REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
17-11-2016	Final Report	1-Sep-2015 - 31-Aug-2016

4. TITLE AND SUBTITLE
Final Report: The Statistical Physics of Stochastic Optimal Control and Learning

5a. CONTRACT NUMBER
W911NF-15-1-0478

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHORS
Evangelos Theodorou

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
Georgia Tech Research Corporation
505 Tenth Street NW

Atlanta, GA 30332 -0420

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
67758-EG-CF.1

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT
This workshop aims to identify interconnections between the areas of statistical physics, stochastic control theory, and learning/adaptation and investigate open questions and future research directions at the intersection of these fields. The investigation will have a theoretical and an application component. Starting with connections between statistical physics and optimal control, concepts from statistical physics have recently been used towards the development of scalable algorithms for stochastic control. Grounded on the fundamental connections between partial differential equations and stochastic differential equations, the aforementioned tools provide probabilistic

15. SUBJECT TERMS
Stochastic Optimal Control, Learning, Dynamical Systems, Statistical Physics, Non-equilibrium Thermodynamics

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Evangelos Theodorou
UU	UU	UU	UU		19b. TELEPHONE NUMBER 404-894-8197

Report Title

Final Report: The Statistical Physics of Stochastic Optimal Control and Learning

ABSTRACT

This workshop aims to identify interconnections between the areas of statistical physics, stochastic control theory, and learning/adaptation and investigate open questions and future research directions at the intersection of these fields. The investigation will have a theoretical and an application component. Starting with connections between statistical physics and optimal control, concepts from statistical physics have recently been used towards the development of scalable algorithms for stochastic control. Grounded on the fundamental connections between partial differential equations and stochastic differential equations, the aforementioned tools provide probabilistic representations of solutions of partial differential equations and generated new algorithms for stochastic control using forward sampling of stochastic differential equations. On the other side, optimality principles in control theory namely the Pontryagin Maximum principle and Dynamic Programming have been used to generalize fluctuation theorems in stochastic thermodynamics. The use of aforementioned optimality principles has contributed towards a better understanding of the role that feedforward and feedback control laws play in stochastic thermodynamic systems. In machine learning, there have been new probabilistic methods for inference, regression and uncertainty representation. Algorithms for statistical inference have interpretations drawn from statistical physics and stochastic thermodynamics. These interpretations have the potential to result in new algorithms for statistical inference and create a direct connection between statistical physics and machine learning/adaptation.

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

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

Received Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

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

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

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

Received Paper

TOTAL:

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

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

Received Paper

TOTAL:

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

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

0

Graduate Students

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Student Metrics

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

The number of undergraduates funded by this agreement who graduated during this period:	0.00
The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:......	0.00
The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:......	0.00
Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):......	0.00
Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:......	0.00
The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense	0.00
The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:......	0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME	PERCENT_SUPPORTED

FTE Equivalent:
Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

There is a report that was created based on the findings of this workshop. See attachment file.

Technology Transfer

N/A

ARO Workshop on the Statistical Physics of Stochastic Optimal Control and Learning

6/14/16 - 6/15/16

Report

PI: Evangelos Theodorou

The workshop took place at Georgia Institute of Technology in Atlanta for two days, June 14 -15 of 2016. The list of the participants included scientists from the areas of statistical physics, stochastic control and dynamical systems theory. The list of participants is given below:

1. Tryphon Georgiou - University of Minnesota
2. Prashant Mehta - University of Illinois Urbana-Champaign
3. Roger Brocket - Harvard university
4. P.S. Krishnaprasad -University of Maryland
5. Jeremy England - Massachusetts Institute of Technology
6. Christopher Jarzynski - University of Maryland
7. Pedro Ortega - University of Pennsylvania
8. David Sivak - Simon Fraser University
9. Aristotle Arapostathis - UT Austin
10. Uwe Claus Täuber - Virginian Tech
11. Dibyendu Mandal - UC Berkeley
12. Vijay Gupta - University of Notre Dame
13. Nathan Kutz - University of Washington
14. Andrew Lamberski - University of Minnesota
15. Dan Goldman - Georgia institute of Technology
16. Wassim Haddad - Georgia institute of Technology
17. Evangelos Theodorou - Georgia institute of Technology

In addition to the faculty participation, there were three visiting graduate students. The names of these students are:

1. Jeery Demers - University of Pennsylvania
2. Sivaranjani Seetharaman - University of Notre Dame

3. Amirhossein Taghvaei - University of Illinois Urbana-Champaign

There was also a number of graduate students from Georgia Tech who joined the workshop. The topics discussed in this workshop were at the intersection of stochastic control theory, learning, dynamical systems theory and statistical physics. The first day of the workshop included two morning and one afternoon sessions with presentations of the participants. The duration for each presentation was 20 minutes. The first day was concluded with a discussion session on open problems. The second day was dedicated to discussion and brainstorming sessions.

The topics discussed during these brainstorming sessions are provided in the next sections. This list of topics is also based on the presentations during the workshop as well as the information provided by the participants after the workshop.

Topic 1 - Stochastic control of distributions and uncertainty representation

Uncertainty is one of the key difficulties in performing control for stochastic systems at the macro and micro level. In systems engineering (macro-level), there are different sources of uncertainty that include uncertainty due to unknown parameters, external stochastic forces and in general uncertainty due to unknown dynamics. For systems at the micro-level such as open quantum systems, uncertainty is a fundamental phenomenon when these theses systems are in interaction with the external environment or they are under continuous measurement/monitoring. Given the different sources of uncertainty for dynamical systems at the macro or micro level, the question of what does it mean to control a random process arises in a very natural way.

During the workshop there was discussion on control methodologies for control of distributions or some finite dimensional representation of these distributions using moments. The discussion is relevant to systems in engineering and robotics but it is also relevant to systems in physics. The workshop highlighted topics for research in the area of control of distributions and applications.

1. *Linear and nonlinear control of distributions:* Current methods are restricted to linear systems or linearized versions of nonlinear stochastic systems and under the assumption that uncertainty is Gaussian. There is the need for novel control algorithms that can steer distributions in the nonlinear case. These control algorithms could result in feedforward and feedback nonlinear policies.

2. *Uncertainty representation:* There are different forms of uncertainty representation which have advantages and disadvantages. For example there forms of uncertainty representation include stochastic differential equations, polynomial chaos theory, Gaussian processes etc. What is the proper representation and how these methods could be combined? Deterministic representations such as polynomial chaos expand the state dimensionality of the problem. Is sparsity an option in these cases? Gaussian Processes on the other hand are computationally expensive since non-parametric regression is based on inverting a kernel matrix that has the dimensionality equal the number of data samples used. Approximate methods for non-parametric regression will be very important in order to perform control of uncertain dynamics.

3. *Finite versus infinite dimensional representations of stochastic systems:* Deterministic and Stochastic Partial Differential Equations (PDEs,SPDEs) can be used to represent uncertainty propagation for the case of fully observable or partial observable nonlinear stochastic dynamical systems. Control of partial observable stochastic systems has an infinite dimensional representation that is related to control of the Zakai stochastic PDE. In the most general case control of partial observable stochastic

systems is related to control SPDEs and is also related to control of open quantum systems. Despite the prior work on SPDEs, most of this work is related to existence and uniqueness of solution and less on the development of scalable algorithms for control of PDEs and SPDEs. Therefore there is the need for research related to the development of scalable iterative algorithms for control of PDEs and SPDEs. From the algorithmic point of view, this work could combine methods from machine learning on the representation of probability measures and distributions as well as recent advancements in high-performance computing related to parallelization, stochastic and/or neuro-morphic computation. While infinite dimensional representations are important for the case of partial observable stochastic control, they do also appear in areas of statistical physics related to critical phenomena, synchronization and renormalization group theory. This area was also discussed during the workshop and is presented in more detail in the Topic 3 section of this report.

Topic 2 - Adaptation and matching of internal with driven dynamics

Recent theoretical progress in nonequilibrium statistical mechanics has established precise quantitative relationships between thermodynamic fluxes and kinetic parameters controlling dynamics in arbitrarily driven far-from-equilibrium systems. Of particular interest is the fact that one of the major factors affecting the likelihood of a system configuration in a driven stochastic evolution is the history of work absorption leading to the outcome. This means that, in many-body systems where holding other factors like internal energy and diffusive accessibility fixed is a weak constraint, likely dynamical outcomes are expected to appear to have been selected in part for their exceptional work histories. Essentially, the system explores its phase space through a biased random walk in which the most durable, irreversible changes in configuration occur at moments when the system absorbs and dissipates the most work.

The relationship between work absorption and dynamics is particularly of interest in systems where the driving environment presents a ”challenge” to the system. Phenomena such as resonance provide examples where particular choices of internal structure within the system can lead to a matching to external drives that dramatically increases the rate of work absorption. In such scenarios, the relationship between work history and dynamical likelihood implies a relationship between structural history and likelihood, raising the possibility that understanding how system structure controls the flow of work may in turn lead to predictive principles for self-organization in the presence of challenging drives. Some questions of particular interest related to this topic are given below:

- **Emergent Prediction:** Some environments are random, but with predictable correlations. Could this lead to pressure on a many-body system to instantiate predictive computations and behaviors?
- **Emergent Centralization:** Fluctuating environments contain correlations on a wide range of time-scales. Can we understand how slow correlations in a drive couple to slow degrees of freedom in self-assembly?

The notion of internal dynamics may correspond to underlying structure, geometry, basis functions.

Topic 3 - Critical phenomena, synchronization and renormalization group theory

The objective is to build on a powerful and highly successful theory for non-linear stochastic dynamics of cooperative multi-component systems, namely critical dynamics, and develop novel efficient protocols for:

(1) steering multi-critical complex interacting dynamical systems toward certain desired universal scaling behavior;

(2) externally controlling the strength of stochastic fluctuations and intrinsic noise in systems that are driven far from thermal equilibrium and display generic scale invariance;

(3) selectively targeting and stabilizing specific self-generated spatio-temporal patterns in strongly fluctuating reaction-diffusion systems and epidemic models.

This interdisciplinary research is anticipated to lead to the emergence of new concepts and innovative mathematical tools, and should be relevant to a wide range of potential applications that span from materials science, *e.g.*, magnetism and surface growth, to synthetic biology, neuroscience, epidemiology, ecology, and social system dynamics.

The study of critical phenomena has been a major research focus in statistical physics for the past five decades. Physical systems near a critical point develop strong, long-ranged correlations that render traditional approximation schemes for interacting many-particle systems obsolete, yet also induce the emergence of universal features that are independent of many underlying microscopic details. The almost simultaneous development of powerful computer simulations and the renormalization group approach to target and understand scale invariance have led to a thorough characterization of static as well as dynamic critical phenomena. The latter even include continuous dynamic phase transitions and generic scale invariance in systems that are driven far from thermal equilibrium. Whereas a considerable body of research has addressed the detailed quantitative characterization, classification, and mathematical description of stochastic scale-invariant critical dynamics, there have been only few and unsystematic attempts to exert control over complex cooperative systems. We identify the following five central and fundamental questions:

- Can we design efficient control mechanisms that would allow us to influence and drive scale-invariant critical dynamics towards certain targeted universality classes ?
- Is it possible to externally control the strength of system-imminent fluctuations and even intrinsic reaction noise amplitudes to achieve desired observable output ?
- May one exploit large-scale collective behavior of near-critical multiple-agent dynamical systems to exert global control through correlation-facilitated spreading of carefully posited local perturbations ?
- Can we construct control schemes to specifically select for certain fluctuation-induced and noise-stabilized spatio-temporal patterns to generate desired morphologies in interacting many-particle and reaction-diffusion systems ?
- Is it feasible to extend powerful renormalization group ideas and techniques to systematically describe and classify control schemes in near-critical dynamical sytems ?

Topic 4 - Stochastic control at the micro scale

Biomolecular protein machines couple distinct activities at distant sites, responding to their fluctuating environment and exerting (stochastic) control over their targets. These machines run the gamut from rotary electrochemical transducers like ATP synthase to mechanochemical machines like the transport motors kinesin, dynein, and myosin, to information-processing complexes such as DNA/RNA polymerases and ribosomes. Understanding their basic operational principles will require fundamental advances in nonequilibrium statistical mechanics, informed by control theory and learning algorithms. Advances in these areas will lead to broader understanding of the limits in evolved molecular function and the resulting constraints on myriad biophysical mechanisms, and will provide guiding principles for the design of artificial nanoscale machines.

We identified the following major questions at the intersection of control theory and molecular energy and information transduction. What general principles govern the control of molecular machine function, given their complex internal dynamics connecting control points and active sites? How do molecular machines achieve reliable control in the presence of strong thermal fluctuations? What are the limits on the feedback control these machines can achieve when the controllers in question can only store a small number of bits of information?

When two strongly fluctuating microscopic systems are tightly coupled, each influences and is influenced by the other. Can we characterize these interactions in term of measurement and feedback, and in this context what are the thermodynamic implications/costs of control? Are there fundamental tradeoffs between the efficacy of the control, the speed of operation, and the rate of energy dissipation? Have biomolecular machines evolved to operate near the limits implied by such tradeoffs? What are the design principles for constructing artificial molecular machines that operate near these limits?

While the second law of thermodynamics constrains the amount of work that can be extracted during a thermodynamic process, feedback fluctuation theorems quantify the extent to which these constraints are modified by the closed-loop control of thermal fluctuations. What do these theorems imply for stochastic optimal control theory? What efficient methods from control theory can identify the most efficacious control nodes of microscopic systems, and how to control them with both static and dynamic perturbations?

In the physics literature, shortcuts to adiabaticity have emerged as open-loop strategies for accelerating the evolution of classical or quantum systems from one equilibrium state to another. Can these strategies be extended to both open-loop and closed-loop control of stochastic systems, thereby providing linear and nonlinear control methodologies for rapidly steering probabilistic states/distributions in the presence of thermal noise? Applications involve nano-oscilators, nano thermal engines and monitoring of chemical and biological processes.

These control questions naturally raise complementary issues in learning and the evolution or design of controllable systems. How are the internal energetics and dynamics of molecular machines specifically tuned to harness environmental fluctuations? What aspects of their environment have they learned to match? What are the implications of recently discovered connections between response programs that are energetically efficient and those that are informationally parsimonious? Can the principles and methods of machine learning be applied to the design of artificial molecular machines or strategies for self-assembly?

Nature has evolved these molecular-scale machines to accomplish many tasks that are similar to our own technological aspirations, and thus insights into how nature has solved control problems at the nanoscale will point to design principles for novel synthetic machines built from strongly fluctuating soft matter. In particular, the engineering of truly autonomous machines will require the harnessing of nonequilibrium boundary conditions and hence a deeper integration of control theory and active matter, as well as careful study of energetic efficiency.

Topic 5 - Stochastic control for growth processes

Significant progress has been made in control theory by focusing attention on the manipulation and regulation of behavior over extended spatial scales both in (i) physical space -swarms, networks etc; and (ii) phase space- beyond steady states to periodic and chaotic orbits, other invariant manifolds, and bifurcations of the same. In essential ways this progress has been limited to deterministic settings, largely ignoring stochastic phenomena. A similar program may be developed in settings where fluctuations cannot be ignored. One proceeds by selecting model problems, for instance probability laws derived from interactions within a population of agents (such as dissipative stochastic oscillators, biological cells, and economic decision makers). Discovery of fundamental principles that govern a program to manipulate dynamical features of such laws may take a concerted effort marshaling concepts and techniques from stochastic control, geometry of dynamical systems, and statistical physics.

Growth phenomena in complex systems (e.g. the formation of material films, the spread of information) are studied intensively by specialists in materials science, network science, statistical science and other related fields. The question of control (and steering) of such processes has received scant attention. However there are effective models available in this context that can be viewed from the perspective of system theory. Specifically one can ask what robust (universal) features of such models can be steered by the use of adequate control authority. How would one exploit such *controllability* in meeting performance objectives? Some of the growth models take the form of stochastic perturbations of well-known partial differential equations such as the Hamilton-Jacobi-Bellman (HJB) equation of stochastic control. Our understanding of HJB equations may be of relevance to advancing control of growth processes.

Topic 6 - Stochastic control and learning in physics and engineering

How do dynamical systems encode information on their environment, match energetic/information transfer, and embed this information in their available degrees of freedom and ultimately transition from ”reinforcement learning” to a model-predictive/anticipatory mode? There are three axes that could potentially define the frame of reference within which one can address this topic.

Spatial Scale - Micro vs Macro: Studying systems at different spatial scales is key to understanding whether there are any universal principles or laws related to how systems encode information about their environment and adapt/learn to handle dynamically changing situations. There are clearly interesting challenges on what is the underlying representation and how this representation relates to spatial scale. For example during the workshop there was discussion on deriving stochastic controllers for dynamic processes at different time scales using information theoretic concepts. The basic idea relies on the minimization of what physicists call free energy difference based on probability measures that are assigned to stochastic differential equations. This free energy difference corresponds to cost functions that typically appear in stochastic optimal control theory. If one would like to study systems at micro level let say for example a system consisting of millions of nano-particles, the stochastic differential equation may not be the right representation to use. One possibility could be to use SPDEs and so then the questions that arise are

1. How do information theoretic concepts such as free energy, relative entropy and free energy difference generalize to the case of stochastic dynamics systems in infinite dimensional spaces?
2. Can we derive optimal controllers for the stochastic systems in infinite dimensional spaces using these information theoretic measures?
3. What methods or techniques are required to address issues such as scalability? Should we sacrifice global optimality in favor of scalability and numerical efficiency?
4. Can we go from theorems on existence and uniqueness of solutions to iterative algorithms with convergence analysis to compute solutions?

These questions arise when one goes from the scale of a one robot/particle to a scale of millions of nano-robots or particles. In addition, optimization of underlying structure in a thermodynamically efficient way may mean adaptation of SDEs or adaptation of SPDEs depending on what is the spatial scale at which optimization takes place. The algorithms and computational methodologies that will be developed here will have direct applications to mechanical and robotics systems with dynamics represented as PDEs or

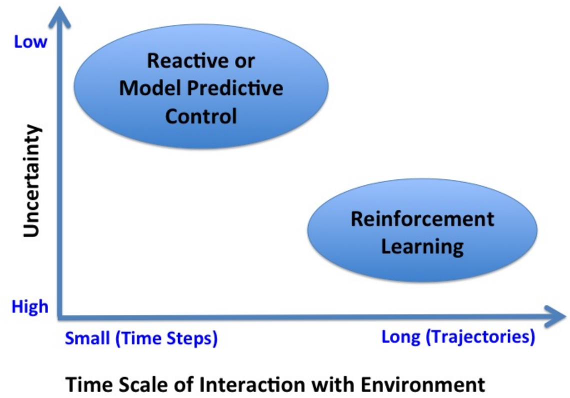

Figure 1: The y axis corresponds to levels of uncertainty. The uncertainty is related to system dynamics model uncertainty as well as uncertainty about the environment. The x axis corresponds to the time scale of interaction of the system under consideration with the environment.

SPDEs. Examples of such systems are soft robotics systems, morphing systems that are characterized by their terrestrial or aerial agility, biologically inspired robotic systems such as the mechanical equivalent of an elephant trunk, etc. Finally, control of SPDEs and PDEs is important for open problems and research questions discussed in Topic 3.

Time Scale - Small vs long time scale and uncertainty: Adaptation and learning at different time scales is very much related to the levels of the uncertainty in the environment (see figure 1). For a controller to be reactive we typically required that the model of the dynamical system is known and also that there is a meaningful cost function for which we known its mathematical representation as a function of the state of the system and the environment. Thus there are two things that have to be known with some certainty, the underlying dynamics of the system under consideration and a model of the world or a model of the interactions of the systems with the world. In a reinforcement learning scenario none of these may be known, namely neither the dynamics nor the environment/cost function. In this sense reinforcement learning promises to find policies and make decisions in model free fashion but in time scales that require extensive interaction with the environment. So there is a trade-off between uncertainty and time scale of decision. In engineered systems very often cost functions or rewards are pre-specified and the question is to learn/adapt the underlying dynamics.

There are also situations in which the cost or reward function is unknown and one has to work with inverse reinforcement learning/inverse optimal control. A hypothesis about biological systems, organisms/micro-organisms or biological structures is that they try to guess the environment within which they operate and find a mapping between their own state and external rewards or external driving signals. So this is more like the inverse reinforcement learning scenario. In all these cases, learning and adaptation could be performed by changing the internal structure of the underlying dynamics and this change can be framed as a stochastic optimal control and statistical inference problem. The underlying dynamics could be represented as SDEs and deterministic/stochastic PDEs.

In conclusion, *there is the need for algorithms that can start from learning and control at the time scale of trajectories or rollouts and then smoothly transition to reactive modes or control policies as more data are collected from the environment and models/representations can be learned.* Reactive controllers are typically controllers that result from optimization at every time step. Since uncertainty is the key factor for making a system to switch adaptation from long to short time scales, the representation of uncertainty plays a very important role, drawing on results in different methodologies in the areas of scientific computing, uncertainty quantification, and statistical machine learning.

Sciences (chemistry-physics) versus engineering: Dynamical systems in physics and chemistry may differ from systems in engineering not only in terms of spatial and time scale but also in terms of actuation and sensing capabilities. In engineering we design systems and decide how many, what kind and where in the design we are are going to place sensors and actuators. In examples in chemistry and physics, the systems are given and therefore our ability to make them more or less controllable/observable is limited. This is a component that should be incorporated in this big idea because it is also related to the underlying structure. There may be cases where engineered systems have to mimic behavior at the micro level. However, it may be possible that we find more efficient ways to adapt the structure of engineered systems for the purposes of learning and adaptation.