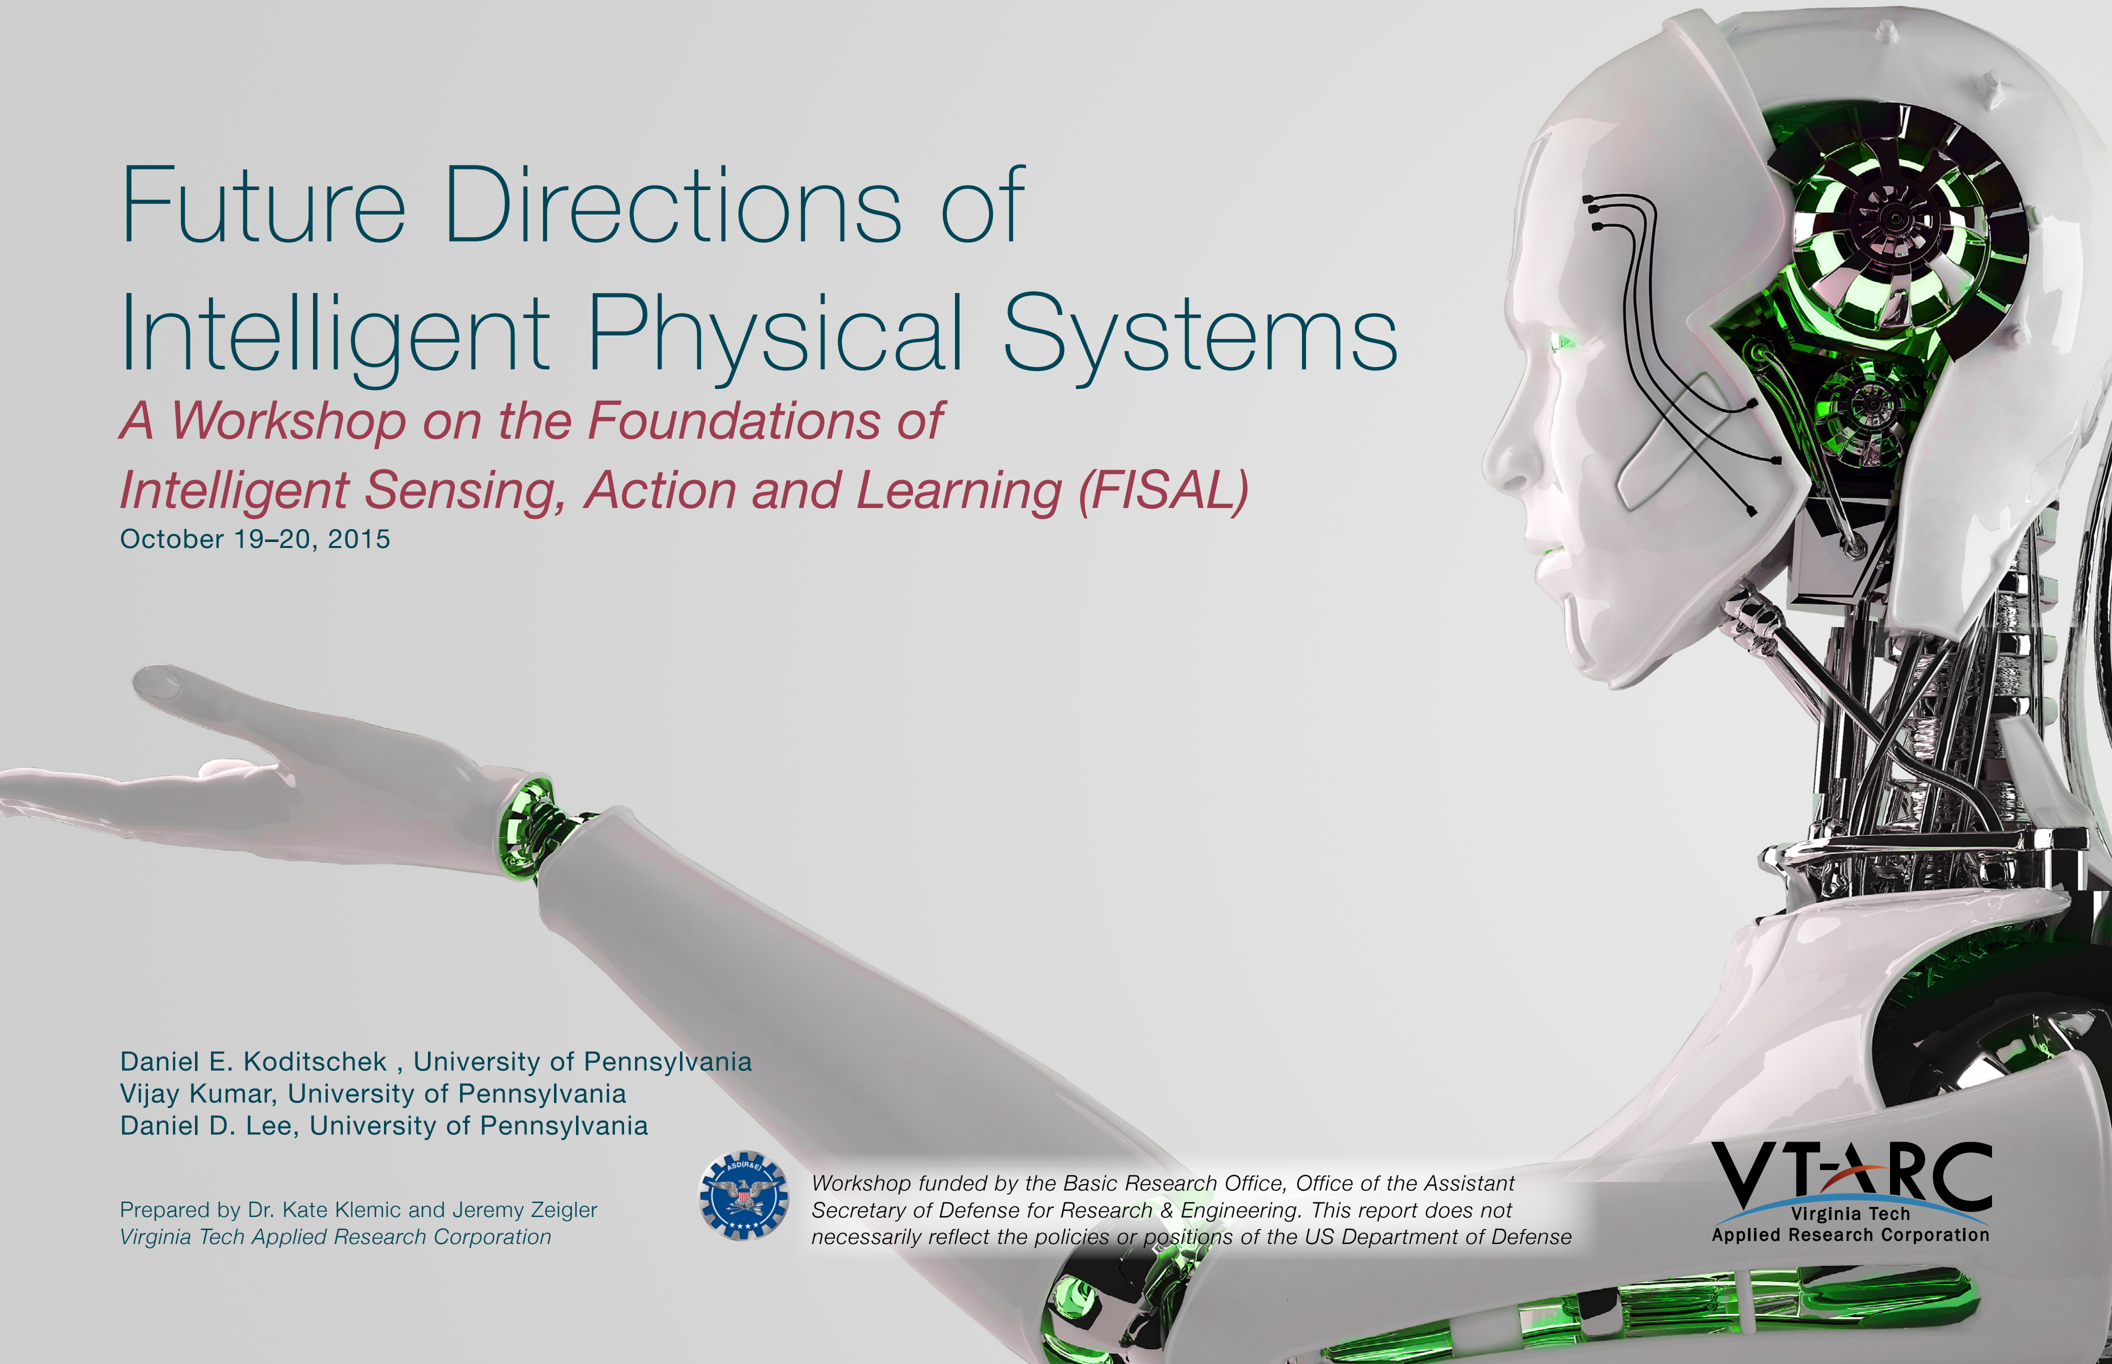

Daniel E. Koditschek , University of Pennsylvania
Vijay Kumar, University of Pennsylvania
Daniel D. Lee, University of Pennsylvania

Prepared by Dr. Kate Klemic and Jeremy Zeigler
Virginia Tech Applied Research Corporation

Workshop funded by the Basic Research Office, Office of the Assistant Secretary of Defense for Research & Engineering. This report does not necessarily reflect the policies or positions of the US Department of Defense

VT-ARC
Virginia Tech
Applied Research Corporation

> Innovation is the key to the future, but basic research is the key to future innovation.
>
> *–Jerome Isaac Friedman, Nobel Prize Recipient (1990)*

Preface

Over the past century, science and technology has brought remarkable new capabilities to all sectors of the economy; from telecommunications, energy, and electronics to medicine, transportation and defense. Technologies that were fantasy decades ago, such as the internet and mobile devices, now inform the way we live, work, and interact with our environment. Key to this technological progress is the capacity of the global basic research community to create new knowledge and to develop new insights in science, technology, and engineering. Understanding the trajectories of this fundamental research, within the context of global challenges, empowers stakeholders to identify and seize potential opportunities.

The Future Directions Workshop series, sponsored by the Basic Research Office of the Office of the Assistant Secretary of Defense for Research and Engineering, seeks to examine emerging research and engineering areas that are most likely to transform future technology capabilities.

These workshops gather distinguished academic and industry researchers from the world's top research institutions to engage in an interactive dialogue about the promises and challenges of these emerging basic research areas and how they could impact future capabilities. Chaired by leaders in the field, these workshops encourage unfettered considerations of the prospects of fundamental science areas from the most talented minds in the research community.

Reports from the Future Direction Workshop series capture these discussions and therefore play a vital role in the discussion of basic research priorities. In each report, participants are challenged to address the following important questions:

- How might the research impact science and technology capabilities of the future?
- What is the possible trajectory of scientific achievement over the next 10–15 years?
- What are the most fundamental challenges to progress?

This report is the product of a workshop held October 19–20, 2015 at the University of Pennsylvania in Philadelphia, PA on Future Directions in Intelligent Physical Systems. It is intended as a resource to the S&T community including the broader federal funding community, federal laboratories, domestic industrial base, and academia.

Executive Summary

The nascent science of Intelligent Physical Systems entails the study of *agents* that can act upon their physical environment by use of perceptual and reasoning processes that also enable them to interact with human and other partners. This new discipline promises future generations of machines that exhibit unprecedented abilities to work for and with us, while fostering new understanding of physical science and engineering and of biology and human psychology.

On October 19–20, 2015, a workshop was held at the University of Pennsylvania to examine the prospects of Intelligent Physical Systems and to scope a research trajectory for the next two decades. Some two dozen prominent researchers in robotics, machine learning and perception, as well as allied areas of systems theory (control, signal processing) and life science (integrative biology, cognitive science) gathered to discuss and debate the opportunities and challenges of the field. They proposed a trajectory of research to overcome the challenges and meet the opportunities. This report is the outcome of those discussions, presented as "Foundations of Intelligent Sensing, Action and Learning (FISAL)".

The workshop participants divided the field of Intelligent Physical Systems into three components:

Action – the study of how an agent's physical and behavioral properties enable it to act and manipulate the world by strategically integrating *form and function*.

Reason – the study of how an agent's sensing and computational properties enable it to represent and predict the world around it.

Interaction – the study of how an agent exists in a bigger world, following the goals and priorities of human partners.

The participants reviewed the history of synthetic science, the evolution of computer science, the origins of the information technology revolution and cybernetics—its creation, degeneration and perhaps near future revival. They reviewed the fundamental breakthroughs, especially in statistical inference, robotics, machine learning, and computer vision that have led to recent commercial successes, like self-driving cars and service robots. However, they find these technologies to be fundamentally limited in their ability to integrate reasoning about the physical world with higher level abstract knowledge.

The participants determined that the goal of the field is to develop systems with intelligent action, reason and interaction that can address complex phenomena, act in environments with complex physics, reason about complex, dynamic worlds and carry out missions as part of large, complex teams.

The participants identified three central gaps in current understanding that requires new research to meet this goal:

Symbolic gap – between the design of bodies and behaviors that impedes both our ability to allocate tasks between form and function in a given environment, as well as our agents' abilities to solve novel problems.

Semantic gap – between symbolic representations meaningful to humans and presently available abstractions of the sensorimotor flows that an embodied agent exchanges with its physical environment.

Representational gap – between the common sense understanding of such abstractions as intent, trust, and motivation that underlie human collaboration and any available computationally effective framework for their symbolic expression and manipulation.

New research paths are expected to produce fundamental discoveries and insights over the next decades. Particular emphasis will be on new machine architectures, with specific awareness of the environment and social interaction. New research will focus on:

Architecture – developing situational and self-awareness, multi-tasking, and representational capabilities that push machine architectures toward the new capabilities.

Environment – increasing architectural capacity relative to environmental complexity to offer an agent and its user a greater understanding of the agent's fitness for that surrounding.

Social interaction – developing new theories of intent, trust and collaboration that lie at the heart of the physically mediated social dynamics necessary to field a human-machine team.

These research areas will include improvements in system modeling, testbeds/training data, benchmarking, verification methods, soft robotics, better manipulation/sensors, and enhanced collaboration across disciplines.

There was significant discussion about the fundamental nature of the field, whether it is a science at all, how it should be formed, and what institutional structure should be formed. In particular, some participants believe that a formal science is required so that the research rests on sound mathematical underpinnings while others are convinced that empirical discovery is the only reliable path forward.

Participants agreed that advancement of the field will require sustained programs that invite, and even require, collaboration between experts in broadly diverse fields, as well as promote introspection concerning the mix of formal theory and novel experiment and the role of each toward the foundations of physical intelligence.

Introduction

Physical intelligence research traces its origins to postwar *cybernetics* research of the mid-20th century that emphasized mathematical foundations and computational abstraction of specific physical settings. By the late 20th century this research had dispersed into a set of specializations in engineering (control, communications, signal processing) and computing (AI, machine learning). The rise of robotics now offers a compelling physical substrate on which those disciplines might join to lay new foundations that target synthesis rather than mere analysis of physical intelligence. For example, the robotic forklift in *Figure 1* exemplifies modern architectures for intelligent sensing, action, learning, and the importance of their interactivity. It is capable of interpreting natural language instructions from a human partner in the context of the robot's knowledge of what it can see and do. This allows people to work alongside the robot just as they would human teammates.

"What I cannot create, I do not understand."

–Richard Feynman

Imagine a team of earth moving machines that "understand" tunnels, ditches, foundations, and so on, as well as the sensitivity of their safe excavation to ambient soil or rubble properties. In addition to their particular expertise, they are also endowed with a generalist "common sense" view of their environment and their intended role within it. They can be gestured at—or even pushed and shoved—by human co-workers who must abruptly alter strategy in a first responder setting. They likely have legs for negotiating broken, unstable terrain, but their bucketed, tool and sensor-tipped appendages are in no way zoomorphic. Nevertheless, their users' experience is something akin to interacting with a team of rescue-intent, tightly cooperative dogs, mules and oxen.

The scientific breakthroughs underlying such technology are still distant. For example, the ability of a dog-like scouting probe to recognize that water has begun to seep into the forward point of an emerging excavation site bespeaks a breakthrough in our present *semantic gap* in how intelligent systems reason about the world. A leap beyond present learning methods is needed to enable the rapid inference and computationally effective semantic representation of a "water present" condition from minute traces of evidence. New science is needed to enable intelligent physical systems to reason about and learn complex models of the world from tiny amounts of data *(Figure 2)*. Or the ability of that machine to scramble out of a narrow side gulley and find its way back to a point of contact in GPS-denied and communications-disrupted settings. A theory of intelligent action is required to reason about shifting and sliding surfaces to enable *proprioceptively* adept agile body motions that can "surf up" the rubble and "invent" new self-catapulting maneuvers *(Figure 3)*. Further, the ability of that machine's *mental model* to anticipate the humans' intended excavation path through this far forward point greatly exceeds strategic and teaming representations presently understood within the HRI community. Finally, such a machine's knowledge of self-capability coupled with the ability to estimate terrain difficulty and weigh the odds of failure against the urgency of the team's mission is beyond our present understanding of how to characterize environmental complexity relative to a task domain.

Figure 1 – Intelligent robotic forklift responds to human language instruction *[1]*

While distant, the scientific basis underlying these advances is by no means unimaginable. The participants find it quite plausible that many of the near term steps proposed below can initiate paths toward the critical antecedent discoveries and insights over the next decade. For example, the research proposed below in “Future Topics in Physical Intelligence Research” focuses on the situational and self-awareness, multi-tasking capabilities, and representational advances key to pushing our machine architectures toward the capabilities just imagined. The direction of inquiry proposed by the participants aims to achieve a titration of architectural capacity relative to environmental complexity and to offer steps toward an agent’s situational calculus that seems clearly antecedent to the future scenario under consideration. The theories of intent, trust and collaboration appear to lie at the heart of the physically mediated social dynamics central to our ever fielding such a human-machine team with confidence in such life-critical settings.

As the discipline grows, the modes of inquiry and standards of intellectual advancement will need to be defined. A group of participants embrace a formal notion of “synthetic science” and its mathematical underpinnings, while others argue for an empirically driven synthetic science advanced by discovery as the only reliable path forward. (A review of the opposing views expressed by participants is detailed in Appendix IV). Those arguing for a formal grounding stress the need for verification algorithms that can take advantage of actual data (in terms of finding counter examples, reasoning about data coverage, and suggesting additional targeted data collections), as well as extending theoretical foundations to achieve the ability to reason about the robustness of systems that adapt and learn. Those favoring a largely empirically driven agenda warn of the potential stultification of ideas by prematurely advanced theory and argue “in defense of hacks.”

On the topic of appropriate institutional structures, there was a considerable diversity of opinion. Some favored “bottom up” settings: creating undergraduate degree programs to force out fundamental problems agreements and disagreements; building standardized testbeds and measuring success through their degree of utilization. Others favored “top-down” structures with unified thematic focus capable of formalizing what is already known and pushing ahead on specific topics via collaboration by established researchers.

Participants broadly agree that new fundamental theories and science are needed in order to develop the next generation of Intelligent Physical Systems that will address society’s expectations for safe, useful, and soundly functioning, embodied and socially embedded technologies. **The core challenge is to develop theories of intelligent action, reason and interaction that can address complex phenomena, guiding the development of robustly working systems that act in environments with complex physics, can reason about complex, dynamic worlds and can carry out missions as part of large, complex teams.**

The research will not only realize future machines possessing unprecedented abilities in perception, action and lifelong learning, but will also help elucidate the relation of physical science and engineering to biology and human psychology. The following sections will elaborate on the challenges and opportunities of the field of Intelligent Physical Systems, and how the obstacles can be overcome.

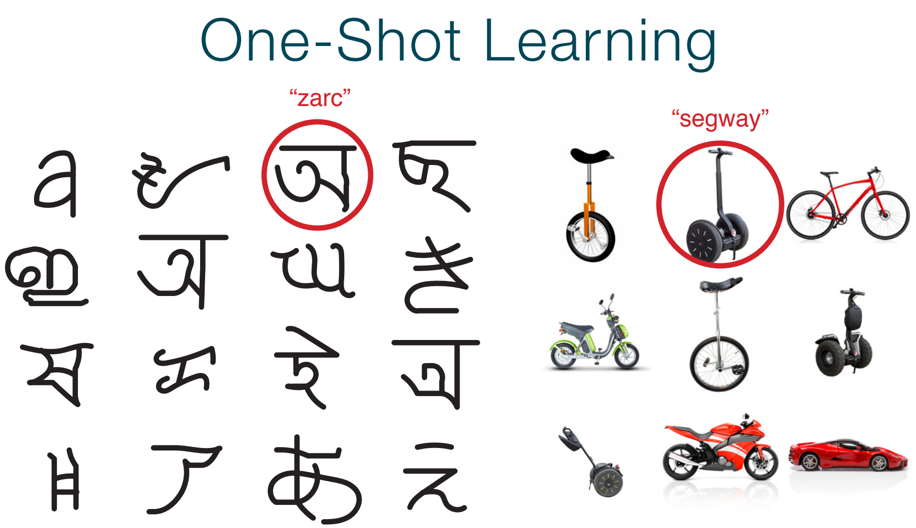

Figure 2 – Advances in machine learning have yielded systems capable of acquiring novel concepts from very few examples. [2]

Figure 3 – Intelligent physical systems can invent new ways of using their bodies to engage the ever-changing, complex environment, much like squirrels, for example, that adaptively catapult off unfamiliar structures, bettering their leap onto a food bearing perch. [3]

Foundations of Intelligent Sensing, Action and Learning

The field of Intelligent Physical Systems has seen some success in developing theories to address specific and narrow problems, exemplified, for instance, by the recent emergence of industrial and consumer products that exhibit some attributes of intelligent action and reasoning. The workshop provided a valuable opportunity for experts from engineering, computing and life sciences to consider ongoing research across these disciplines and to determine how reformulation and coordination of efforts can establish the Foundations of Intelligent Sensing, Action and Learning (FISAL). This section articulates the science base identified by workshop participants to achieve the next generation of *intelligent, interactive, social, sensorimotor systems* comprised of three main components: action, reason and interaction.

Action

Action is the study of how an agent's physical and behavioral properties enable it to move and manipulate the world by strategically integrating form and function. For example, the morphology of the gecko setae *(Figure 4)* is exquisitely suited for interaction with its environment. This is a characteristic concern of roboticists with strong links to integrative biology and *animal cognition*.

Intelligent action requires theories of how a physical body can act in the spatial and temporal context of a physical environment. The participants use the term "sensorimotor system" in place of "robot" to illustrate the strong consensus that studying instances of such agency in both artificial and living systems is essential to fundamental progress *(Figure 5)*. It contrasts with traditional AI research which has characteristically abstracted away the underlying physical embodiment in exchange for intuitive representations and the computational efficiency they afford. **New fundamental theories are required that can model complex physical embodiment and elucidate the relationship of form with function to explain the effects on intelligence.**

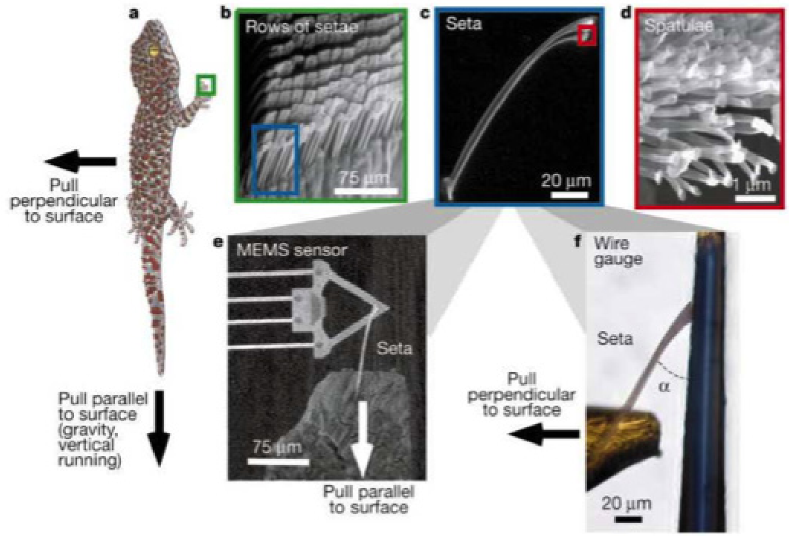

Figure 4 – Animal bodies are exquisitely suited to their behavioral repertoire in the context of specific habitats. The accomplished interaction of morphology with environment accounts in no small part for their physical intelligence. The Gecko's sticky toes offer one example of how extraordinary form (these remarkably built hierarchical structures offer mechanically programmable adhesion arising from multiple, complex subsystem interactions across six orders of magnitude length scales) confers intelligent function. [4]

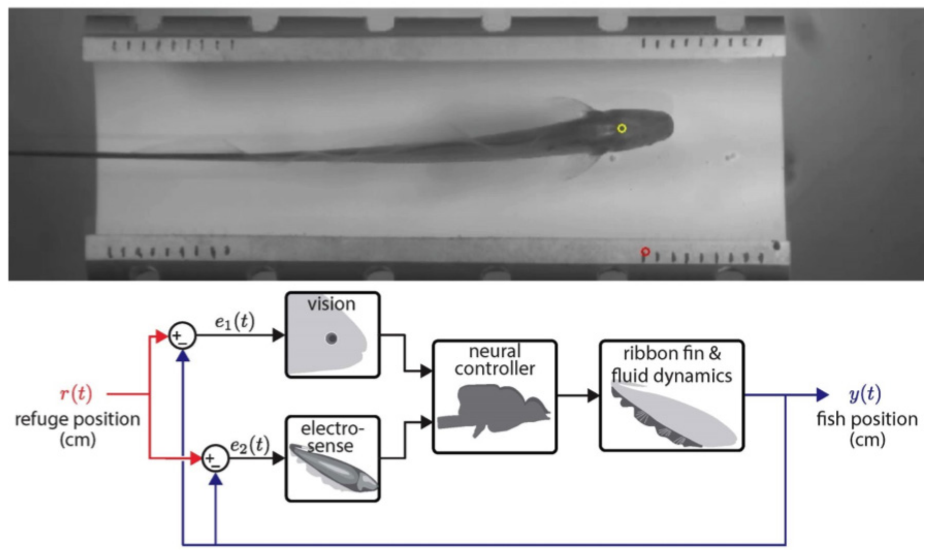

Figure 5 – Probing an animal's sheltering behavioral system by breaking into the sensorimotor feedback loop on the lab bench. [5]

Reason

Reason is the study of how an agent's sensing and computational properties enable it to represent and predict the world around it. This is an aspect of study traditionally emphasized by researchers in AI and learning *(Figure 6)*.

Intelligent reason requires theories of how to combine physical embodiment and abstract knowledge. Here, the use of the ubiquitous yet elusive term "intelligent" draws attention to the particular difficulties of defining, measuring and achieving robustly sustainable embodied complex autonomous operation in real time and space. The participants contrast this principle with traditional robotics research which has characteristically focused on physical embodiment at the expense of higher-order knowledge, substantially limiting the interface of autonomy to more complex tasks *(Figure 7)*. **New fundamental theories are required that can model and explain how to reason about the complex combination of physical embodiment and "common sense" knowledge and behavior.**

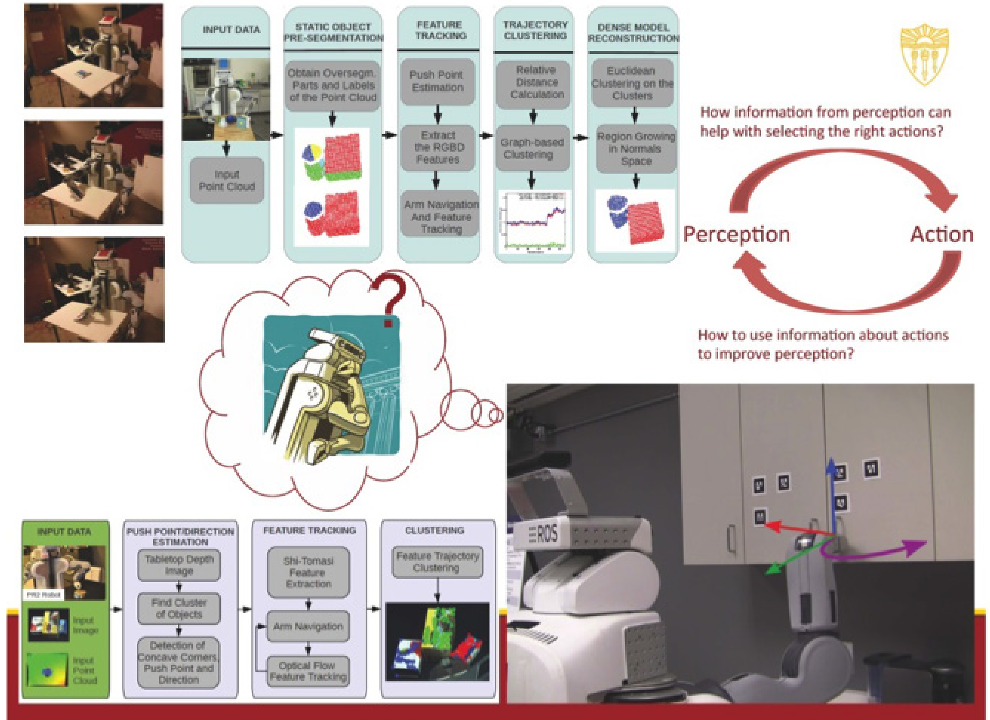

Figure 6 – AI systems for mobile manipulation in unstructured environments are built from layers of feature extraction, learning, model building and planning modules. [6]

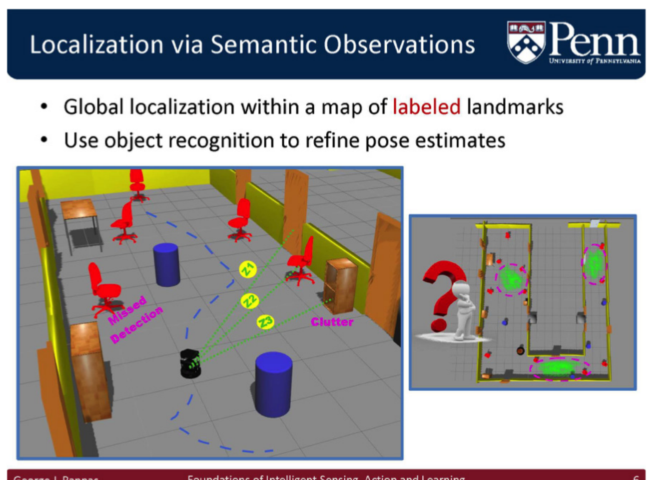

Figure 7 Recent work combines physically embodied higher-order "semantic" knowledge with theoretical guarantees and computationally tractable algorithms suitable for real time implementation. [7]

Interaction

Interaction is the study of how an agent exists in a bigger world, following the goals and priorities of human partners, a particular focus of the human-robot interaction (HRI) community *(Figure 8)*.

Intelligent interaction requires theories of how an intelligent system must behave as part of a larger, populated world. The participants use the terms "interactive" and "social" to underscore the overwhelming consensus that it is exactly the exchange between such agents and their environments—including other agents—that offers the greatest scope for new science and technological advance, but, at the same time, provokes the most challenging intellectual problems. Too many so-called intelligent physical systems have ignored the "human in the loop" in exchange for simplicity of design but leading to operational failure. **New fundamental theories are required to explicitly capture complex human behaviors, expectations and communication in the design of intelligent physical systems.**

Figure 8 *Left*– **Advances in Human-Robot Interaction entail focused interdisciplinary collaboration of roboticists with multiple disciplines within engineering and computing, but also tie directly into animal cognition, social psychology, cognitive science, and developmental psychology, among other social sciences.** [8 L] *Right* – **Example of robots with the ability to request help from humans and other remote resources.** [8 R]

Successes, Opportunities, and Challenges of Physical Intelligence Research

Current research in Intelligent Physical Systems falls into three primary areas: embodied intelligence, AI and learning, and social intelligence. This section summarizes the participants' view of the current state of these areas, as well as, their opportunities and challenges.

Embodied Intelligence

A subgroup of workshop participants focused on the role of the embodying mechanical system as a significant source of intelligence and hence autonomy in animals, as well as machines. The participants strongly argued that it is essential to recognize that embodied intelligence is more than an "implementation detail" to be abstracted away within higher level reasoning. By contrast, linking intelligence to action in the physical world through an embodiment—complete with sensors, actuators, and the capacity for motivation, learning and recall—is crucial for enabling complex, highly dynamic and creative action. This view contrasts with the main tradition within AI that continues to couch intelligence in largely computational terms: as search (e.g. modern chess engines) and function approximation (e.g. convolutional neural networks). It is inspired in part by insights from integrative and evolutionary biology offering the century-long investigation of organisms as the indivisible unit of biological behavior. Such inquiry explores an animal's agency as arising only through the embodiment of a whole, intact body in the physical world. Many roboticists who find this view essential to progress in our field have pursued deep, and still developing, collaborations with integrative biologists to the substantial benefit of both disciplines.

For example, over the last 20 years, a field of "neuromechanics" is emerging from this intertwining of integrative biology with the rise of robotics. Again, this development may be contrasted in analogy to the way in which the AI tradition of computation as intelligence deeply influenced systems neuroscience, which until recently treated the *brain* as a *biological computer*, and *intelligence* solely in terms of functional neural circuits. **The neuromechanics view recognizes that neural computation is inherently linked to work in the physical world, and that sensory feedback and neural computation can only be understood in their closed-loop (embodied) context.** This new synergy between biology and robotics is revolutionizing both fields, leading to a new, deeper understanding of animal intelligence while creating machines with steadily expanding physical competences that promise a path toward autonomy. The ability to translate ideas between neuromechanics and robotics has benefited from the mathematics of dynamical systems and control theory. These formalisms provide a means to describe synthetic systems, behaviors, and computations using a hierarchy of models in a systematic manner for both analysis and synthesis.

A next generation of scientists and engineers has turned this approach back toward biology, using that systems theory to explore neural systems with greater and greater fidelity. Their approach to "closed-loop neuroscience" is generating new ways of augmenting an animal's experience in real-time in order to tease out the target of an animal's motor outputs, while simultaneously decoding the neural activity that underlies them. Such closed-loop alternations of the apparent "physics" that translates neural activity at one level (e.g. motor output) to subsequent neural signals (e.g. sensory inputs) promises a critical capability that, over the next 5–10 years, will enable us to understand how an animal directs its motor commands so that it can obtain the sensory information it needs to form a cognitive, symbolic representation of its environment, suitable for decision making.

Ultimately, this leads to a conception of animal agency as an embodied system navigating the physical world: the ability to spatiotemporally "stamp" its sensory signals with meaningful context and then act on its world to influence its next sensations in the desired manner. Recently, the Nobel Prize in Physiology or Medicine was awarded based on the neural correlates of this "cognitive map" theory. Whereas this notion—that sensory signals modulate an internal representation of place—should, in principle, work equally well for a merely passive observers, the opportunity in the next 10–20 years lies in understanding how the neuromechanics of autonomous mobility and manipulation drive the active deployment of the sensorium to facilitate a

> "Linking intelligence to action in the physical world through an embodiment—complete with sensors, actuators, and the capacity for motivation, learning and recall—is crucial for enabling complex, highly dynamic and creative motion."

contextualized set of actionable memories embedded in a cognitive scaffold of space and time. **In turn, reflecting those insights back onto the sensorimotor design of robots gives the promise of machines that begin to exhibit agency heretofore only found in animals.**

AI and Learning

Many of the recent successes of intelligent physical systems such as self-driving cars and service robots have been driven by fundamental breakthroughs in statistical inference, machine learning and computer vision. However, the state of the art in these technologies has not yet delivered a sufficiently high level of capability. Even the most advanced intelligent systems do not yet have the ability to carry out complex, long-duration missions without human oversight, and cannot adapt to substantial changes in the world or deal with more than minimally off-nominal situations.

Intelligent physical systems are fundamentally limited by the inability to integrate their reasoning about the physical world with higher level abstract knowledge. The vast majority of robotic systems reason in terms of the basic geometry of the world around them and very simplistic models of world dynamics rather than in terms of symbolic representations that represent discrete entities with function and complex dynamics. While AI has a long tradition of reasoning about the world with higher level, symbolic representations, **the open challenge has been how to bridge the "semantic gap" between the high-level symbolic representations and the lower-level sensor signals and motor actions of a robot.**

Recent breakthroughs in these fields such as "*deep learning*", object recognition and symbol grounding have the promise of closing this gap. Decades of work in AI and computational perception are advancing machine ability to infer abstract relationships starting from raw sensor inputs such as pixel intensity values *(Figure 9)*. Recent progress in learned models of object segmentation and recognition has been shown to allow much more general object manipulation in cluttered, natural environments than was previously possible. Similarly, recent progress in learned models of symbol grounding has enabled natural language interaction that allows a robot and a human partner to collaborate more naturally than was previously possible. However, these results only indicate the potential for advances in the level of autonomy: major questions must be answered and theories must be developed in order to achieve robots that can reliably carry out complex missions autonomously.

A science of combined physical and abstract reasoning will advance a number of essential capabilities in intelligent systems. For a robotic system that is executing a long-duration, complex mission autonomously by reasoning about the world and learning over time, there is a critical need to be able to specify and ensure correct behavior.

"The ability to bring "common sense" knowledge to bear on representations of the physical world will lead to dramatically more data-efficient learning and especially knowledge and skill transfer and generalization across domains."

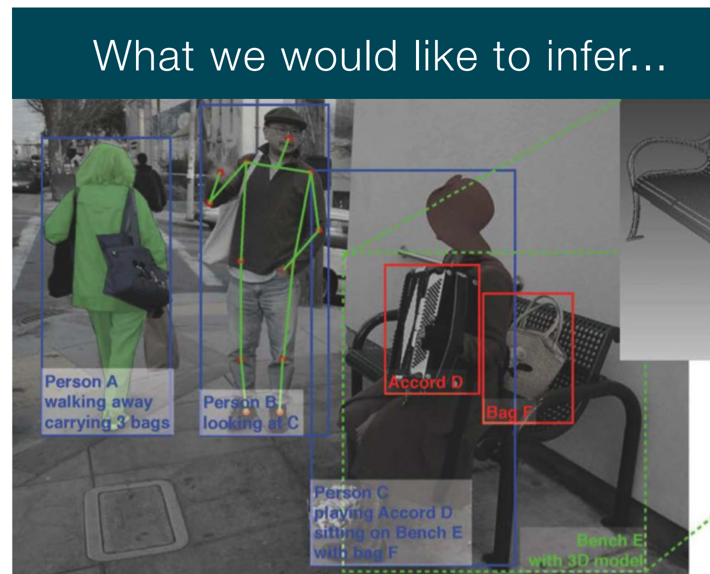

Figure 9 – Will person B put some money into Person C's tip bag? [9]

> "the open challenge has been how to bridge the "semantic gap" between the high-level symbolic representations and the lower-level sensor signals and motor actions of a robot."

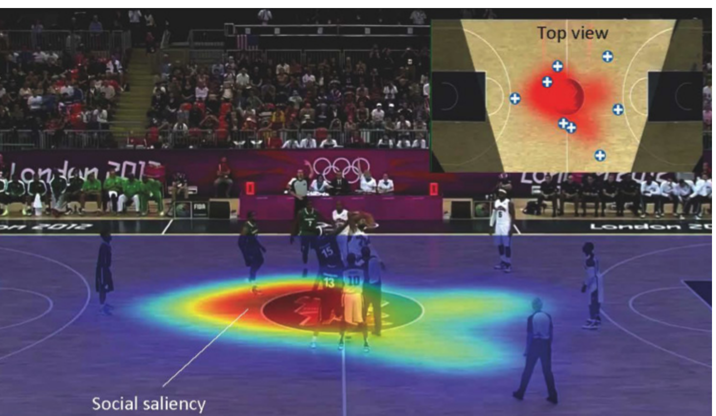

Figure 10 – General artificial mental models may potentially emerge from early developments like this system that learns to predict social saliency from social interaction data captured by first person cameras. *[10]*

Existing techniques in verification and safety assurance cannot be applied to systems that incorporate modern perception or to systems that learn and adapt over time. Major advances in system modelling are required to even represent the performance of the system, as well as advances in algorithms for analysis to ensure safe and reliable performance. Second, while statistical inference and learning are becoming critical capabilities for enabling high levels of autonomy, the state of the art algorithms are usually data-intensive. Data is often cheap for purely in silicon applications such as web-based image recognition or semantic content extraction, but is considerably more expensive for a physically embodied agent. **The ability to bring "common sense" knowledge to bear on representations of the physical world will lead to dramatically more data-efficient learning and especially knowledge and skill transfer and generalization across domains.** In a robotic context, it is not even necessarily clear what constitutes the boundaries of a "domain" or how to define domains with respect to autonomous knowledge and skill transfer. Third, even with the best verification techniques and data-efficient learning, robots have limited ability to model the behavior of the people around them. Inferring hidden states of the world from partial sensor data has been a major thrust of perception and learning in robotics for many years, but there has been limited success at inferring human intentional states, as we now consider in detail.

Social Intelligence

There is a prevailing need for robots in human environments. For an autonomous system to exist in a populated, social context, major advances are required in modelling and inferring human behavior and intentional goals, both of the people themselves but also of what is expected of robot behavior—including the possibility that they must infer and follow novel social norms (culture and environment dependent), rather than necessarily merely model those endowed by their human users. **Robotics has finally reached the state of readiness for the beginnings of such a push toward interacting with people in human environments.**

Two major drivers have caused the current surge in human-robot interaction (HRI) research and development. The first driver is technological, and includes the recent leap in perception through affordable 3D vision for human activity tracking, as well as the development of ever smaller, safer, and, increasingly, softer robot bodies. For example, recent work in computer vision develops computational *mental models* by learning to predict social saliency, the likelihood of joint attention to an input image or video using the social interaction data captured by first person cameras *(Figure 10)*.

The second driver is socio-economic, resulting from societal factors (aging population, tech-savvy youth, and safety and health challenges), creating economic opportunities that are producing a surge in industry investment in robotics development (currently focused on autonomous driving and drones, but expanding into manufacturing and home automation).

HRI contexts vary drastically, from structured ones, such as factories, roadways, airports, and hospitals, to less structured ones, such as streets, public areas, office environments, and retirement homes, to the ultimate unstructured environments: homes. In all cases, HRI involves a combination of real-time perception (of the environment and humans), understanding of not only the current state and ongoing activity, but also intentions of the human participants, and autonomous (or semi-autonomous) response that is safe, timely, natural, ethical, engaging, collaborative, and effective relative to the goals of the interaction context. HRI also includes one-on-one, one-to-many, and many-to-many human-robot interactions, which span a variety of relevant background literature (animal behavior, distributed coordination, economic models, etc.) and models for communication and coordination.

The yearly Robocup soccer, an adversarial game between teams of autonomous robots, represents one of the earliest (and still enduring) empirical settings

> “Robotics has finally reached the state of readiness for the beginnings of such a push toward interacting with people in human environments.”

for studying and synthesizing unstructured exchanges between physically embodied agents *(Figure 11)*.

Current research into HRI is divided into two separate, non-interacting subfields and associated communities: physical/contact-based HRI and social/non-contact based HRI. Physical HRI encompasses medical robotics, haptics, manufacturing, some service robotics, rehabilitation, and assistive robotics. Non-physical HRI includes socially assistive robotics, educational robotics, social robotics, some service robotics, and entertainment robotics. In spite of this dichotomy of research areas, most real-world problems involve both physical and non-physical aspects of the HRI problem. The two communities therefore need to be brought together to address real-world challenges that make obvious the need for their synergy. More broadly, progress in HRI requires much closer collaboration between robotics, machine vision, machine learning (ML), and AI. Currently, those areas are also largely separate, not yet tackling the same challenges.

Figure 11 – Robocup soccer explores unstructured exchanges between physically embodied agents. [11]

HRI is inherently interdisciplinary at multiple levels. Not only does it require cross-disciplinary work within engineering, as discussed above, but it also ties directly into animal cognition, social psychology, cognitive science, and developmental psychology, among other social sciences. Failing to connect to those disciplines results in naive work or “reinventing the wheel”, but bringing the disciplines together requires sufficient investment, typically lacking in most funding programs. Finally, a major challenge for HRI progress is the need for accessible data sets and evaluation scenarios to ground the work in real world contexts of interest. Because of privacy concerns surrounding the use of human data, and the complexity of deploying robots in real-world human environments, currently very few HRI research projects actually use realistic multi-modal interaction data (featuring audio, video, possibly physiologic data, background data, etc.) and are tested in real-world environments outside of the lab or highly controlled warehouse. The final challenge to progress is the current lack of affordable platforms available for research use in the US. However, with the surge in robotics interest by industry/startups, this landscape may change in the next decade or so.

Research in HRI advanced drastically after the introduction of affordable 3D vision (Kinect, PrimeSense) and the associated models of human activity, facilitating recognition and tracking needed for HRI. As outlined above, similar drastic leaps in capability could be achieved by removing some of the key barriers, including training data sets, evaluation testbeds and environments, and synergies with machine vision and machine learning in particular. **For contact-based HRI to advance, major leaps in manipulation as well as soft robotics will be necessary to develop systems that are safe and capable in everyday human environments.**

Future Topics in Physical Intelligence Research

Workshop participants discussed specific directions of Physical Intelligence Research over the next few decades. They categorized these examples of productive future research programs into three areas—Architecture, Environmental Interaction and Tasking or Human Social Interaction—and also gave examples of tasks and assays suitable for benchmarking progress in physical intelligent systems. **The participants were generally optimistic that pursuing research programs along the lines of these examples has the strong potential to yield critical antecedent discoveries and insights over the next 10 years that will result in fieldable intelligent, interactive, social, sensorimotor systems within two decades.**

Architecture

Situational and Self Awareness: Research projects of this sort aim to create aims to create robots that are able to predict their own chance of success or failure, can synthesize an explanation of their own behavior, are capable of extrapolation and generalization, able to transfer a learned task across sensor modalities, and can synthesize textual description of a scene.

Management of Multiple Tasks Work in this area will produce architectures capable of supporting a computationally effective notion of a task domain as a set of multiple, distinct tasks that fit together in some usefully composable manner. It will yield systems with flexible goal structures, interleaving control networks that allow accomplishment of different tasks. It will focus on effective task representations, or task domains, and parameterize the space of controllers or behavior-designs (so robot can explore).

Graceful Handling of Uncertainty Research in this area will produce designed systems with "common sense". It will find general approaches to operate with hugely imperfect knowledge as distinct from "savant" approaches to handling highly structured representations of uncertainty in highly specific problem settings. It will find scalable, approximately optimal ways of dealing with uncertain dynamical systems.

Representation This research will develop an architecture that understands what a word means. Research in this area will find the sweet spot between design and learning—a spectrum ranging from black box statistical methods for converting raw data into actuated motions (skills) through methods for the learning or construction of primitives whose compositions into tasks can be learned, or, at the other end of the spectrum, architectures relying upon constructed primitives whose compositions into tasks are explicitly programmed. It advances our insight into the tradeoff between model-driven and data-driven methods.

Environmental Interaction

Empirical Benchmarking Research in this area will develop principled validation/verification and standards for physical systems that operate at large (in unstructured and instrumented natural and synthetic environments). It will develop task benchmarks (e.g. robot planning) that can be applied and compared/contrasted across different environments. It will find environments and tasks relative to which (all, none, some) system components can be provably validated by simulation without need for physical experiments.

Environmental Complexity Research projects in this area aim to determine what class(es) of models offer the most appropriate mix of mathematical tractability and physical accuracy in representation of the sensorimotor experience of an embodied agent. Such work seeks to determine whether there are intrinsic measures (and if so what are they) or if it is task or agent specific. It aspires to give conditions for failure or impossibility of a task in an environment and characterize the conditions under which a class of architectures (e.g., vision-based algorithm) will achieve some class of tasks according to some quantified measure of success or error. For example, it might give environmental conditions sufficient for a separation theorem to hold for some class of perception and action abstractions.

Agent Capacity Work of this nature seeks to define autonomy as a relation of tasks to environments with a measured degree of autonomy as a function of the complexity of the system, environment and tasks. It aspires to characterize animal autonomy and give a formal characterization of and conditions for exaptation. For example, it might aim to provide a physical reachability proof for a legged machine in a specified environment, to formalize the notion of "affordance" (e.g., a "placeability" notion encompassing all places a robot can put an object) and to give a proof that a given manipulation system can access all (or some characterized subset of) the affordances presented by a given set of environments.

Tasking (Human Social Interaction)

Theory and Practice of Intent Research in this area will characterize the meaning and psychology of human intent, develop mathematically and computationally effective models and use them to define a corresponding notion for physical systems. It will develop a general theory which maps human intention to robot intention and develops machine inference of human intent from observation. It formalizes the relationship between the amount of human interaction required and the capability of the robot to infer intent.

Theory and Practice of Trust Work in this area is needed to characterize the meaning and psychology of human trust, developing mathematically and computationally effective models that will define properties of physical systems in terms of the systems' capacity to act and interact according to human intention. It develops a scale of trust and empirical measures to characterize it.

Theory and Practice of Coordination This research aims to characterize the meaning and psychology of human strategy, to develop mathematically and computationally effective models and use them to extend notions of coordination beyond physical configuration space into spaces of strategy and intent. This work will define and develop empirical metrics for these deeper versions of coordination, especially for teams of heterogeneous systems and develop theory that enables the proof of coordination for a given level of shared intent.

Physical Intelligence Assays and Benchmarking

The participants agreed upon a sample of benchmarking tasks that could measure progress toward the emerging new systems' achievement of physical intelligence.

Sample Thresholds of Agent Capacity Threshold capabilities examples include grasping any object that fits into a cubic foot volume, hypothesizing the mechanical properties or affordances of any such object, playing soccer as team-mate or even as coach, performing autonomously in two different task-domains, completing an automated transfer of task domain mastery—e.g., use web to recover data instructions from web to acquire new skill, and building a localization module that is accurate to some spec, such as, anywhere within one degree/millimeter/second.

Sample Assays of Interaction These might include robotic attendants that improve quality of life measures, are personalized to fit or improve an individual (or team) performance (e.g., over multi-week training), can improve a child's learning performance over stated developmental period, and determine by interaction if dynamic entities in the environment are cooperative human, non-cooperative human, reactive, non-reactive, or stationary.

Conclusion

The emerging field of intelligent physical systems is poised to deliver artifacts of unprecedented capability and human benefit, as well as contribute fundamental insights to a broad range of disciplines spanning the life sciences, engineering and mathematics. Workshop participants are convinced that near-term focus on fundamental problems of architecture, environment and interaction can lead to novel systems and breakthrough theories within a decade that could enable fielded artifacts of the sort imagined in the introduction within the next twenty years. **Machines that are physically adept, socially aware, and intrinsically motivated to achieve their users' intent will revolutionize quotidian human toil and dramatically amplify human capacity for emergency response.**

There was a broad consensus among the meeting participants that the present juncture represents a key moment wherein programs requiring substantive encounters and resources focused on forging long term alliances between these constituent fields have the potential to dramatically refocus and help cohere their presently disparate trajectories. Alongside the brief account of the three principal communities involved in research on physical intelligence, this report touches on how the traditional tension between proponents of "embodied" and "reasoned" intelligence might be brought into stronger correspondence by the imperative of "social" interaction. Similarly, the contrast between the formalist and empiricist views of such a synthetic science suggests that programs promoting productive collaborations between their exponents can move the field along farther and faster than can either group in isolation.

The meeting presented a diversity of traditions and approaches to physical intelligence, exhibiting multiple fields rich in strong researchers with exciting records of achievement that augur similarly productive futures. However, closing the gaps between present capabilities and the introductory vignette of a broadly useful common-sense synthetic physical intelligence within the next two decades will require far stronger and coherent alliance between them.

A theoretical paradigm for physical agency is needed to unify the exciting advances in machine learning and perception with the burgeoning understanding of animal cognition and human psychology alongside the slowly emerging formalism and practice of programmable work. The central barriers to understanding and design require sustained investments to achieve a synthetic science of physical intelligence with all its technological promise.

The pursuit of a science of Physical Intelligence is of critical interest, as established repeatedly, and by multifarious observations during the workshop described in this report. A broad consensus of participants would agree that two primary motivators underly this view. The first is that such a science holds the promise of creating a future generation of machines that will display unprecedented abilities in perception, action and lifelong learning. The second is the potential for new understanding of the relation between physical science and engineering and that of biology and human psychology. These advances in fundamental understanding will increase our ability to extract performance from such machines and, at least as importantly, to develop rational explanations for when and why further improvement is not possible.

It is natural to ask "why now?" Why is the timing right? With the attention that industry is giving to consumer-oriented products that rely on some aspects of autonomy, the scientific community can leverage hardware and other commercial assets to experiment and accelerate scientific understanding and elucidate underlying principles—something industry will not pursue. Of course, this flowering of commercial attention bespeaks breakthroughs in statistical inference, machine learning and computer vision whose insights are not a substitute for, but rather offer new help toward developing the foundations of intelligent physical systems. At the same time, the commercial technology is pushing out dramatically improved, cheaper and new tools that promise to revolutionize the empirical side of this science including the recent leap in perception through affordable 3D vision for human activity tracking, as well as the development of ever smaller, safer, and, increasingly, softer robot bodies. Finally, a strong socio-economic incentive to develop autonomous physical agents is emerging in consequence of widely reported societal factors such as an aging population, tech-savvy youth, and safety and health challenges.

The workshop participants were optimistic that fundamental research into intelligent physical systems of the kind described in this report will transform our understanding of intelligence and lead to critically beneficial applications within the next two decades.

Glossary

Agency (or Agent) – a computational architecture possessing some perceptual apparatus and ability to act on its environment for the purpose of achieving an autonomous goal.

Animal Cognition – the science of the internal capacity for representation and reasoning across a wide diversity of species, including invertebrates.

Cybernetics – the mathematical study of behavioral governance mechanisms in animal and artificial systems.

Deep Learning – Computational properties of abstract neuronal models that have influenced the origins of, successively, computer science, AI, and machine learning over the past fifty years.

Form and Function – interaction of morphology with environment to achieve function—the same anatomical components and even the same physical elements typically contribute multiple, distinct simultaneous behavioral benefits.

Haptics – this term applies generally to the science and engineering of physical devices that bring a virtual sense of “feel” (touch, heft, response, and so on) for humans interacting with a remote environment.

Mental Model – internal representations of an agent’s environment ranging from relatively narrow instrumental constructions to the farthest reaching foundations of intelligence.

Proprioception – That component of an agent’s sensory endowment located within the body but excited by its own contact with or movement within the environment.

Situational Awareness – the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future.

Semantic Gap – the conflicts between meanings expressible in one as against another (formal or informal) symbol system.

Appendix I
Intelligent Physical Systems

Nora Ayanian – *http://www-bcf.usc.edu/~ayanian/*
University of Southern California, *ayanian@usc.edu*
Department of Computer Science
PhD (2011), Mechanical Engineering, University of Pennsylvania

Nora Ayanian leads the Automatic Coordination of Teams (ACT) Laboratory and is also a member of the Robotics and Autonomous Systems (RASC) Center and Viterbi Arid Climates and Water Research Center.

Her research focuses on creating end-to-end solutions for coordinating teams of robots that start from truly high-level specifications and deliver code for individual robots in the team, such as using simple multi-touch inputs to control a team of UAVs. Ayanian brings a unique approach to multi-robot systems, creating unified solutions that address task assignment, path planning, and control that are broadly applicable across all aspects of multi-robot systems and mobile sensor networks. Her solutions provide guarantees of convergence and safety on real robotic systems. Ayanian is also a co-founder and current co-chair of the IEEE Robotics and Automation Society Technical Committee on Multi-Robot Systems.

She is a recipient of an NSF CAREER award, Best Paper in the Robotics Track at International Conference on Automated Planning and Scheduling (ICAPS), the Hanna Reisler Mentorship Award for mentorship of female undergraduate students in research, Best Student Paper at the International Conference on Robotics and Automation (ICRA), and a National Science Foundation Graduate Research Fellowship. Ayanian was also named to the inaugural "Mic50" by Mic.com (2015), "40 under 40 Professors who Inspire" by NerdWallet (2014), "AI's 10 to Watch" IEEE AI Magazine (2013).

Yuliy Baryshnikov – *https://publish.illinois.edu/ymb/*
University of Illinois, Urbana-Champaign, *ymb@uiuc.edu*
Departments of Mathematics and Electrical and Computer Engineering
PhD (1987), Applied Mathematics, Institute of Control Sciences, Moscow

Yuliy Baryshnikov is Professor of Mathematics and Electrical and Computer Engineering at University of Illinois at Urbana-Champaign. He earned his PhD in applied mathematics, from Institute of Control Sciences in Moscow. He spent his Humboldt Research Fellowship University of Osnabruck in Germany, and then worked as a faculty member in the Netherlands, UK and France, before joining Bell Labs in Murray Hill, NJ in 2001. In 2011 he resigned from his position as a department head there and moved West, to become professor of mathematics and electrical and computer engineering at the University of Illinois at Urbana-Champaign.

His research interests include probability theory, singularities, dynamical systems, and combinatorics. Among applied areas his favorites are sensor networks, nonlinear control, mathematical economics, self-assembly.

Noah Cowan – *http://limbs.lcsr.jhu.edu/people/cowan/*
Johns Hopkins University, *ncowan@jhu.edu*
Department of Mechanical Engineering
PhD (2001), Electrical Engineering and Computer Science, University of Michigan

Noah Cowan is an Associate Professor of Mechanical Engineering at Johns Hopkins University . He leads the Locomotion in Mechanical and Biological Systems (LIMBS) Lab and is co-director of the Laboratory for Computational Sensing and Robotics (LCSR). The LIMBS Lab conducts original experiments and computational analyses on both biological and robotic systems, with a focus on applying concepts from dynamical systems and control theory to garner new insights into the principles that underly neural computation. Dr. Cowan's research program has been recognized by a Presidential Early Career Award in Science and Engineering (PECASE) and a James S. McDonnell Complex Systems Scholar award. Dr. Cowan also received the William H. Huggins Excellence in Teaching award and the Dunn Family Award for Excellence in Mentoring.

Robert Full – *http://polypedal.berkeley.edu/*
University of California, *Berkeley, rjfull@berkeley.edu*
Department of Integrative Biology, Poly-Pedal Lab
PhD (1984), Biology, State University of New York – Buffalo

Robert J. Full is a Chancellor's and Goldman Professor in the Department of Integrative Biology at the University of California at Berkeley. He is founder and director of CiBER, the Center for interdisciplinary Bio-inspiration in Education and Research and Director of the Poly-PEDAL Laboratory.

His research program in comparative physiology and biomechanics has shown how examining a diversity of animals can lead to the discovery of general principles. His fundamental discoveries in animal locomotion have inspired the design of novel neural control circuits, artificial muscles, autonomous legged search-and-rescue robots and the first, synthetic self-cleaning dry adhesive inspired by his discovery on how geckos stick.

Professor Full is the Editor-in-Chief of the journal Bioinspiration and Biomimetics, and the Principal Investigator on an NSF Integrative Graduate Education and Research Traineeship on Bio and Bio-inspired Motion Systems Operating in Complex Environments that is training the next generation of biologists and engineers to collaborate in mutually beneficial relationships. Professor Full received a Presidential Young Investigator Award, was named a Mentor in the Life Sciences by the National Academy of Sciences and is a Fellow of the American Association for the Advancement of Science.

Robert Ghrist – *https://www.math.upenn.edu/~ghrist/*
University of Pennsylvania, *ghrist@math.upenn.edu*
Departments of Mathematics & Electrical/Systems Engineering
PhD (1995), Applied Mathematics, Cornell University

Robert Ghrist is the Andrea Mitchell PIK Professor of Mathematics and Electrical & Systems Engineering at the University of Pennsylvania. He is a recognized leader in the field of Applied Algebraic Topology, with publications detailing topological methods for sensor networks, robotics, signal processing, tracking, network discovery, and more. He is the author of a leading textbook on the subject (Elementary Applied Topology, 2014).

His prior work in leading the DARPA DSO SToMP project and participating in several DoD MURIs is complemented by NSF CAREER, NSF PECASE, and SciAm50 awards. He is also a dedicated expositor and communicator of Mathematics, with a popular MOOC on Calculus at Coursera.

Lucia Jacobs – *http://jacobs.berkeley.edu*
University of California, Berkeley, *jacobs@berkeley.edu*
Department of Psychology and Helen Wills Neuroscience Institute
PhD (1987), Ecology and Evolution, Princeton University

Lucia F. Jacobs leads the Jacobs Lab of Cognitive Biology at Berkeley where she is a Professor in Department of Psychology and Institute of Neuroscience, as well as a member of the Institute for Cognitive and Brain Science.

The focus of her research is the ecology and evolution of navigating choices: how animals make choices about what and where to eat, how to navigate and map new terrains and how generally to integrate diverse sources of information to make adaptive decisions in uncertain environments. Animal species include humans, search dogs and rodents (domestic and wild). Her theoretical work on navigation focuses on the evolution of limbic structures (hippocampus, olfactory systems) and their integrated role in spatial navigation.

She is a recipient of a NSF CAREER award, a Hellman Junior Faculty Award, a Prytanean Faculty Award and a Mary Rennie Epilepsy Award. She has presented her work in a Herbert Spencer Lecture at the University of Oxford, a Santa Fe Public Lecture and the 2013 Michigan State Distinguished Lecturer in Cognitive Science.

Dan Koditschek – *http://kodlab.seas.upenn.edu/Kod/Projects*
University of Pennsylvania, *kod@seas.upenn.edu*
Department of Electrical and Systems Engineering
PhD (1983), Electrical Engineering, Yale University

Daniel E. Koditschek is the Alfred Fitler Moore Professor of Electrical and Systems Engineering, within the University of Pennsylvania School of Engineering and Applied Science. Koditschek received his bachelor's degree in Engineering and Applied Science and his M.S. and Ph.D. degrees in Electrical Engineering in 1981 and 1983, all from Yale University. He served on the Yale Faculty in Electrical Engineering until moving to the University of Michigan a decade later. In 2005, he moved to Penn as Chair of the recently formed Electrical and Systems Engineering Department, and remained in that position through 2012.

Koditschek's research interests include robotics and, more generally, the application of dynamical systems theory to intelligent mechanisms. His more than 200 archival journal and refereed conference publications have appeared in a broad spectrum of venues ranging from the Transactions of the American Mathematical Society through The Journal of Experimental Biology, with a concentration in several of the IEEE journals and related transactions. Various aspects of this work have received mention in general scientific publications such as Scientific American and Science as well as in the popular and general lay press such as The New York Times and Discover Magazine.

Dr. Koditschek is a member of the AMS, ACM, MAA, SIAM, SICB and Sigma Xi and is a Fellow of the IEEE and the AAAS. He holds secondary appointments within the School of Engineering and Applied Science in the departments of Computer and Information Science and Mechanical Engineering.

Hadas Kress-Gazit – *http://verifiablerobotics.com/people/pi*
Cornell University, *hadaskg@cornell.edu*
Department of Mechanical and Aerospace Engineering
PhD (2008), Electrical Engineering, University of Pennsylvania

Hadas Kress-Gazit is an Associate Professor at the Sibley School of Mechanical and Aerospace Engineering at Cornell University and the director of the Verifiable Robotics Research Group.

Her research focuses on formal methods for robotics and automation and more specifically on creating verifiable robot controllers for complex high-level tasks using logic, verification, synthesis, hybrid systems theory and computational linguistics. She received an NSF CAREER award in 2010, a DARPA Young Faculty Award in 2012 and the Fiona Ip Li '78 and Donald Li '75 Excellence in teaching award in 2013.

Vijay Kumar – *http://www.kumarrobotics.org/*
University of Pennsylvania, *vijay.kumar@seas.upenn.edu*
Departments of Mechanical Engineering and Applied Mechanics, Computer and Information Science, and Electrical and Systems Engineering
PhD (1987), Mechanical Engineering, Ohio State University

Vijay Kumar is the Nemirovsky Family Dean of Penn Engineering with appointments in the Departments of Mechanical Engineering and Applied Mechanics, Computer and Information Science, and Electrical and Systems Engineering at the University of Pennsylvania. Kumar's group works on creating autonomous ground and aerial robots, designing bio-inspired algorithms for collective behaviors, and on robot swarms. They have won many best paper awards at conferences, and group alumni are leaders in teaching, research, business and entrepreneurship. Kumar is a fellow of ASME and IEEE and a member of the National Academy of Engineering.

Vijay Kumar has held many administrative positions in the School of Engineering and Applied Science, including director of the GRASP Laboratory, chair of Mechanical Engineering and Applied Mechanics, and the position of the Deputy Dean. He served as the assistant director of robotics and cyber physical systems at the White House Office of Science and Technology Policy.

Steven LaValle – *http://msl.cs.uiuc.edu/~lavalle/*
University of Illinois, Urbana-Champaign, *lavalle@uiuc.edu*
Department of Computer Science
PhD (1995), Electrical Engineering, University of Illinois

Steven M. LaValle is Professor of Computer Science in the Department of Computer Science at the University of Illinois. He received his Ph.D. in Electrical Engineering from the University of Illinois in 1995. From 1995–1997 he was a postdoctoral researcher and lecturer in the Department of Computer Science at Stanford University. From 1997–2001 he was an Assistant Professor in the Department of Computer Science at Iowa State University.

His research interests include robotics, virtual reality, sensing, planning algorithms, computational geometry, and control theory. He is mostly known for his introduction of the Rapidly exploring Random Tree (RRT) algorithm, which is widely used in robotics and other engineering fields. He was also an early founder and chief scientist of Oculus VR, acquired by Facebook in 2014, where he developed patented tracking technology for consumer virtual reality and led a team of perceptual psychologists to provide principled approaches to virtual reality system calibration, health and safety, and the design of comfortable user experiences. He also authored the books Planning Algorithms, and Sensing and Filtering, and is currently writing Virtual Reality.

Dan Lee – *https://www.grasp.upenn.edu/*
University of Pennsylvania, *ddlee@seas.upenn.edu*
Department of Electrical and Systems Engineering
PhD (1995), Condensed Matter Physics, Massachusetts Institute of Technology

Daniel Lee is the UPS Foundation Chair Professor in the School of Engineering and Applied Science at the University of Pennsylvania. He received his B.A. summa cum laude in Physics from Harvard University in 1990 and his Ph.D. in Condensed Matter Physics from the Massachusetts Institute of Technology in 1995. Before coming to Penn, he was a researcher at AT&T and Lucent Bell Laboratories in the Theoretical Physics and Biological Computation departments.

He is a Fellow of the IEEE and AAAI and has received the National Science Foundation CAREER award and the University of Pennsylvania Lindback award for distinguished teaching. He was also a fellow of the Hebrew University Institute of Advanced Studies in Jerusalem, an affiliate of the Korea Advanced Institute of Science and Technology, and organized the US-Japan National Academy of Engineering Frontiers of Engineering symposium. As director of the GRASP Laboratory and co-director of the CMU-Penn University Transportation Center, his group focuses on understanding general computational principles in biological systems, and on applying that knowledge to build autonomous systems.

Kevin Lynch – *http://www.mccormick.northwestern.edu/research-faculty/directory/profiles/lynch-kevin.html*
Northwestern University, *kmlynch@northwestern.edu*
Department of Mechanical Engineering
PhD (1996), Robotics, Carnegie Mellon University

Kevin Lynch is Professor and Chair of the Mechanical Engineering Department at Northwestern University. He earned a BSE in Electrical Engineering from Princeton University and a PhD in Robotics from Carnegie Mellon University. He is a member of the Neuroscience and Robotics Lab (nxr.northwestern.edu) and the Northwestern Institute on Complex Systems (nico.northwestern.edu).

Dr. Lynch's research focuses on dynamics, motion planning, and control for robot manipulation and locomotion, self-organizing multi-agent systems, and functional electrical stimulation for restoration of human function.

He is a Senior Editor of the IEEE Robotics and Automation Letters, co-author of "The Principles of Robot Motion" (MIT Press, 2005) and "Embedded Computing and Mechatronics" (Elsevier, 2015), an IEEE fellow, and the recipient of the IEEE Early Career Award in Robotics and Automation, Northwestern's Professorship of Teaching Excellence, and the Northwestern Teacher of the Year award in engineering.

Jitendra Malik – *http://www.cs.berkeley.edu/~malik/*
University of California, Berkeley, *malik@berkeley.edu*
Department of Electrical Engineering, Computer Science Division
PhD (1985), Computer Science, Stanford University

Jitendra Malik is the Arthur J. Chick Professor in the Computer Science Division of the Department of Electrical Engineering and Computer Sciences also serves on the faculty of the department of Bioengineering, and the Cognitive Science and Vision Science groups. He received his PhD in Computer Science from Stanford University in 1985. During 2002–2004 he served as the Chair of the Computer Science Division and during 2004–2006 as the Department Chair of EECS.

Professor Malik studies computer vision. He develops models and algorithms that given an image, infer properties of the objects, people, and places in the world that gave rise to the image. He is also interested in the computational modeling of human vision. With colleagues, he has helped develop concepts and techniques such as anisotropic diffusion for image de-noising, normalized cuts for clustering and segmentation, high dynamic range imaging, ecological statistics of perceptual grouping, and machine learning approaches to visual recognition.

He has published nearly 200 papers, and mentored more than 50 doctoral and postdoctoral students. Jitendra Malik received the Distinguished Researcher Award from IEEE PAMI-TC and the K.S. Fu Prize of the International Association of Pattern Recognition. He is a member of the National Academy of Sciences and the National Academy of Engineering, and a Fellow of the American Academy of Arts and Sciences.

Maja Mataric – *http://www-robotics.usc.edu/~maja/*
University of Southern California, *mataric@usc.edu*
Departments of Computer Science, Neuroscience and Pediatrics
PhD (1994), Computer Science and Artificial Intelligence,
Massachussetts Institute of Technology

Maja Mataric is Professor and Chan Soon-Shiong Chair in the Computer Science Department, Neuroscience Program, and the Department of Pediatrics at the University of Southern California. She is founding director of the USC Robotics and Autonomous Systems Center (RASC), co-director of the USC Robotics Research Lab and Vice Dean for Research in the USC Viterbi School of Engineering. She received her PhD in Computer Science and Artificial Intelligence from MIT in 1994, MS in Computer Science from MIT in 1990, and BS in Computer Science from the University of Kansas in 1987.

Her Interaction Lab's research into socially assistive robotics is aimed at endowing robots with the ability to help people through individual non-contact assistance in convalescence, rehabilitation, training, and education. Her research is currently developing robot-assisted therapies for children with autism spectrum disorders, stroke and traumatic brain injury survivors, and individuals with Alzheimer's Disease and other forms of dementia.

She is a Fellow of the American Association for the Advancement of Science (AAAS), Fellow of the IEEE and AAAI, and recipient of the Presidential Awards for Excellence in Science, Mathematics & Engineering Mentoring (PAESMEM), the Anita Borg Institute Women of Vision Award for Innovation, Okawa Foundation Award, NSF Career Award, the MIT TR35 Innovation Award, and the IEEE Robotics and Automation Society Early Career Award. She served as the elected president of the USC faculty and the Academic Senate. At USC she has been awarded the Viterbi School of Engineering Service Award and Junior Research Award, the Provost's Mentoring Award and Center for Interdisciplinary Research Fellowship, the Mellon Mentoring Award, the Academic Senate Distinguished Faculty Service Award, and a Remarkable Woman Award. She is featured in the science documentary movie "Me & Isaac Newton", in The New Yorker ("Robots that Care" by Jerome Groopman, 2009), Popular Science ("The New Face of Autism Therapy", 2010), the IEEE Spectrum ("Caregiver Robots", 2010), and is one of the LA Times Magazine 2010 Visionaries.

Prof. Mataric is the author of a popular introductory robotics textbook, "The Robotics Primer" (MIT Press 2007), an associate editor of three major journals and has published extensively. She serves or has recently served on a number of advisory boards, including the National Science Foundation Computing and Information Sciences and Engineering (CISE) Division Advisory Committee, and the Willow Garage and Evolution Robotics Scientific Advisory Boards.

Lisa Miracchi – *https://philosophy.sas.upenn.edu/bio/miracchi*
University of Pennsylvania, *miracchi@gmail.com*
Department of Philosophy
PhD (2014), Philosophy, Rutgers University

Lisa Miracchi is an Assistant Professor of Philosophy at the University of Pennsylvania. She received her Ph.D. in Philosophy and Certificate in Cognitive Science from Rutgers University, New Brunswick in 2014, and her A.B. in Philosophy from Harvard in 2009. In 2014–2015, she was a Bersoff Assistant Professor/ Faculty Fellow at NYU's Philosophy Department, and was associated with NYU's Center for Mind, Brain, and Consciousness.

Her research interests are mainly in foundational theoretical questions about the nature of the mind (including perception, emotion, thought, knowledge, agency, and intelligence) and how to best understand and provide scientific explanations of these phenomena.

Robin Murphy – *http://faculty.cs.tamu.edu/murphy/*
Texas A&M University, *murphy@cse.tamu.edu*
Department of Computer Science and Engineering
PhD (1992), Computer Science, Georgia Institute of Technology

Robin Murphy is the Raytheon Professor of Computer Science and Engineering and the Director of the TEES Center for Emergency Informatics and the Center for Robot-Assisted Search and Rescue at Texas A&M.

She is a founder of the field of disaster robotics, participating in over 20 deployments including 9/11 World Trade Center, Hurricane Katrina and Fukushima Daiichi as covered in her recent TED talk. She has over 150 scientific publications on artificial intelligence, human-robot interaction, and robotics including the award winning Disaster Robotics which catalogs ground, aerial, and marine robot use by responders at 34 events worldwide and synthesizes best practices and tactics. She is an IEEE Fellow, winner of the 2014 ACM Eugene L. Lawler Award for Humanitarian Contributions Within Computer Science and Informatics, an "Innovator in AI" and "Agent of Change" by TIME, an "Alpha Geek" by WIRED Magazine, an Economist "Drone Ranger," one of the "Most Influential Women in Technology" by Fast Company, one of the Top 25 Doers, Dreamers and Drivers for 2015 by Government Technology Magazine, and #14 on the list of the 30 Most Innovative Women Professors Alive Today. Dr. Murphy has served on the Defense Science Board, including co-chairing the 2012 study on the role of autonomy in DoD, the USAF Scientific Advisory Board, and the Board on Army Science and Technology.

George Pappas – *http://www.georgejpappas.org/*
University of Pennsylvania, *pappasg@seas.upenn.edu*
Department of Electrical and Systems Engineering
PhD (1998), Electrical Engineering and Computer Sciences, *University of California at Berkeley*

George Pappas is the Joseph Moore Professor and Chair of the Department of Electrical and Systems Engineering at the University of Pennsylvania. He also holds a secondary appointment in the Departments of Computer and Information Sciences, and Mechanical Engineering and Applied Mechanics. He is member of the GRASP Lab and the PRECISE Center. He has previously served as the Deputy Dean for Research in the School of Engineering and Applied Science.

His research focuses on control theory and in particular, hybrid systems, embedded systems, hierarchical and distributed control systems, with applications to unmanned aerial vehicles, distributed robotics, green buildings, and biomolecular networks. He is a Fellow of IEEE, and has received various awards such as the Antonio Ruberti Young Researcher Prize, the George S. Axelby Award, the O. Hugo Schuck Best Paper Award, the National Science Foundation PECASE, and the George H. Heilmeier Faculty Excellence Award.

Alejandro Ribeiro – *https://alliance.seas.upenn.edu/~aribeiro/wiki/*
University of Pennsylvania, *aribeiro@seas.upenn.edu*
Department of Electrical and Systems Engineering
PhD (2007), University of Minnesota

Alejandro Ribeiro is the Rosenbluth Associate Professor at the Department of Electrical and Systems Engineering. He received his B.Sc. degree in Electrical Engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, in 1998. From 1998 to 2003, he was a member of the technical staff at Bellsouth Montevideo. He recieved a M.Sc. and Ph.D. degree in Electrical Engineering from the Department of Electrical and Computer Engineering at the University of Minnesota, in 2005 and 2007.

His research interests are in the applications of statistical signal processing to the study of networks and networked phenomena. His focus is on structured representations of networked data structures, graph signal processing, network optimization, robot teams, and networked control. Dr. Ribeiro received the 2014 O. Hugo Schuck best paper award, the 2012 S. Reid Warren, Jr. Teaching Award, the NSF CAREER Award in 2010, and paper awards at the 2015 Asilomar Conference on Signals Systems and Computers, the 2013 American Control Conference, as well as, the 2006 and 2005 International Conferences on Acoustics, Speech and Signal Processing. Dr. Ribeiro is a Fulbright scholar and a Penn Fellow.

Nicholas Roy – *http://groups.csail.mit.edu/rrg/*
Massachussetts Institute of Technology, *nickroy@csail.mit.edu*
Departments of Aeronautics and Astronautics, Computer Science and Artificial Intelligence Laboratory
PhD (2003), Robotics Institute, Carnegie Mellon University

Nicholas Roy is an Associate Professor in the Department of Aeronautics & Astronautics at the Massachusetts Institute of Technology. He leads the Robust Robotics Group at MIT and is a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL). He received his Ph. D. in Robotics from Carnegie Mellon University in 2003.

His research interests include aerial robotics and mobile autonomy, planning under uncertainty, machine learning and human-computer interaction. His recent work includes developing machine learning techniques for aggressive, dynamic flight through unknown, unstructured environments. Additionally, he has led the development of methods for symbol grounding and natural language understanding for robotic systems.

He is a recipient of an NSF CAREER award and an RAS Early Career Award. He and his students have won best paper and best student paper prizes at Robotics Science and Systems (RSS), the International Conference on Robotics and Automation (ICRA), the International Conference on Micro Air Vehicles (ICMAV), the International Conference on Multimodal Interfaces (ICMI) and the International Conference on Automated Planning and Scheduling (ICAPS). He and his students have won the International Competition on Micro-Air Vehicles (IMAV) and the AUVSI Aerial Robotics Competition. He was the founder of Project Wing at Google [x].

Ruslan Salakhutdinov – *http://www.cs.toronto.edu/~rsalakhu/*
Carnegie Mellon University, *rsalakhu@cs.cmu.edu*
Department of Machine Learning
PhD (2009), Department of Computer Science, University of Toronto.

Ruslan Salakhutdinov is Associate Professor to the Machine Learning Department at Carnegie Mellon University. He received his PhD in computer science from the University of Toronto in 2009 and spent two post-doctoral years at the Massachusetts Institute of Technology Artificial Intelligence Lab and then was Assistant Professor in the Departments of Statistics and Computer Science at the University of Toronto until Spring 2016.

Dr. Salakhutdinov's primary interests lie in deep learning, machine learning, and large-scale optimization. His main research goal is to understand the computational and statistical principles required for discovering structure in large amounts of data. He is an action editor of the Journal of Machine Learning Research and served on the senior programme committee of several learning conferences including NIPS and ICML.

He is an Alfred P. Sloan Research Fellow, Microsoft Research Faculty Fellow, Canada Research Chair in Statistical Machine Learning, a recipient of the Early Researcher Award, Connaught New Researcher Award, Google Faculty Award, and is a Senior Fellow of the Canadian Institute for Advanced Research.

Stefan Schaal – *http://www-clmc.usc.edu/~sschaal/*
University of Southern California, *sschaal@usc.edu*
Departments of Computer Science, Neuroscience, and Biomedical Engineering
PhD (1991), Technical University of Munich

Stefan Schaal is Professor of Computer Science, Neuroscience, and Biomedical Engineering at the University of Southern California, and a Founding Director of the Max-Planck-Insitute for Intelligent Systems in Tuebingen, Germany. He is also an Invited Researcher at the ATR Computational Neuroscience Laboratory in Japan, where he held an appointment as Head of the Computational Learning Group during an international ERATO project, the Kawato Dynamic Brain Project (ERATO/JST). Before joining USC, Dr. Schaal was a postdoctoral fellow at the Department of Brain and Cognitive Sciences and the Artificial Intelligence Laboratory at MIT, an Invited Researcher at the ATR Human Information Processing Research Laboratories in Japan, and an Adjunct Assistant Professor at the Georgia Institute of Technology and at the Department of Kinesiology of the Pennsylvania State University.

Dr. Schaal's research interests include topics of statistical and machine learning, neural networks, computational neuroscience, functional brain imaging, nonlinear dynamics, nonlinear control theory, and biomimetic robotics. He applies his research to problems of artificial and biological motor control and motor learning, focusing on both theoretical investigations and experiments with human subjects and anthropomorphic robot equipment.

Dr. Schaal has co-authored over 300 papers in refereed journals and conferences. He is a co-founder of the "IEEE/RAS International Conference and Humanoid Robotics", and a co-founder of "Robotics Science and Systems", a highly selective new conference featuring the best work in robotics every year. Dr. Schaal served as Program Chair at these conferences and he was the Program Chair of "Simulated and Adaptive Behavior" (SAB 2004) and the "IEEE/RAS International Conference on Robotics and Automation" (ICRA 2008). Dr. Schaal is has also been an Area Chair at "Neural Information Processing Systems" (NIPS) and served as Program Committee Member of the "International Conference on Machine Learning" (ICML). Dr. Schaal serves on the editorial board of the journals "Neural Networks", "International Journal of Humanoid Robotics", and "Frontiers in Neurorobotics".

Dr. Schaal is a member of the German National Academic Foundation (Studienstiftung des Deutschen Volkes), the Alexander von Humboldt Foundation, the Society For Neuroscience, the Society for Neural Control of Movement, the IEEE, and AAAS.

Mac Schwager – *http://web.stanford.edu/~schwager/*
Stanford University, *schwager@stanford.edu*
Department of Aeronautics and Astronautics
PhD (2009), Mechanical Engineering, MIT

Mac Schwager is an assistant Professor of Aeronautics and Astronautics at Stanford University, and is the director of the Multi-robot Systems Lab (MSL) at Stanford.

His research interests are in distributed algorithms for control, perception, and learning in groups of robots, autonomous aerial vehicles, autonomous cars, and animals. He has worked on distributed persistent surveillance for networks of UAVs monitoring large-scale environments, adaptive distributed coverage control for sensing robots, cooperative manipulation, coordination and formation flying for groups of UAVs using onboard vision, agile control of UAV swarms, and human-swarm interfaces. He is also interested in provably safe multi-vehicle feedback interactions in autonomous driving.

He received the NSF CAREER award in 2014, and the Early Career Research Excellence Award from the Boston University College of Engineering in 2015. He received two best conference paper finalist awards at the International Conference on Robotics and Automation (2008, 2011), and received the best conference paper award at the 2008 International Conference on the Simulation of Adaptive Behavior.

Jianbo Shi – *https://www.cis.upenn.edu/~jshi/*
University of Pennsylvania, *jshi@seas.upenn.edu*
Department of Computer & Information Science
PhD (1998), University of California, Berkeley

Jianbo Shi is a Professor of Computer & Information Sciences and a member of the GRASP Laboratory at the University of Pennsylvania. He received his BA in Computer Science and Mathematics from Cornell University. He received his Ph.D. degree in Computer Science from University of California at Berkeley in 1998, for his thesis on normalize cuts image segmentation algorithm. He joined The Robotics Institute at Carnegie Mellon University in 1999 as a research faculty, where he lead the Human Identification at Distance(HumanID) project, developing vision techniques for human identification and activity inference. In 2004, he received a US National Science Foundation CAREER award on learning to see—a unified segmentation and recognition approach.

His current research focuses on human behavior analysis and image recognition-segmentation. His other research interests include image/video retrieval, and vision based desktop computing. His long-term interests center around a broader area of machine intelligence, he wishes to develop a "visual thinking" module that allows computers not only to understand the environment around us, but also to achieve higher level cognitive abilities such as machine memory and learning.

Stefano Soatto – *http://web.cs.ucla.edu/~soatto/*
University of California, Los Angeles, *soatto@cs.ucla.edu*
Department of Computer Science
PhD (1996), Control and Dynamical Systems, California Institute of Technology

Professor Soatto received his Ph.D. in Control and Dynamical Systems from the California Institute of Technology in 1996; he joined UCLA in 2000 after being Assistant and then Associate Professor of Electrical and Biomedical Engineering at Washington University, and Research Associate in Applied Sciences at Harvard University. Between 1995 and 1998 he was also Ricercatore in the Department of Mathematics and Computer Science at the University of Udine—Italy. He received his D.Ing. degree (highest honors) from the University of Padova-Italy in 1992.

Dr. Soatto is the recipient of the David Marr Prize (with Y. Ma, J. Kosecka and S. Sastry of U.C. Berkeley) for work on Euclidean reconstruction and reprojection up to subgroups. He also received the Siemens Prize with the Outstanding Paper Award from the IEEE Computer Society for his work on optimal structure from motion (with R. Brockett of Harvard). He received the National Science Foundation Career Award and the Okawa Foundation Grant. He is Associate Editor of the IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) and a Member of the Editorial Board of the International Journal of Computer Vision (IJCV) and Foundations and Trends in Computer Graphics and Vision.

Gaurav Sukhatme – *http://www-robotics.usc.edu/~gaurav/*
University of Southern California, *gaurav@usc.edu*
Departments of Computer Science and Electrical Engineering
PhD (1997), Computer Science, University of Southern California

Gaurav S. Sukhatme is Dean's Professor of Computer Science (joint appointment in Electrical Engineering) and is currently Chairman of the Computer Science department at the University of Southern California. He leads the USC Robotic Embedded Systems Lab and is Associate Director of the USC Robotics and Autonomous Systems Center (RASC).

Sukhatme's research is in multi-robot systems and robot networks with applications to environmental robotics and on-body systems. He has published extensively in these and related areas. Sukhatme has served as PI on numerous NSF, DARPA and NASA grants. He was a Co-PI on the Center for Embedded Networked Sensing (CENS), an NSF Science and Technology Center. He is a fellow of the IEEE and a recipient of the NSF CAREER award and the Okawa foundation research award. He is one of the founders of the Robotics: Science and Systems conference. He was program chair of the 2008 IEEE International Conference on Robotics and Automation and the 2011 IEEE/RSJ International Conference on Robots and Systems. He is the Editor-in-Chief of Autonomous Robots and has served as Associate Editor of the IEEE Transactions on Robotics and Automation, the IEEE Transactions on Mobile Computing, and on the editorial board of IEEE Pervasive Computing.

Manuela Veloso – *http://www.cs.cmu.edu/~mmv/*
Carnegie Mellon University, *mmv@cs.cmu.edu*
Department of Computer Science
PhD (1992), Computer Science Department, Carnegie Mellon University

Manuela M. Veloso is the Herbert A. Simon University Professor in the Computer Science Department and Machine Learning Department at Carnegie Mellon University, with courtesy appointments in the Robotics Institute, and Electrical and Computer Engineering Department. She researches in Artificial Intelligence and Robotics. She founded and directs the CORAL research laboratory, for the study of autonomous agents that Collaborate, Observe, Reason, Act, and Learn, www.cs.cmu.edu/~coral.

Professor Veloso is IEEE Fellow, AAAS Fellow, AAAI Fellow, and the past President of AAAI and RoboCup. Professor Veloso and her students have worked with a variety of autonomous robots, including mobile service robots and soccer robots. The CoBot service robots have autonomously navigated for more than 1,000km in multi-floor office buildings.

Co-chairs

Daniel E Koditschek
Vijay Kumar
Daniel D. Lee

Contributing authors to the report

Noah Cowan
Maja Mataric
Nicholas Roy
Gaurav Sukhatme
Hadas Kress-Gazit
Lisa Miracchi

Government observers

Dr. Lynne Parker, NSF
Dr. Marc Steinberg, ONR
Dr. Tristan Nguyen, AFOSR
Dr. Greg Chirikjian, NSF
Dr. Behzad Kamgarparsi, ONR
Dr. Jiwei Lu, ASD(R&E)/BRO
Dr. David Han, ASD(R&E)/BRO

University of Pennsylvania Rapporteurs

Nikolay Atanasov
Subhrajit Bhattacharya
Philip Dames
Dan Guralnik
Marcell Missura
Amanda Prorok
Paul Reverdy

VT-ARC Rapporteurs

Brian Hider, Project Manager
Virginia Tech Applied Research Corporation, brian.hider@vt-arc.org
Thomas Hussey, Senior Consultant
Virginia Tech Applied Research Corporation, twhussey@flash.net
Elmer Yglesias, Research Scientist
Virginia Tech Applied Research Corporation, kate.klemic@vt-arc.org

Appendix II
Workshop Organization

The two-day workshop was organized to encourage lively discussion and debate and to maximize the interaction of the participants. After short introductory presentations from each participant, the remainder of the workshop was devoted to a series of small group breakouts with combinatorial mixing of participants in each session.

The short introductory presentations aimed to present the speaker's insights on:

- Foundations of his/her discipline
- Recent breakthroughs in field (and own research)
- Future research directions of field

The breakout sessions addressed:

- What new research areas/opportunities are expected in the next decade?
- What capabilities are achievable in 5 years? 10 years? 15 years?
- How do other research areas ideas/interests synergize with ours?
- What intellectual and/or infrastructure investments are necessary to advance the field?
- Are there any missing disciplines, methods, ideas from the workshop?

Appendix III
Bibliography for Figures

[1] S. Teller, M. R. Walter, M. Antone, A. Correa, R. Davis, L. Fletcher, E. Frazzoli, J. Glass, J. P. How, A. S. Huang, J. H. Jeon, S.Karaman, B. Luders, N. Roy and T. Sainath, "A voice-commandable robotic forklift working alongside humans in minimally-prepared outdoor environments," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on,* 2010, pp. 526–533.

[2] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, "Human-level concept learning through probabilistic program induction," *Science*, vol. 350, no. 6266, pp. 1332–1338, 2015.

[3] N. Hunt, J. Jinn, T. Libby, L. F. Jacobs, and R. J. Full, "Learning to Launch: targeted leaping from a dynamic obstacle in squirrels," presented at the Society of Integrative and Comparative Biology Annual Meeting & Exhibition Final Program and Abstracts, West Palm Beach, FL, 03–Jan–2015.

[4] K. Autumn, Y. Liang, T. Hsieh, W. Zesch, W. P. Chan, R. Kenny, and R. J. Full, "Adhesive force of a single gecko foot-hair," *Nature*, vol. 405, pp. 681–685, 2000.

[5] S. Sefati, I. D. Neveln, E. Roth, T. R. Mitchell, J. B. Snyder, M. A. MacIver, E. S. Fortune, and N. J. Cowan, "Mutually opposing forces during locomotion can eliminate the tradeoff between maneuverability and stability," *Proc. Natl. Acad. Sci.*, vol. 110, no. 47, pp. 18798–18803, 2013.

[6] L. Righetti, M. Kalakrishnan, P. Pastor, J. Binney, J. Kelly, R. C. Voorhies, G. S. Sukhatme, and S. Schaal, "An autonomous manipulation system based on force control and optimization," *Auton. Robots*, vol. 36, no. 1–2, pp. 11–30, 2014.

[7] N. Atanasov, M. Zhu, K. Daniilidis, and G. J. Pappas, "Localization from semantic observations via the matrix permanent," I*nt. J. Robot. Res.*, vol. 35, no. 1–3, pp. 73–99, Jan. 2016.

[8 l] J. Fasola and M. J. Mataric, "Modeling dynamic spatial relations with global properties for natural language-based human-robot interaction," in *RO-MAN, 2013 IEEE*, 2013, pp. 453–460.

[8 r] R. Ventura, B. Coltin, and M. Veloso, "Web-based remote assistance to overcome robot perceptual limitations," in AAAI Conference on Artificial Intelligence (AAAI-13), Workshop on Intelligent Robot Systems. AAAI, Bellevue, WA, 2013.

[9] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, "Contour Detection and Hierarchical Image Segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 5, pp. 898–916, 2011.

[10] H. Soo Park and J. Shi, "Social saliency prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 4777–4785.

[11] D. Nardi, I. Noda, F. Ribeiro, P. Stone, O. von Stryk, and M. Veloso, "RoboCup Soccer Leagues," *AI Mag.*, vol. 35, no. 3, pp. 77–85, Fall 2014.

Appendix IV

Is Physical Intelligence a formal science?

Unsurprisingly, given the acknowledged infancy of the field, there are many other aspects of intelligent physical systems science that elude broad consensus. **Perhaps the most fundamental controversy revealed by this meeting surrounds the very notion of what we might mean by the term "science" at all.** On the one hand, a large cross section of participants finds a compelling opportunity to establish a synthetic science that uses mathematical reasoning to formalize refutable empirical hypotheses about what can be built to operate autonomously in the physical world—and then applies computational and engineering tools to bring those formally expressed designs into the world as useful artifacts. On the other hand, a substantial group expresses deep skepticism about the present (or, in some cases, long term) utility of formal reasoning beyond the basic principles of empiricism, urging that there is no substitute for discovery by building, failing, generating new designs by dint of intuition and experience, then rebuilding, and continually repeating this cycle of design by creatively inspired and experimentally tested tinkering. That said, in both camps, there is broad excitement about the social value and intellectual prospects for pursuing either path or both in some combination.

View A: Physical Intelligence as a Formal Synthetic Science

Many researchers are excited by the prospect of a synthetic science of physical intelligence with firm mathematical foundations that undergird a body of empirical practice and help motivate its expansion. A synthetic science is one whose artifacts comprise its fundamental arbiter of understanding, following Feynman's dictum, quoted above. The example of computer science looms large. Turing's mathematical rendering of what we mean by "algorithms" motivated the quest for physical substrates on which to instantiate them. The social impact of these physically embodied computations gave rise to an explosion of empirical improvements in their operation and fabrication. These advances in practice motivated further mathematical consideration of fundamental limits (e.g., P=NP?) stimulating still broader empirical innovation—a virtuous spiral that launched the information technology revolution which continues apace. For example, McCarthy's empirical discovery that the lambda calculus could undergird a programming language gave rise to a new generation of type theorists and a topological model of semantics, leading eventually to the recent startling proposal for a revised foundation of mathematics itself using homotopy type theory. In turn, the concomitant elevation of what we mean by "program" to the status of a mathematical object gives promise of powerful new tools to secure the embattled software engineering industry.

The birth of physical intelligence as a synthetic science will necessarily take a different course from that of computing since it is fundamentally concerned with the informatics of work. Its origins in *cybernetics* underwent a late twentieth century dispersal into the specialized applied mathematical disciplines of control, communications, information theory and signal processing.[1] The rise of robotics offers a compelling physical substrate on which those disciplines might join with statistics and computing to help lay new foundations targeting synthesis rather than mere analysis of physical intelligence. Missing still are the unifying principles, but a number of underlying needs can be discerned from the lists of "next step" problems in Section 3.2 that suggest a general consensus on what aspects of missing theory might have greatest near term impact.

Defining and quantifying task complexity and agent capacity in relation seem overdue for empirically refutable formal hypotheses. Characterizing the notion of a task domain ought to benefit from a formalization of evolutionary history that teaches us intelligence has been driven by social imperatives such as arise in predator/prey relationships and foraging/agriculture activities. Representations of intent, motivation, and trust within and between physical agents are needed that admit some formalized notion of dynamics and its concomitant estimation and control theories. Formal approaches to design and comparison of agent reward structures mediating interaction with the environment (e.g., optimizing vs. satisficing) or other agents (e.g., decentralized competitive economic exchange vs. centralized cooperative hierarchy) or humans (e.g., trading off independence for subservience to assigned intent) remain to be articulated. In all these dimensions of theory, the opportunity and challenge is to identify how the physically situated mechanism performing work upon its environment invites and necessitates a new, formal view of agency.

Perhaps the most urgent motivation for pursuing a formalism to underlie the field of intelligent physical systems is the potential impact of mischief or even merely mistakes. When a desktop computer crashes we may be irritated but when a life-critical technology fails unexpectedly it is increasingly the case that lives are at stake. The availability of a theorem with regard to some desired or undesired behavior of course does not imply that we must observe what is guaranteed but rather, that we are handed a crucial debugging tool since the failure of the guarantee implies that the presumed hypothesis was wrong.

[1] Many discussions at the meeting involved references to past episodes in the history of synthetic science: computer science and the origins of the information technology revolution; cybernetics—its creation, degeneration and perhaps near future revival. There is likely to be great value in pursuing more careful scholarship respecting such questions.

View B: Physical Intelligence as a Growing Constellation of Empirical Principles

Shaping, experimentation, and empiricism have traditionally played a central role in Engineering Science, particularly in the formative stages of a discipline. Some researchers argue that a synthetic Science of Autonomy—a quest to elucidate the foundations of Intelligent Sensing, Action and Learning—must embrace this tradition in order to be successful. This position is based on two arguments. Both are statements of considered opinion shared to varying extent by a significant subset of the Workshop participants. **This group of scholars is skeptical that it is presently advisable (some question whether it is even possible) to define and pursue a science of autonomy with underlying principles expressed using mathematical formalism.**

The first argument is based on timing. It holds that the state of the field is currently in its mid-to-late childhood where primary reliance on mathematical formalism as a means to facilitate progress is not the appropriate way to proceed. "Premature mathematization," in the words of one workshop attendee, is the "kiss of death." Instead, foundational work is best pursued via support and promotion of a growing constellation of empirical principles—to be formalized mathematically later—when there is more fundamental experimentally grounded insight.

The second argument is based on the assertion that foundational understanding is not synonymous with mathematical formalization. Foundational understanding may be expressed in a core set of (non-mathematical) principles or ideas. This, somewhat more provocatively, forces us to acknowledge the possibility that a science of autonomy may be destined to never be the kind of science that is grounded in mathematical formalization. In this, it may resemble for example, molecular biology, where large swathes of world-class science are done with little or no recourse to mathematical formalism. It is instructive to remember that when that subject was in its infancy, a cohort of physicists turned to molecular biology with the traditional mathematical tools of their (centuries old) trade. Success, both measured in early breakthroughs (e.g., the structure of the double helix and the central dogma due to Watson and Crick) and sustained progress decades later (e.g., retroviruses due to Baltimore) went to the model builders, experimenters, and empirical thinkers—those who explicitly rejected the pursuit of a set of mathematical underpinnings for the subject. One might be tempted to think that this analogy is flawed because molecular biology is a natural science, unlike a science of autonomy. To those who may believe this, a second analogy is worth considering. The internet—arguably one of the most complex and useful artifacts that humans have created—lacks an all-encompassing mathematical foundation.[2] Some have argued that it is always destined to be thus—that this is not a "bug" but a "feature."

In its most succinct form this side of the debate might be characterized by an argument in defense of "hacks." That debate is not yet settled, but it is at least uncontroversial to assert that artifacts of stupendous complexity and tremendous practical utility can be constructed without first creating an underlying mathematical formalism.

[2] This claim might be contested by considering the discovery of a natural stability theory underlying the decentralized congestion algorithms that underly internet traffic control.