

AFRL-IF-RS-TR-2006-156
Final Technical Report
May 2006



THE DYNAMICS OF LEARNING AND THE EMERGENCE OF DISTRIBUTED ADAPTATION

Santa Fe Institute

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K359/00

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2006-156 has been reviewed and is approved for publication

APPROVED: /s/

ROBERT J. PARAGI
Project Engineer

FOR THE DIRECTOR: /s/

JAMES W. CUSACK
Chief, Information Systems Division
Information Directorate

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE MAY 2006		3. REPORT TYPE AND DATES COVERED Final Jun 00 – Dec 05
4. TITLE AND SUBTITLE THE DYNAMICS OF LEARNING AND THE EMERGENCE OF DISTRIBUTED ADAPTATION			5. FUNDING NUMBERS C - F30602-00-2-0583 PE - 62301E PR - TASK TA - 00 WU - 03	
6. AUTHOR(S) James P. Crutchfield				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Santa Fe Institute 1399 Hyde Park Road Santa Fe NM 87501			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFSB 3701 North Fairfax Drive 525 Brooks Road Arlington Virginia 22203-1714 Rome New York 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2006-156	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Robert Paragi, IFSB, 315-330-3547, Robert.Paragi@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This project developed fundamental theory and novel algorithms for adaptive learning in autonomous collective-agent systems. The first goal was to develop a new mathematical framework for analyzing the dynamical mechanisms that support learning in a range of novel information processing substrates. The second goal was to adapt this single-agent learning theory and associated learning algorithms to the distributed setting in which a population of autonomous agents communicate to achieve a desired group task. The results provide a mathematically sound basis for quantitatively measuring 1) the degree of individual-agent intelligence; 2) the significance of agent-environment interaction; and 3) the emergence of cooperation in agent collectives. A novel feature was that the theoretical approach synthesized recent results in the areas of machine learning, statistical inference, statistical mechanics, nonlinear dynamics, and pattern formation theory and then applied them to single-agent learning and adaptive learning in agent collectives. The central impact of the project will be its systematic and quantitative approach to predicting behavioral complexity in agents and in their environments. This will provide a foundation for measuring agent and agent-collective intelligence which, in turn, should allow for systematic engineering and monitoring of these systems.				
14. SUBJECT TERMS Autonomous agents, software agent system dynamics, intelligent agents				15. NUMBER OF PAGES 42
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Contents

I. Summary	1
II. Introduction	2
A. Dynamics of Learning	3
B. Emergence of Distributed Adaptation	3
III. Methods, Assumptions, and Procedures	4
A. Brief Overview of the Approach	4
B. Relationship to Existing Theory	5
IV. Results and Discussion	6
A. Task 1: Dynamics of Learning	7
1. Computational Mechanics: Foundations	8
2. Computational Mechanics: Algorithms	8
3. Single-Agent Learning	10
4. Discovering Spatial Patterns	10
5. Finite-Data Scaling of Pattern Discovery	11
B. Task 2: Emergence of Distributed Adaptation	11
1. ϵ -Machine Nets as Models of Agent Collectives	12
2. Finite-Population Dynamical Systems	13
3. Evolution of Finite-Memory Agents	13
4. Self-Organizing Agent Population Dynamics	13
5. Theory and Analysis of Emergent Collective Adaptation	14
6. Persistent Chaos in High Dimensions	14
7. Dynamics of Multiagent Games	14
8. Multiple Agents Servicing Multiple Tasks	15
9. Causal Synchrony	15
C. TASK Demonstration Project	15
1. MAS Dynamics Simulator	16
2. Robotic MAS Development System (RoMADS)	16
3. RoMADS Sensor Nets	17
4. Open Experiment Framework Related Experiments	17
5. System Development Status	18
6. RoMADS Experiments on Multiagent Systems	18
V. Conclusions	19
A. Sponsored Workshops	20
B. Publications and Reports	20
C. Presentations	23

D. Presentations Hosted by the Project	30
References	31

I. SUMMARY

Vision: To address the main concerns of the Taskable Agent Software Kit (TASK) portion of DARPA’s Agent-Based Computing Program we developed fundamental theory and novel algorithms for adaptive learning in autonomous collective-agent systems. There were two major goals. The first was a mathematical framework for analyzing the dynamical mechanisms that support learning in a range of novel information processing substrates. The second was adapting this single-agent learning theory and associated learning algorithms to the distributed setting in which a population of autonomous agents communicate to achieve a desired group task. The results provide a mathematically sound basis for quantitatively measuring (i) the degree of individual-agent intelligence, (ii) the significance of agent-environment interaction, and (iii) the emergence (or lack thereof) of cooperation in agent collectives.

Innovative Ideas and Tools: A novel feature of the projects is that the theoretical approach synthesized recent results in the areas of machine learning, statistical inference, statistical mechanics, nonlinear dynamics, pattern formation theory, and the developing area of complex systems and then applied them to (i) single-agent learning and (ii) adaptive learning in agent collectives. The project also had unique and regular access to experts in a wide range of disciplines—such as, biology, economics, cognitive science, and sociology—that historically have investigated collective behavior. The phenomena, insights, and results in these disciplines was particularly helpful in grounding the project’s approach to fundamental theory.

Central Hypothesis and Experiments: The first part of our central working hypothesis was that intelligent behavior of autonomous agents requires them (individually) to adaptively estimate the complexity of their environment (including other agents) and build internal dynamical models that capture the patterns and regularities which they detect. The second part of the hypothesis was that measures of cooperation in agent collectives can be built out of the theory describing individual agent capabilities.

The basic effort established the mathematical foundations of what constitutes useful patterns and regularities in environmental stimuli and how these can be adaptively modeled and used as a basis for action and decision-making by autonomous agents. The hypothesis was evaluated by developing mathematical models, rigorously bounding their minimal information processing and computational properties, developing analytical approximations that predict single-agent and agent-collective behavior, and comparing these results with extensive simulations for refining the mathematical framework and for final validation the approach. We developed an experimental distributed robotic platform to test these results.

Expected Impact: A sizable literature exists on agent-based systems.[1–14] Given this and the diversity of proposals for designing and implementing intelligent agents and the wide range of collective behaviors exhibited by interacting agents, [15–47] we believe the central impact of the project will be its systematic and quantitative approach to predicting behavioral complexity in agents and in their environments. This will provide a foundation for measuring agent and agent-collective intelligence which, in turn, should allow for systematic engineering and monitoring of these systems.

II. INTRODUCTION

This project addressed two questions: How can we design agent collectives to achieve a desired global function? And, how can we do this so that the collective behavior is robust? To address various aspects of these questions, we focused on several projects whose overall intent was to develop novel algorithms for adaptive learning in autonomous agent systems. There were two major goals. The first was developing a new dynamical framework for analyzing the nonlinear processes that support learning and adaptation in individual agents. The second was generalizing this single-agent learning theory to the distributed setting in which a population of agents communicate to achieve a group task.

A practical objective in these was to develop methods for detecting, monitoring, and quantifying the degree of cooperation that emerges in interacting multiagent systems. Another was to develop algorithms that would allow individual adaptive agents to build models of, and so come to predict, their environments.

A brief chronology will help show how the project addressed these goals over its lifetime.

From its start in June 2000 to July 2001, the project was involved in establishing the Dynamics of Learning group at the Santa Fe Institute. The primary research focus was on single-agent learning and adapting methods from computational mechanics to that problem.

From August 2001 to July 2002, a bridge was established between the dynamics of single-agent learning and the notions of randomness and structural complexity from computational mechanics. In particular, the outlines of an algorithm to incrementally infer causal states was worked out. The algorithm was a key component in our analysis of single-agent learning. In short, the goal of an agent's adaptive dynamics is to infer a causal model (ϵ -machine, see below) of its environment.

From August 2002 to July 2003, the project investigated how complicated a multiagent system's dynamics can be. We developed new tools, both analytical and software, for exploring this kind of behavioral diversity and showed that the full range of nonlinear behaviors from stable oscillations to deterministic chaos can occur. We also established a modeling framework that allowed us to analyze the mechanisms that lead to such behaviors.

From August July 2003 to July 2004, the key effort was to find methods to grapple with the complicated behaviors that multiagent system's exhibit. In addition to developing a dynamical systems theory analysis for these, a central part was the development of our Robotic MultiAgent Development System (RoMADS) experimental platform. RoMADS allows for real, physical experimentation with robotic vehicle collectives.

From August 2004 to December 2005, the project's termination, the focus was on the TASK Demonstration project and developing a new version of the RoMADS system. We also focused on innovations in the dynamical systems modeling of interacting agents that allow one to study the dynamical stability (and instability) of a much wider range of agent adaptation algorithms than previously.

A. Dynamics of Learning

Some time ago, we introduced an approach, called *computational mechanics*, [48, 49] for detecting information processing capabilities embedded in nonlinear processes. Computational mechanics methods allow one to delineate (i) how much historical information is stored in a process’s current state, (ii) how this memory is organized, and (iii) how this organization produces future behavior. We refer to these three properties as a process’s *intrinsic computation*.

An important component of computational mechanics is its constructive approach to intrinsic computation: it provides algorithms for estimating intrinsic computational properties. This leads naturally to the present project in that these algorithms, and the theory describing them, give a principled approach to learning optimal models of nonlinear processes.

Thus, the central approach was to use computational mechanics to develop a theory of how a dynamical system, as part of its behavior, can learn a model of its “environment”; that is, of signals imposed on it. This type of dynamical learner was our model for the processes controlling intelligent agents. One of the key ideas here is that of *intrinsic representations* introduced by us. [48, 49] The learning theory we developed showed how an agent can infer these representations and use them to predict the environment’s future behavior and, in this way, to take advantage of environmental regularities.

It is important to emphasize that learning in a natural environment is not like the learning of language by humans. The latter is certainly an important topic; it provides many of the metaphors for learning in artificial intelligence. Learning in a natural environment, however, is not constrained by a human language syntax nor even by the need for symbolic codification. Therefore, learning in a natural environment cannot rely on innate or required grammatical or other linguistic constraints. In fact, the analogues of such constraints, if they exist, must be learned from the environment. These constraints reflect real-world geometry, physical law, and a wide-range of nonlinear pattern-forming processes that nature expresses.

We believe that analogies with natural language acquisition have led to the wrong conceptual and mathematical framework for (i) learning in natural environments and (ii) distributed adaptation as expressed by collectives of adaptive agents. These problems are one reason we took an alternative approach to learning and adaptation. Specifically, computational mechanics does not suffer from the problems inherent in symbolic approaches to intelligent behavior and its design.

B. Emergence of Distributed Adaptation

Once the basic theory of dynamical learning was fleshed out, we began—in a principled, analytical, and quantitative way—to investigate the interaction of individual learners in various types of structured populations. We refer to these as *agent collectives*. This approach—first focusing on individual learning and then collective dynamics—greatly reduced the number of parameters (and conceptual degrees of freedom) so that we could cleanly identify those problems and properties that are the result only of collective interactions. Our goal was, after all, to determine what agent collectives can do that is novel,

unique, and ultimately better (or worse) than other alternatives to performing a given task.

To be more explicit, using the computational mechanics approach we addressed the following three questions.

1. What is “distributed” adaptation?
2. How can we detect and design the “emergence” of distributed adaptation?
3. What is a true “collective”? Does the group do more than the sum of the individual agents?

The following sections briefly expand on these questions so that we can detail our chosen projects on the dynamics of learning and the emergence of distributed adaptation and their results. The next section addresses how our approach is different from and also similar to current approaches to learning, on the one hand, and distributed artificial intelligence, on the other. The report then gives a synopsis of the individual problems and results developed during the project. It concludes with a few general remarks. The appendices provide lists of workshops the project organized, the publications produced, and the presentations given by group members on the project. Finally, the last appendix gives a list of visiting scholars’ presentations to the Dynamics of Learning group.

III. METHODS, ASSUMPTIONS, AND PROCEDURES

Our approach used concepts and analytical tools from statistical and nonlinear physics, theoretical computer science, and communication theory. What is novel about the approach was the integration we achieved of these disparate fields for the problem of detecting adaptation in individual agents and cooperation in multiagent systems.

In the past, detecting emergent cooperation has been done subjectively. Our goal was to understand the phenomenon well enough so that we could develop analytical and predictive theories of emergent cooperation and algorithms that automatically (and objectively) quantify and monitor cooperation.

We believe that new design methods for multiagent systems will follow directly from our results. Moreover, we believe that our results will lead to diagnostic tools for large-scale distributed systems that consist of semi-autonomous intelligent agents.

The wide applicability of our approach is one of its most promising aspects. First, the basic problem of emergent cooperation occurs in many natural and artificial systems. Second, the results should be applicable to autonomous robotic vehicles, for example, as well as to software agents on the web.

A. Brief Overview of the Approach

Having sketched the basic problems, it is worthwhile to outline what we saw as the distinctive concepts underlying our approach.

1. *Informational complexity of learning*: We developed, using recent results in machine learning and computational learning theory,[3, 12, 13, 50–53] a theory that describes the constraints on and requirements for learning models of environments with varying degrees of complexity.
2. *Dynamical setting*: We adapted these results and the theory of trade-offs into a dynamical setting to develop algorithms for time-dependent, incremental, and adaptive learning.
3. *Distributed adaptation*: We then extended the analysis of learning dynamics to a distributed setting, in which collections of largely autonomous adaptive agents interact via some communication network.
4. *Evolutionary and ecological setting*:[9, 14, 15, 17, 35, 50, 54–58] We merged these results with evolutionary adaptation by developing the appropriate mathematical population dynamics. In particular, we investigated these results for finite collections of interacting finite-memory agents.

B. Relationship to Existing Theory

The project was only possible due to recent advances in several fields. Computational mechanics borrows heavily from several areas. To get a sense of the overall approach, it is useful to briefly mention several of these fields and what they contributed.

1. *Dynamical systems theory*:[59–61] Geometric analysis of state-space structures that guide nonlinear systems in their temporal behavior; the control of chaotic dynamical systems using very small signals to manipulate and suppress inherent instabilities.[62]
2. *Statistical mechanics*:[63–66] Models of how macroscopic behavior emerges from microscopic equations of motion. Can we detect and then characterize, at a “thermodynamic” level, collective (and, hopefully, useful) behavior in populations of adaptive agents?
3. *Modern statistical inference*[67, 68] for nonlinear processes: Trade-offs between sample size and model complexity. This, to mention just one example, has not been extended to spatial or distributed collections of adaptive agents.
4. *Pattern formation theory, phase transitions, and critical phenomena*:[69, 70] How dynamical and statistical ensemble mechanisms lead to cooperative behavior.
5. *Computational learning theory*, including aspects of machine learning and AI:[3, 12, 13, 50–53, 55, 71] Recent progress in understanding the resource trade-offs in various types of learning paradigm: artificial neural networks, hidden Markov models, and more abstract settings, such as Valiant’s PAC learning,[72] but within distribution-dependent learning paradigms.

Information theory[73] and computation theory[74, 75] also played important roles, though largely by providing basic tools and observables for analysis.

IV. RESULTS AND DISCUSSION

We reformulated and refined the main questions given above in the following ways—ways that motivated our selection of subprojects.

1. How little intelligence can autonomous agents use and still achieve a given level of functionality—individually and collectively?
2. What additional, global capabilities (if any) do agent collectives have?
3. Does global capability allow individual agents to be simpler still?
4. How do we detect (quantitatively) that a given designed collective is, in fact, more functional than another?
5. What is the balance of local functionality and distributed cooperation?
6. How much flexibility is there in achieving such a balance?
7. How do different communication topologies affect global function and an agent’s minimal required intelligence?
8. Can we predict that a given collective is more functional, adaptive, or robust than another?

For each question and related subproject a number of ancillary, but common issues arise:

1. How do we define a workable, measurable notion of adaptation?
2. How do we define a workable, measurable notion of functionality?
3. How do we define a workable, measurable notion of emergence?
4. What observables are appropriate for experimental detection of these properties and for theory development?
5. How can we define robustness in learning, adaptation, and evolution?
6. What are the resource trade-offs between, e.g.,
 - (a) Sample size versus model complexity,
 - (b) Response time versus model complexity,
 - (c) Collective size versus global function, and
 - (d) Collective size versus individual intelligence?

We selected several projects to address various aspects of the above questions and ancillary issues. The chosen projects cannot decisively answer these questions, if the latter are construed too broadly. With this and the size of the overall effort in mind, we focused on the following subprojects according to (i) their importance to the main questions, (ii) the likelihood of success in tackling them, and (iii) our existing expertise and available computing resources. These were grouped into two broad efforts, labeled Task 1 and Task 2.

Task 1. For the dynamics of agent learning the selected projects are the following:

1. Computational Mechanics: Foundations.
2. Computational Mechanics: Algorithms.
3. Single-Agent Learning.
4. Discovering Spatial Patterns.
5. Finite-Data Scaling of Pattern Discovery.

Task 2. For analyzing the emergence of distributed adaptation in agent collectives the projects are the following.

1. ϵ -Machine Nets as Models of Agent Collectives.
2. Finite-Population Dynamical Systems.
3. Evolution of Finite-Memory Agents.
4. Self-Organizing Agent Population Dynamics.
5. Theory and Analysis of Emergent Collective Adaptation.
6. Persistent Chaos in High Dimensions.
7. Dynamics of Multiagent Games.
8. Multiple Agents Servicing Multiple Tasks.
9. Causal Synchrony.
10. Robotic MultiAgent Development System (RoMADS).

A. Task 1: Dynamics of Learning

We continued to develop our multidisciplinary approach to computation in dynamical systems and to inferring models (specifically, ϵ -machines) of nonlinear processes, as the central mathematical framework for the dynamics of single-agent learning. We refined our existing tools in ways that facilitated our analyzing the intrinsic computation in distributed

dynamical systems. The new emphasis, however, was to understand the resource trade-offs in changing model classes to produce optimal models and forecasts from temporal and spatio-temporal data series.

These refinements were necessary for practical reasons. In autonomous agent learning sensory data is limited in quantity, has low signal-to-noise ratio, and is nonstationary. Thus, efficient algorithms for ϵ -machine reconstruction are essential. Developing these algorithms, in turn, requires a deeper appreciation of the dynamics of learning and adaptation. Aside from the pragmatic goals, the related projects elucidated a number of fundamental resource trade-offs present in learning and adaptation.

1. *Computational Mechanics: Foundations*

The basis for this was the notion of an optimal representation afforded by causal states and ϵ -machines. We had already established a suite of theorems that provide the mathematical foundations of computational mechanics.[48, 76] Specifically, we proved that as one ranges over alternative choices of “state” the ϵ -machine has a three-fold optimality: (i) no set of alternative states is more informative about future observables than the causal states; of those choices of state that are as predictive as the causal states, none has (ii) a smaller statistical complexity nor (iii) a smaller entropy rate over the internal states. Much of the early theoretical development was concerned with continuing exploration and refinement of the mathematical foundations of computational mechanics. We are particularly interested in analyzing the effects of measurement distortion on these types of optimality.

During this project we developed means of explicitly calculating the average-case growth rates of most standard complexity measures (algorithmic information, thermodynamic depth, logical depth, sophistication, minimum description length, and the like) from the causal state representation of a complex system. This work provided the first means for explicit calculation of most of these measures and showed that many of them are asymptotically equivalent (e.g., thermodynamic depth and sophistication).

We successfully extended computational mechanics to use measure-theoretic foundations.[77] This will allow for a much wider range of applications of our notions of causal structure—for example, to statistical mixture processes and continuous-valued processes.

2. *Computational Mechanics: Algorithms*

There were a number of new algorithms that we developed as we adapted computational mechanics to single-agent and multiagent learning.

a. Agent adaptation through incremental learning We completed our new incremental learning—Causal State Splitting Reconstruction (CSSR or “scissor”)—algorithm. This infers causal models from time series by adaptively splitting candidate causal states. We proved its convergence on the optimal model using a combination of conditional-independence and large-deviations arguments. We also performed a large-scale Monte Carlo simulation to

experimentally demonstrate the algorithm’s feasibility.[78, 79]

b. Real-time dynamical learning algorithms The goal in this was to update and modify the structure of an estimated ϵ -machine as measurement data is sensed—i.e., in real-time. This is a key step in practical implementations of ϵ -machine reconstruction for autonomous agents. Additionally, this type of incremental learning formed the basis of our studies of adaptation in ϵ -machine nets, which are described below.

With an incremental version of ϵ -machine reconstruction, a number of questions about adaptive learning arise that we attempted to answer. Outstanding ones were the following.

1. Is incremental learning necessarily faster than our current batch reconstruction methods?
2. To which learning problems is incremental learning best adapted?
3. How does the on-line property compromise the descriptive capability of ϵ -machines?
4. Is incremental learning more parsimonious in its use of available resources and input data?

The goal was to embed the on-line learning algorithm into a dynamical system itself, so that the system’s attractor has a subspace that is a model of an applied signal. One important aspect of analyzing such a real-time dynamical learner is the need for a formal definition of adaptation. This is necessary to monitor how well the learning is being performed, as well as how robust it is to fluctuations and nonstationarity. Several extensions of existing dynamical systems measures—e.g., conditional Lyapunov exponents or conditional statistical complexity—were explored as a basis for a quantitative definition of adaptation and the rate of adaptation.

c. Hierarchical ϵ -machine reconstruction We investigated how to innovate new model classes from inadequate ones. Innovating new model classes had been worked out for the transition from finite-state machines to nested-stack automata, from stochastic finite-state machines to infinite-state stochastic machines, and from cellular automata to cellular transducers.[49] This subproject was reduced to coding up an algorithm that facilitated the innovation of hierarchical model classes. Building on the learning resource-interdependence work, we investigated the trade-offs between representation and modeling efficiency and the associated resource requirements for innovating new classes of representation.

d. Spectral reconstruction of causal architecture We developed a new method—the ϵ -Machine Spectral Reconstruction (ϵ MSR or “emissary”) algorithm—for inferring causal models from spectral (rather than temporal) data. We successfully applied this to experimental scattering spectra from complex materials.[80–83] In addition to working on the theoretical basis, we extended the applications beyond solid-state materials to spectral sensors for intelligent agents. Initial results on solid-state materials have been extensively expanded and compared with established models. They show a considerable improvement in both predictive value and ability to infer physical parameters.

3. Single-Agent Learning

We developed a theory that describes the consequences of an agent’s inability to model a structured environment: to the extent that the environment’s structural complexity (measured quantitatively) exceeds the representational or storage capacity of the learning agent, then the environment appears more random (in proportion to the difference).

a. Information Theory of Agent Learning We adapted our previous information-theoretic characterization of the structural complexity of nonlinear processes. The generality of the approach allowed us to analyze a number of learning and adaptation problems faced by a wide range of agent types and in a variety of environment—e.g., physical robots and software agents.[84]

b. Synchronizing to the Environment We developed a rigorous analysis of the process by which an agent comes to know the state of its environment.[85] In this work we introduced a new information theoretic quantity, which we called *transient information*, that measures the total uncertainty experienced while an agent becomes synchronized to its environment. This quantity is important, compared to existing information-theoretic measures of environmental complexity, in the situation when an agent must act immediately on incomplete information about its environment. We applied the general theory to a widely encountered problem in communications and control system design in multi-agent systems: how to synchronize most efficiently and noise-tolerantly to a periodic signal. We developed a general theory that should be easily applied to a number of specific problems in multi-agent communication and coordination.

4. Discovering Spatial Patterns

We made significant progress on computational mechanics for multidimensional information sources.[86, 87] This included extending the formal theory of epsilon-machines and causal states to arbitrary dimensional spaces and, also, developing and validating a reconstruction algorithm that can build spacetime epsilon-machines from 1+1 D dimensional data. This established our conjecture that computational mechanics can be extended to multidimensional data in a constructive manner. The results are useful for intelligent agents that must infer aspects of spatial structure in their environment.

We developed methods for inferring causal models from multidimensional spatiotemporal data. These consist of a suite of computational tools for automatically generating hypotheses concerning emergent structures in spatial systems and for automatically and efficiently testing such hypotheses. The result is a method to discover novel patterns, not seen before. We compared the methods to conventional ones of statistical mechanics in the case of dynamic spin systems. We reformulated the problem of identifying patterns in two and higher dimensions in a way that promises to overcome many prior limitations.[86]

We developed new algorithms for building the nonlinear filters (transducers) that are used for detecting embedded structures.[87] Having committed to one version of the previously developed formalism, we analyzed examples to see if it can be reduced to practice.

5. *Finite-Data Scaling of Pattern Discovery*

The relationship between a system’s information processing characteristics—e.g., its statistical complexity and entropy rate—determine in large measure the difficulty of learning a model of it. This difficulty is directly reflected in the compute time and space and in the amount of data required. Our goal was a general scaling theory of these trade-offs along the lines found in phase transition theory.[64, 88]

We extended our current methods to study the effect of finite resources on the optimality of reconstructed ϵ -machines. This is an important problem from several viewpoints. The first viewpoint concerns the practical problem of implementation efficiency. Given fixed amounts of time and computational resources (instruction set and storage), what level of approximation can the inference technique be expected to attain? The second viewpoint can be described in biological terms: organisms during their evolution and development are faced with exactly the same trade-off between modeling, predicting, and taking advantage of observed regularities in environment states. These trade-offs arise from limited metabolic resources and fixed capacities to perceive and model the environment. Analytic results on these trade-offs go some distance toward delineating the structural and information processing requirements for such adaptive organisms.

The main questions centered on how to balance the cost of additional machine components (causal states and transitions) against improved prediction and overfitting. At each stage the ϵ -machine’s structure should be the simplest possible so that the available data yields the best statistical predictions of environment behavior and so that the model has the highest probability of generalizing to new data. Large models, in contrast, spread the given data more thinly over a large number of internal parameters, thereby increasing the estimation error. Simultaneously, the machine’s ability to predict future observations should be maximized.

To investigate how these considerations are balanced, we measured the ϵ -machine’s structure via the statistical complexity. The prediction error was measured by the length of the machine-encoded data. This length was to be minimized along with the ϵ -machine complexity. Our optimality criterion was a version of various minimum modeling principles put forth by Rissanen,[89] Wallace,[90] and others.[91, 92] (This approach should not be confused, however, with one-pass data compression techniques, such as Ziv-Lempel coding,[93] in which, even for finite-state sources, the effective model is infinite.[89, 90])

B. Task 2: Emergence of Distributed Adaptation

In large-scale distributed systems the elementary components—agents, robots, neurons, ants, firms, and so on—behave as if they themselves are complicated dynamical processes. Many of their important and defining characteristics are the product of nonlinear interactions and emergent cooperative behavior. How do they individually and collectively process information? How do their dynamical behaviors and susceptibility to noise interact to support and limit their computational capabilities?

Our goal in approaching these questions, at the most basic level, was analogous to that found in Shannon’s pioneering work in developing a theory of “communication”.[73, 94]

Shannon explicitly disavowed using his mathematical theory to analyze the meaning, semantics, utility, and information content carried in communication channels. Instead, his theory considered only the amount of information and, due to this restriction, he was able to derive very basic constraints on reliable information transmission. The quantitative aspects of our computational mechanical approach stand in the same relation to distributed dynamical systems that Shannon’s theory did to the particular domains of communication that a channel can carry. At this level, we wished to understand the basic constraints between the (statistical) complexity of an agent-agent “communication channel” and the rate of reliable information processing of agent collectives. More importantly, though, the computational mechanics approach is distinct from Shannon’s original goals in that computational mechanics explicitly provides a procedure for detecting the effective “codes” or “protocols” that carry the significant information within a channel. Furthermore, it differs in that it also relates these codes to (i) known types of computation, (ii) state space structures that guide the dynamical behavior of the underlying system, and (iii) semantic and functional aspects of those structures.

1. ϵ -Machine Nets as Models of Agent Collectives

Our investigation of how individual agents and agent collectives perform computations involved analyzing simulations of related nonlinear dynamical systems and of a new type of ϵ -machine-based distributed model—called an ϵ -machine net.

Underlying the problems and themes noted above, one discerns a common theoretical model: finite-state or infinite-state stochastic ϵ -machines with input sensors and output effectors. We refer to these as ϵ -transducers. They are the mathematical articulation of the model implied by the computational-mechanical analysis of a single learning agent. In particular, ϵ -transducers are defined so as to capture all of the information processing and computing that is embedded in a learning agent. The goal, of course, is that individual agents be adaptive, in the senses discussed above. This is naturally modeled within the computational mechanics framework by including, in one form or another, the ϵ -machine reconstruction algorithm as part of the internal agent control dynamics. In this way, adaptive agents can be represented by ϵ -machine learners. As it reads its input stimulus, each ϵ -transducer learner changes structure according to a real-time dynamical version of ϵ -machine reconstruction, as described above. The net result is a purely computation- and information-theoretic model of the active, adaptive processing performed by an agent. Thus, agent collectives can be modeled as a network of interacting ϵ -transducers; a model that we refer to as an ϵ -machine net. This is a collection of ϵ -transducers with inputs and outputs connected in a fixed or possibly dynamic communication topology. The global, adaptive behavior that emerges from a given network architecture can be modeled in this way as a network of adaptive ϵ -machine learners.

2. *Finite-Population Dynamical Systems*

Collectives, in any practical setting, consist of finite populations of agents. Small, here, is measured relative to how sparsely a finite population coarse-grains the infinite-population state space. Finite population effects, which we have investigated previously in evolutionary and ecological dynamics,[95–98] dominate collective behavior. We showed that these effects can have significant, even counter-intuitive, dynamical consequences. Our studies of finite-population effects for general nonlinear dynamical systems demonstrated novel behaviors that arise from the interaction of the coarse-grained state space, sampling stochasticity, and nonlinear state-space structures (e.g., attractor types, stable-unstable manifolds, and so on).

3. *Evolution of Finite-Memory Agents*

An important type of collective is an evolving population of replicating agents that modify and transmit their internal control dynamics to others. This setting is useful for several reasons. First, one can use evolutionary methods to minimize the size of individual control systems or to optimize communication topologies. Since we were interested in the simplest real-time agents that are functional, typically only finite memory and finite time are available to implement a decision or control strategy. For this reason finite-memory ϵ -machines and ϵ -transducers are the most natural representation for this situation. Given this, the second reason an evolutionary setting is useful is that one way to find the smallest such agents is to use evolutionary search methods.

We had developed some expertise with evolving spatially extended (cellular) automata.[99–101] In this project we applied similar evolutionary methods, but to evolve finite-state machines.[102–107] We developed a theory of the evolutionary search for evolving ϵ -machines as we have done for our previous studies of evolving cellular automata and epochal evolutionary search dynamics.[108]

4. *Self-Organizing Agent Population Dynamics*

We completed our first projects in the emergence of structural complexity in groups of interacting finite-state agents—the *finitary process soup*. One result is that evolutionary pressures build structure hierarchically and prefer relatively less complex individuals that are generalists and can participated in a wider range of interactions. This result has important implications for the automated design of control systems.[109, 110]

We extended our work on self-organization in groups of interacting finite-state agents to spatially distributed agents. We showed that space adds a new and important kind of information storage and this, in turn, significantly affects the kinds of coordination and structure seen in agent populations.

5. *Theory and Analysis of Emergent Collective Adaptation*

In mathematical terms, agent collectives are high-dimensional stochastic dynamical systems. Though they may hold in principle, many of the basic techniques and concepts from dynamical systems theory do not directly apply to agent collectives and other complex systems. There are a number of difficulties that confront the application of nonlinear dynamics to high-dimensional systems generally, such as distinguishing nonstationary from stationary and transient from asymptotic behaviors. What does carry over from the study of low-dimensional dynamical systems is the fact that geometric structures in state space still determine the available information processing in a high-dimensional system.

How can distributed systems compute robustly? We feel there is some basic feature that has been missed to date that explains this capacity as found in neurobiological systems, as an important and instructive example. There, the observation that transient, not asymptotic, behavior determines the largest fraction of image classification and recognition in vision suggests that one would only need to roughly design the basin boundaries guiding the behavior of a neurobiological system in order to effect a given type of image processing. This comes at the cost of accuracy of the implementation—a trade-off that would have to be studied. In any case, we believe that this basic question can be elucidated by applying existing techniques to high-dimensional distributed dynamical systems and that the answer will delineate, at the same time, the functionality and robustness that agent collectives can manifest.

6. *Persistent Chaos in High Dimensions*

To explore general properties of collective adaptation, we completed the central theoretical results and simulation analysis on high-dimensional dynamical systems models of learning—recurrent neural networks.[111] We found that instability, in the form of deterministic chaos, is a common and persistent class of dynamical behavior. The implications for multiagent systems, which are necessarily high dimensional, is clear. It may be very difficult to design away instabilities in collective behavior.

7. *Dynamics of Multiagent Games*

We completed a new mathematical framework—MultiAgent Dynamical Systems (MADS)—for modeling the dynamics of multiagent games. We established new results on the global behavior of these high-dimensional dynamical systems, including sufficient conditions for stability and instability and limits on the achievable degrees of complexity and of emergent coordination.[112, 113] We extended this work to substantially more general methods to derive equations of motion from any type of local adaptation mechanism used by agents.

8. *Multiple Agents Servicing Multiple Tasks*

As a concrete example of this, we reformulated important multiagent systems (MAS) notions (“strategy”, “adaptation“, “cooperation”) in purely dynamical-systems terms. The goal was to describe and predict the emergence of spontaneous order in an MAS in terms of high-dimensional bifurcation theory. In particular, we developed a model for the TASK Open Experiment Framework (OEF) problem, what we called *Multiple Agents Servicing Multiple Tasks* (MASMT) and completed its basic theoretical description. We developed an interactive graphical interface to the simulation code that allows users to explore the system’s parameter space and its nonlinear dynamics. We also developed code and ran simulations to analyze the attractor-basin structure of the MASMT dynamical systems. The code streams for the interactive graphical interface to the MASMT and general multiagent dynamical systems (MADS) simulation codes were merged.

9. *Causal Synchrony*

a. Causal Synchrony in Multiagent Systems We developed ways for measuring and mapping causal synchrony in distributed agent systems. In particular, we collaborated with Kristina Lerman (TASK PI) and Aram Galstyan, of the University of Southern California’s Information Sciences Institute, to measure the spatiotemporal distribution of synchrony in the network minority game and determine its relation to the emergence of coordination and global efficiency.

b. Neural Causal Synchrony We applied computational mechanics to measure synchrony and coherence in heterogeneous multiagent networks. For the initial phase of developing and then testing these measures on well characterized network systems, we collaborated with a computational neuroscience group at the University of San Francisco on applications to neuronal networks. We examined extensive simulations of biophysically realistic networks.

C. **TASK Demonstration Project**

We developed an autonomous vehicle demonstration for the TASK Open Experiment Framework. This was done as a supplemental effort to the original Dynamics of Learning project. The goal was to demonstrate coordination and adaptation in a multiagent system (MAS) on a range of tasks specified in the Open Experiment Framework problem.

Specifically, we developed a robotic platform—the Robotic MultiAgent Development System (RoMADS)—that was used to experimentally validate our theoretical predictions, derived from computational mechanics, about learning, adaptation, and cooperation in multiagent systems. We constructed a collective of robotic vehicles and a high-efficiency database system that permits extensive, detailed analysis of their behavior. We developed a set of benchmark learning algorithms for individual robots and the robot collective including reinforcement learning algorithms.

1. MAS Dynamics Simulator

The goal here was to develop an interactive simulation of dynamical models of MAS coordination and adaptation. This consisted of the following components.

1. Automatically generate symbolic equations of motion from our reinforcement-learning MAS dynamics framework.
2. Easily integration of these into our differential equation solvers.
3. An interactive GUI that guides exploration and quantitative analysis of emergent behaviors.

2. Robotic MAS Development System (RoMADS)

The goal here was to develop a test bed for basic MAS behaviors using physical vehicles. RoMADS was developed to full functionality and was used for multiagent behavioral experiments and design. It included a video server and object-tracking system, web-based control and design system, and a high-performance database engine for data (especially video) storage and analysis.

The system consisted of several hardware and software components:

1. RoboCars: a collection of small robotic, wirelessly controlled vehicles;
2. RoboCar distributed control software;
3. The GROUNDS data logging and messaging system;
4. A Linux cluster for backend analysis and real-time simulation of learning algorithms; and
5. An WWW interface to all control and monitor functions.

Six RoboCars were made from standard Lego Mindstorms kits with additional sensing hardware. They were programmed using the LeJOS Java operating system for the Mindstorms “brick” (embedded controller). The interaction arena in which the cars ran was a 4’ x 8’ foot experiment table with predictable surface and walls.

The RoboCars had two differential DC drive motors that provided three speeds of forward and backward motion as well as turn-in-place direction changing. For the “native” sensors they had four bumpers to detect collisions and proximity to other cars, a light sensor used for locating other vehicles or stationary targets, distance and compass sensors that provided moderately accurate position estimates, and an internal battery voltage indicator to trigger recharging behavior.

The cars accepted simple motion commands and returned telemetry over their built-in infrared link which was connected to a RoboCar control system host running on an external PC level processor.

A RoboCar software module supplied basic intelligence for each of the RoboCars. The software, written in pure Java, allowed heterogeneous devices to be integrated into a MAS and then controlled and tracked in a common manner. It also provided hierarchical levels of control, sequencing, and adaptive behavior for each device, as well as access to all data structures for storage, analysis, and reloading. Using this system each device could be returned to any check-pointed state in the course of a set of experiments. Along with managing the live data, all commands and results were logged to local files or the external GROUNDS host for online monitoring and offline analysis.

The *General Robot Observations Using Nonlinear Dynamical Systems* (GROUNDS) system was hosted by a J2EE (Java 2 Enterprise Edition) system that was distributed across a number server platforms for load balancing and failover as necessary. The J2EE system supplied a container for Java Beans (user code modules), a small relational database (which integrates with any JDBC compliant DBMS), and a web server. Using this system all robot data was logged to the database and retrieved by external programs using standard protocols including HTTP (web), RMI (native Java), and CORBA (standard C++). GROUNDS also acted as a switching center, connecting external programs directly to control individual RoboCars. The switching center implemented higher-level communication and coordination behaviors between RoboCars and other devices in a controlled manner. As a side effect, it also was a straightforward matter to allow external servlets access to the RoboCars to run web based experiments. Thus, the entire system was made open to (authorized) use from anywhere on the net.

The Linux cluster was used to develop and execute the analysis and metrics software described above, for the RoMADS system, as well as the simulations.

The J2EE web server executed GROUNDS servlets that provide access to all monitoring and control functions over the Internet. This included a webcam viewing the experiment table, access to logged data, and direct access to the RoboCars as needed.

3. *RoMADS Sensor Nets*

The goal was to demonstrate sensor data integration techniques.

We integrated MICA Motes into the RoMADS systems to augment vehicle communication and sensory inputs, both on the RoboCars and in intelligent stationary agents. A Mote was yet another device to RoMADS where it logged data to GROUNDS. The data was then distributed to the RoboCars by external GROUNDS programs.

4. *Open Experiment Framework Related Experiments*

The goal here was to demonstrate multiagent behaviors with physical vehicles and sensor nets under the following guidelines.

1. Stationary and mobile targets: Tracking and servicing using adaptive algorithms.
2. Quantitative analysis: Collect data over many experiments to measure MAS efficiency, emergent intelligence (statistical complexity), and coordination (causal synchrony).

3. Develop mathematical dynamical models (MAS reinforcement learning framework) to predict experimentally observed MAS behaviors.
4. MAS experiments:
 - (a) Compare individual and group behaviors when mapping an area.
 - (b) Compare behaviors in simple tracking and servicing activities.
 - (c) Demonstrate robustness and adaptability to a change in the number of targets or agents.
 - (d) Provide a 2+1/2 D demonstration of system integration, using a second level on the experiment arena to simulate aerial surveillance of ground vehicles.

5. System Development Status

We evaluated the Intel/UC Berkeley MICA2 Motes and their Tiny OS software for integration into our distributed robotic system. These additions were geared to the TASK Demonstration Project.

We designed a new sensor I/O board for the MICA Mote controllers and fabricated and tested them. They were integrated into the RoboCars.

We designed and built a third generation of the RoboCar platform that relies completely on the MICA Mote controller. The new version was smaller, 5" diameter compared to the 8" diameter of the previous version. This allows doubling the number of robotic agents running simultaneously within the limited space of the experimental arena. Equally important, the run-time quadrupled from about 2 hours to 8 hours. Finally, the communication bandwidth was greatly increased from a single 9600 bps serially polled, shared channel to a 64 kps, packet-based CDMA channel. This substantially facilitated scaling RoMADS to more agents.

One remaining, but important task is to estimate how communication speeds scale with increasingly large numbers of vehicles sharing the radio channel. It turned out, for example, that the MICA2 controller communication speeds were not sufficiently high to support extrapolated future needs. This led us to upgrade to the MICAZ microcontroller which is based on ZigBee 250 kbps radio.

Along these lines, the PI established a collaboration with Sun Microsystems on its SunSpots initiative for sensor networks. We investigated adapting the hardware platform for our Multiagent Development System to the all-Java-based controllers that Sun Microsystems Labs has pioneered as the basis of a fourth generation of RoboCar. The benefits include seamless integration into the MADS (almost all-)Java-based system and another factor of 2 (estimated) in communication bandwidth. The main drawback is that this controller is still underdevelopment and not yet commercially available.

6. RoMADS Experiments on Multiagent Systems

At the TASK PI meetings we reported encouraging results on an algorithm for tabula rasa learning running on robotic vehicles. These showed that individual robots and

robot collectives can learn specific useful behaviors—such as, wall following and collision avoidance—*without these behaviors being programmed in by hand*.

We completed empirical tests of one of our tabula rasa learning algorithms on our experimental multiagent robotic platform. We experimented with learning in groups of two, three, and four robotic vehicles in experimental arenas of varying geometry. We quantitatively characterized the previously observed behaviors that (i) learning takes longer in groups due to the increased complexity of the statistics seen by each agent, (ii) nonetheless the robots eventually learn basic navigational skills, and (iii) they learn new behaviors that allow them to avoid crowding.

V. CONCLUSIONS

The main goals of this project were to understand the dynamics of learning in single agents well enough so that we could analyze and design agent collectives that perform a desired global task. Our approach to these goals was to extend the computational mechanics framework in a way that provided a dynamical model of learning and adaptation—a model that could be analyzed quantitatively and theoretically. We also used computational mechanics to develop methods to monitor the emergence of coordinated behavior in agent collectives and used these to test and verify global functionality, or lack of this functionality, and to compare distributed cooperation to alternative approaches to performing a given task.

The main advantages of the approach are its grounding in a rigorous theory of how nonlinear processes intrinsically compute (i.e., its grounding in computational mechanics) and the attention it pays to well developed theories of emergent collective behavior, such as found in pattern formation theory and statistical mechanics. We believe that the project laid the foundation for a new and novel view of agent collectives and of the issues involved in their design. The main disadvantages would be that in trying to develop predictive analytical theories for the behavior of learning and adaptive systems we may not be able to solve them in closed-form. In fact, we showed that their collective behaviors are extremely rich—exhibiting the panoply of dynamical behaviors that one now expects from large-scale, complex dynamical systems.

APPENDIX A: SPONSORED WORKSHOPS

A number of workshops were organized by the project as part of its research and outreach efforts.

1. The PI co-organized the founding workshop for his SFI Network Dynamics Program, “Structure and Dynamics in Complex Adaptive Networks”, Santa Fe, New Mexico, 10-12 August 2000.
2. The PI co-organized the SFI Business Network Meeting on “Network Dynamics”, Santa Fe, New Mexico, 22 March 2001.
3. The PI co-organized the workshop “The Internet as a Complex Adaptive System”, Santa Fe Institute, Santa Fe, New Mexico, 28-30 March 2001.
4. The PI and Stephanie Forrest (UNM, TASK PI) helped organize the April 2001 TASK PI Meeting in Santa Fe, New Mexico, and at SFI.
5. The PI organized the October 2002 TASK PI Meeting in Santa Fe, New Mexico, and at SFI. This included the PI and his group presenting a half-day tutorial on the Dynamics of Learning project and computational mechanics, statistical physics, and information-theoretic methods for multiagent systems.
6. The PI co-organized a successful workshop on “Collective Cognition: Mathematical Foundations of Distributed Intelligence,” bringing together workers in statistical physics, mathematical learning theory, computer science, economics, sociology, and other fields that are concerned with distributed information processing and cognition. (For a complete list of participants, talk titles and abstracts, and other information on the workshop, see <http://cse.ucdavis.edu/~dynlearn/colcog>.) The workshop addressed key TASK issues of distributed control, adaptation in heterogeneous environments, and spontaneous leadership.

APPENDIX B: PUBLICATIONS AND REPORTS

1. “When Evolution is Revolution—Origins of Innovation”, in *Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function*, James P. Crutchfield and P. Schuster, editors, Santa Fe Institute Series in the Science of Complexity, Oxford University Press, New York (2001) 104-138.
2. “Dynamics of Evolutionary Processes”, in *Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function*, James P. Crutchfield and P. Schuster, editors, Santa Fe Institute Series in the Science of Complexity, Oxford University Press, New York (2001) 1-22. (James P. Crutchfield and P. Schuster).

3. James P. Crutchfield and Peter Schuster, editors, "Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function", Oxford University Press, New York (2001).
4. James P. Crutchfield and David P. Feldman, "Synchronizing to the Environment: Information Theoretic Constraints on Agent Learning", *Advanced in Complex Systems* (2001) submitted. Santa Fe Institute Working Paper 01-03-020 and arXiv.org/abs/nlin.AO/0103038.
5. Wim Hordijk, Cosma Rohilla Shalizi and James P. Crutchfield, "An Upper Bound on the Products of Particle Interactions in Cellular Automata," *Physica D* 154 (2001): 240-258.
6. Cosma Rohilla Shalizi and James P. Crutchfield, "Computational Mechanics: Pattern and Prediction, Structure and Simplicity," *Journal of Statistical Physics* 104 (2001): 819-881.
7. Yuzuru Sato, Eizo Akiyama and J. Doyne Farmer, "Chaos in Learning a Simple Two Person Game," *Proceedings of the National Academy of Sciences USA* **99**:7 (2002) 4748-4751.
8. Cosma R. Shalizi and James P. Crutchfield, "Information Bottlenecks, Causal States, and Statistical Relevance Bases: How to Represent Relevant Information in Memory-less Transduction", *Advances in Complex Systems* **5**:1 (2002) 1-5.
9. Yuzuru Sato and James P. Crutchfield, "Coupled Replicator Equations for the Dynamics of Learning in Multiagent Systems", *Physical Review E* **67**:1 (2003) 40-43.
10. Dowman P. Varn, Geoff S. Canright, and James P. Crutchfield, "Discovering Planar Disorder in Close-Packed Structures from X-Ray Diffraction: Beyond the Fault Model", *Physical Review B* **66**:17 (2002) 156-159.
11. David P. Feldman and James P. Crutchfield, "Regularities Unseen, Randomness Observed: Levels of Entropy Convergence", *CHAOS* **13**:1 (March 2003) 25-54.
12. Francisco Jimenez-Morales, Melanie Mitchell, and James P. Crutchfield, "Evolving one-dimensional cellular automata to perform a non-trivial collective behavior task: One case study", *Computational Science-ICCS 2002, Part I, Proceedings* 2329 (2003) 793-802.
13. Cosma R. Shalizi, Kristina Lisa Shalizi, and James P. Crutchfield, "Pattern Discovery in Time Series, Part I: Theory, Algorithm, Analysis, and Convergence, Santa Fe Institute Working Paper 02-10-060; arXiv.org/abs/cs.LG/0210025 (2002).
14. Kristina Lisa Shalizi, Cosma R. Shalizi and James P. Crutchfield, "Pattern Discovery in Time Series, Part II: Implementation, Evaluation, and Comparison", <http://cse.ucdavis.edu/~cmg/CompMech/papers/pditsii.html> (2002).

15. Dowman P. Varn, Geoff S. Canright, and James P. Crutchfield, “Inferring Pattern and Disorder in Close-Packed Structures from X-ray Diffraction Studies, Part I: epsilon-Machine Spectral Reconstruction Theory”, submitted to Physical Review B; Santa Fe Institute Working Paper 03-03-021; [arXiv.org/abs/cond-mat/0302585](http://arxiv.org/abs/cond-mat/0302585) (2002).
16. Dowman P. Varn, Geoff S. Canright, and James P. Crutchfield, “Inferring Pattern and Disorder in Close-Packed Structures from X-ray Diffraction Studies, Part II: Structure and Intrinsic Computation in Zinc Sulphide”, submitted to Physical Review B; Santa Fe Institute Working Paper 03-03-022; [arXiv.org/abs/cond-mat/0302587](http://arxiv.org/abs/cond-mat/0302587) (2002).
17. Cosma R. Shalizi and David J. Albers, “Symbolic Dynamics for Discrete Adaptive Games.” SFI Working Paper 02-07-031. Submitted to Physical Review E. <http://arxiv.org/abs/cond-mat/0207407>.
18. James P. Crutchfield and Peter Schuster, “Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function”, Santa Fe Institute Series in the Sciences of Complexity, Oxford University Press, New York (2003).
19. David P. Feldman and James P. Crutchfield, “Structural Information in Two-Dimensional Patterns: Entropy Convergence and Excess Entropy”, Physical Review E **67** (2003) 051103.
20. David P. Feldman and James P. Crutchfield, “Synchronizing to Periodicity: The Transient Information and Synchronization Time of Periodic Sequences”, Advances in Complex Systems **7**:3-4 (2002) 329–355.
21. Michael Schippling, “An Overview of Distributed Robotics”, cse.ucdavis.edu/~dynlearn/RoMADS/RobishLiteratureReview.html.
22. D. P. Varn and James P. Crutchfield, “From Finite to Infinite Range Order via Annealing: The Causal Architecture of Deformation Faulting in Annealed Close-Packed Crystals”, Physical Letters A **234**:4 (2004) 299–307.
23. N. Ay and James P. Crutchfield, “Reductions of Hidden Information Sources”, Journal of Statistical Physics **210**:3-4 (2005) 659-684.
24. James P. Crutchfield and O. Gernerup, “Objects that Make Objects: The Population Dynamics of Structural Complexity”, Journal of the Royal Society Interface **3** (2005) 345-349. Santa Fe Institute Working Paper 04-06-020; [arxiv.org: nlin.AO/0406058](http://arxiv.org/nlin.AO/0406058).
25. Dave Albers and Clint Sprott, “Structural Stability and Partial Hyperbolicity in Large Dynamical Systems”, Physical Review E (2004) submitted.
26. Erik Talvitie, “Inferring the Structure of Space”, SFI REU Project Technical Report (2004).
27. Selah Lynch, “Multisymbol epsilon-Machine Reconstruction”, SFI REU Project Technical Report (2004).

28. John Albers, “Multiple Agent Dynamical Systems”, SFI REU Project Technical Report (2004).
29. David Feldman and James P. Crutchfield, “Transients in Periodic Systems”, *Advances in Complex Systems* **7**:3-4 (2004) 329-355.
30. Carl S. McTague and James P. Crutchfield, “Automated Pattern Detection—An Algorithm for Constructing Optimally Synchronizing Multi-Regular Language Filters”, *Theoretical Computer Science* (2006) in press. Santa Fe Institute Working Paper 04-09-027. [arxiv.org e-print cs.CV/0410017](http://arxiv.org/e-print/cs.CV/0410017).
31. David Albers, Clint Sprott, and James P. Crutchfield, “Persistent Chaos in High Dimensions”, *Physics Review Letters* (2005) submitted.

APPENDIX C: PRESENTATIONS

1. James P. Crutchfield, “The Evolutionary Unfolding of Complexity”, Colloquium, Computer Science Department, University of Western Ontario, London, Ontario, Canada.
2. James P. Crutchfield, “Intrinsic Computation”, Stanford University, Stanford, California.
3. James P. Crutchfield, “Complexity”, Computational Economics Summer School, Santa Fe Institute.
4. James P. Crutchfield, “Intrinsic Computation”, Sociology Department, Columbia University, New York, New York.
5. James P. Crutchfield, “Dynamics of Learning and Distributed Adaptation”, TASK KickOff Meeting, Charleston, South Carolina. 3 October 2000.
6. James P. Crutchfield, “Objets d’Bits: Thinking about Emergent Structures in Cellular Automata”, Center for the Study of Complex Systems, University of Michigan, Ann Arbor, 27 October 2000.
7. James P. Crutchfield, “Objets d’Bits: Thinking about Emergent Structures in Cellular Automata”, Center for Studies in Biology and Physics, Rockefeller University, New York, 7 November 2000.
8. James P. Crutchfield, “Objets d’Bits: Thinking about Emergent Structures in Cellular Automata”, Mathematics Department, California State University, Northridge, Los Angeles, 20 November 2000.
9. Cosma R. Shalizi, “ ϵ -Transducers: Computational Mechanics of History-Dependent Transduction”, SFI/Chinese Academy of Sciences (PRC) Working Group Meeting, Santa Fe, 15 August 2000.

10. James P. Crutchfield, "Pattern Discovery", Hewlett-Packard Research Laboratories, Palo Alto, California, 8 March 2001.
11. James P. Crutchfield, "Function and Robustness", Packard Program on Robustness, Founding Workshop, Santa Fe Institute, Santa Fe, New Mexico, 16-17 March 2001.
12. James P. Crutchfield, "Network Dynamics", SFI Business Network Meeting, Santa Fe, New Mexico, 22 March 2001.
13. James P. Crutchfield, "Network Dynamics", Workshop on The Internet as a Complex Adaptive System, Santa Fe Institute, Santa Fe, New Mexico, 29 March 2001.
14. Cosma R. Shalizi, "Spatiotemporal Emergent Structures from Causal Architecture," Center for the Study of Complex Systems, University of Michigan-Ann Arbor, 22 March 2001.
15. James P. Crutchfield and Cosma R. Shalizi, "Dynamics of Learning and the Emergence of Distributed Adaptation", Site Visit by Rome AFB (Jamie Lawton and Robert Paragi) Santa Fe Institute, Santa Fe, New Mexico, 16 April 2001.
16. James P. Crutchfield, "Synchronizing to the Environment", DARPA TASK PI Meeting, Santa Fe, New Mexico, 17-19 April 2001.
17. James P. Crutchfield, "Causality and Pattern Discovery", Workshop on Determinism, Max-Planck Gessellschaft, Ringberg Castle, Tegernsee, Bavaria, Germany, 4-8 June 2001.
18. Cosma R. Shalizi, "Computational Mechanics and Pattern Discovery", SFI/Max Planck Institute for Mathematics in the Sciences Joint Workshop on "Complexity Science in Eastern Europe/Complexity: Unifying Themes for the Sciences and New Frontiers for Mathematics", Leipzig, Germany, 14-18 May 2001.
19. Cosma R. Shalizi, "Foundations of Complex Systems—Probability, Statistics, and Networks", SFI Complex Systems Summer School, June–July 2001, Santa Fe, NM.
20. James P. Crutchfield, "New Algorithms for Adaptive Learning", DARPA TASK CAHDE REF Meeting, Media Lab, MIT, 27–28 August 2001.
21. James P. Crutchfield, "Causality and Pattern Discovery", DisIntegrative Themes meeting, Santa Fe Institute, 23-27 July 2001.
22. Yuzuru Sato, "Dynamics of learning and coupled replicator equations," presentation, Autumn 2001 meeting of the Physical Society of Japan.
23. Cosma R. Shalizi, "Pattern Discovery in Networks," presentation, 18 July 2001, Air Force Office of Scientific Research workshop "Infospherics: Science for Building Large-scale Global Information Systems," George Mason University, Fairfax, Virginia, 17-19 July 2001.

24. Cosma R. Shalizi, "Pattern Discovery Techniques for Social Science," SFI Computational Economics Summer School, Santa Fe, 20 July 2001.
25. James P. Crutchfield, "Pattern Discovery", Rockefeller Research Center, Bellagio, Italy, 26 October 2001.
26. Dowman P. Varn, "Beyond the Fault Model", poster presentation, Dynamics Days, Baltimore, 4-7 January 2002.
27. James P. Crutchfield and Cosma Rohilla Shalizi, "Spontaneous Leadership", DARPA/TASK meeting in Washington, D.C., 9-11 January 2002.
28. James P. Crutchfield, "Collective Cognition," presentation at Collective Cognition workshop, 21-25 January 2002.
29. Cosma R. Shalizi, "Causal Architecture of Collectives," presentation at Collective Cognition workshop, 21-25 January 2002.
30. James P. Crutchfield, "What Does it Mean to Compute? Origins of Intrinsic Computation", invited presentation at AAAS meeting in Boston, 14-19 February 2002.
31. James P. Crutchfield, "What Does it Mean to Compute? Origins of Intrinsic Computation", seminar at Dartmouth, 18 February 2002.
32. Cosma R. Shalizi, "Causal Synchrony," seminar, Physics Department, University of San Francisco, 26 February 2002.
33. James P. Crutchfield, "Theory Theory", SFI Science Board Symposium, Santa Fe, 4 May 2002.
34. James P. Crutchfield, "Causal Synchrony", Columbia workshop on Networks and Contagion, Columbia University, 23-25 May 2002.
35. James P. Crutchfield, "Causal Synchrony", Czech Academy of Sciences Workshop on Evolutionary Innovation, Prague, 27-31 May 2002.
36. James P. Crutchfield, "What Is Complexity?", SFI Complex Systems Summer School, Santa Fe, 11-12 June, 2002.
37. James P. Crutchfield, "Emergent Coordination in Multi-Agent Systems", TASK PI Meeting, Chicago, 19-21 June 2002.
38. Cosma R. Shalizi, "Causal Synchrony in Networks," poster, SFI Science Board meeting, 3 May 2002.
39. Cosma R. Shalizi, "Network Synchrony and Network Architecture," SFI Business Network meeting on Networks and Supply Chains, 20 May 2002.

40. Cosma R. Shalizi, “Foundational Issues in Complex Systems,” SFI Complex Systems Summer School, 11, 14, and 17 June 2002.
41. Yuzuru Sato, “Dynamics of coupled replicator equations”, Dynamics Days 2002 Europe, Heidelberg, July 2002.
42. James P. Crutchfield delivered a lecture course on computational mechanics to the SFI Complex Systems Summer School, 12-28 July 2002, Budapest, Hungary.
43. James P. Crutchfield participated in the final National Academy of Sciences Computer Science and Technology Board’s Committee on Information Technology and Creativity meeting, 29-3 August 2002, Montreal, Canada.
44. Carl McTague, “Cellular Automata Computational Mechanics: Automated Pattern Discovery”, Dynamics of Learning group seminar, 10 September 2002.
45. Dave Feldman, “Information Theory in Two and Higher Dimensions”, Dynamics of Learning Group seminar, 17 August 2002.
46. James P. Crutchfield tutorial on “Computational Mechanics for Multiagent Systems” for TASK working group on 9 October 2002.
47. David Feldman presented a tutorial on “Information Theory for Multiagent Systems” for TASK working group on 9 October 2002.
48. Michael Schippling, “An Overview of Distributed Robotics”, Dynamics of Learning Group seminar, 29 October 2002; see presentation materials on Dynamics of Learning website (<http://cse.ucdavis.edu/~dynlearn>).
49. David Albers, “The Dynamics of Learning in Multiagent Systems”, Chaos and Complex Systems Group, University of Wisconsin, Madison, 14 October 2002.
50. David Albers, “Structural Stability in High-Dimensional Dynamical Systems”, Chaos and Complex Systems Group, University of Wisconsin, Madison, 15 April 2003 and to the Dynamics of Learning Group, Santa Fe Institute, 14 October 2002.
51. James P. Crutchfield, “Multiple Agents Servicing Multiple Tasks”, DARPA TASK PI Meeting, 19-20 January 2003.
52. James P. Crutchfield, “Patterns and Pattern Discovery—Lecture Course”, Centro de Ricerca Matematica Ennio de Giorgi, Scuola Normale Superiore, University of Pisa, Pisa, Italy, 24-27 February 2003.
53. Michael Schippling, “Adaptively Learning Causal Models by Distributed Robots”, Dynamics of Learning Group seminar. 25 March 2003.
54. James P. Crutchfield, “Is Anything Ever New?”, Rackham Summer Interdisciplinary Institute, University of Michigan, Ann Arbor, 3 April 2003.

55. James P. Crutchfield, "Pattern from Process", Zero-One Spring Lecture Series, Microsoft Corporate Headquarters, Mountain View, 20 May 2003.
56. Dave Albers and Yuzuru Sato, "Multiagent Dynamical Systems: Review and Application to the Multiple Agents Servicing Multiple Tasks Problem," Dynamics of Learning Seminar, 27 May 2003, Santa Fe Institute.
57. James P. Crutchfield, "Dynamics of Multiagent Systems", BISON Consortium Meeting, Telenor, Oslo, Norway, 3 June 2003.
58. James P. Crutchfield, "An Introduction to Computational Mechanics", SFI Complex System Summer School, St. Johns College, Santa Fe, New Mexico, 11-13 June 2003.
59. Yuzuru Sato, "Dominance of Milnor Attractors in Globally Coupled Dynamical Systems", Dynamics of Learning Seminar, 15 July 2003.
60. Dave Albers, "Structural Stability and Partial Hyperbolicity in Large Dynamical Systems", Dynamics of Learning Seminar, 22 July 2003.
61. Erik Talvitie, "Problems in Inferring the Structure of Space", Dynamics of Learning Seminar, 29 July 2003; SFI REU Talk, 8 August 2003.
62. Selah Lynch, "Algorithms for Multisymbol epsilon-Machine Reconstruction", Dynamics of Learning Seminar, 12 August 2003; SFI REU Talk, 13 August 2003.
63. John Albers, "Multiple Agents Servicing Multiple Tasks", Dynamics of Learning Seminar, 26 August 2003; SFI REU Talk, 27 August 2003.
64. Olof Gernerup, "Finitary Process Soup", Dynamics of Learning Seminar, 2 September 2003, Santa Fe Institute.
65. James P. Crutchfield, "Automated Pattern Discovery", Invited Plenary Lecture, SIAM Meeting on Industrial Applications of Complex Systems, Toronto, 13-15 October 2003.
66. James P. Crutchfield, "Dynamical Embodiments of Computation in Cognition", Ron Brachman (IPTO Director) et al visit, Santa Fe Institute, Santa Fe, New Mexico, 22-23 October 2003.
67. James P. Crutchfield, "Complex Systems Theory?", Invited Lecture, SFI Faculty Retreat, Bishop's Lodge, Santa Fe, New Mexico, 24-25 October 2003.
68. James P. Crutchfield, "Intrinsic Computation", Physics Department Colloquium, University of California, Davis, 12 November 2003.
69. James P. Crutchfield, "Automated Pattern Discovery", Engineering Department Seminar, University of California, Davis, 13 November 2003.

70. James P. Crutchfield, “Form and Function: From pattern to semantics and on to function”, Collegium Budapest-SFI Workshop on Form and Function, Budapest, Hungary, 16-19 November 2003.
71. James P. Crutchfield, “Dynamics of Learning in Distributed Robotics”, TASK PI Meeting, Washington DC, 4-5 December 2003.
72. James P. Crutchfield, “Intrinsic Computation and Pattern Discovery”, Colloquium, Physics Department, University of New Mexico, Albuquerque, 12 January 2004.
73. Nihat Ay (SFI), “Information Geometry, Part I”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 20 January 2004.
74. Nihat Ay (SFI), “Information Geometry, Part II”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 20 January 2004.
75. Yuzuru Sato (SFI), “Multiagent Dynamical Systems and Learning”, Arizona Days workshop, Los Alamos National Laboratory, 30 January 2004.
76. Chris Strelhoff (SFI/UIUC), “Statistical Error Estimates for Markov Models of Dynamical Systems”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 10 February 2004.
77. Yuzuru Sato (SFI), “Multiagent Dynamical Systems and Learning”, Research Seminar, Santa Fe Institute, 17 February 2004.
78. Dave Albers (SFI/Madison), “Multiple Agents Servicing Multiple Tasks, Part I”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 24 February 2004.
79. Nihat Ay (SFI), “Causal State Reduction of Complex Processes”, Research Seminar, Santa Fe Institute, 25 February 2004.
80. James P. Crutchfield, “Intrinsic Computation and Nanotechnology”, invited lecture, Conference on Nanotechnology, Biotechnology, Information Technology, and Cognitive Science, New York, 26 February 2004.
81. Dave Albers (SFI/Madison), “Multiple Agents Servicing Multiple Tasks, Part II”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 2 March 2004.
82. Nihat Ay (SFI), “Information Geometry, Part III”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 23 March 2004.
83. James P. Crutchfield, “Theory Theory”, invited lecture, Seminar on Computation in Natural Systems, Center for Nonlinear Studies, Los Alamos National Laboratory, 29 March 2004.
84. Nihat Ay (SFI), “Information Geometry, Part IV”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 30 March 2004.

85. James P. Crutchfield (SFI), “Practical Computational Mechanics”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 6 April 2004.
86. James P. Crutchfield, “Pattern Discovery and Automated Theory Building”, Physics-Astronomy Colloquium, Northwestern University, Chicago, Illinois, 9 April 2004.
87. James P. Crutchfield (SFI), “Complex Systems Theory?”, Science Board Spring Meeting, Santa Fe Institute, 7 May 2004.
88. Olof Gornerup (SFI), “The Emergence of Hierarchical Structures in Evolution: The Finitary Process Soup”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 1 June 2004.
89. James P. Crutchfield, Dynamics of Learning Group (SFI), “Overview of the Dynamics of Learning: Introductions to the REU Summer Interns”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 8 June 2004.
90. Dave Albers (SFI/Madison), “Outstanding Problems in the Statistical Mechanics of Nonequilibrium”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 22 June 2004.
91. John Albers (SFI/Madison), “Adjoint Sensitivity Analysis”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 29 June 2004.
92. James P. Crutchfield, “MultiAgent Dynamical Systems: Completion of the Theoretical Framework and Simulator”, TASK PI Meeting, Washington, DC, 5 August 2004.
93. James P. Crutchfield and Michael Schippling, “RoMADS—Robotic MultiAgent Development System” and “MultiAgent Dynamical Systems”, TASK Demonstration, Washington, DC, 4 August 2004.
94. David Albers, “Robust Chaos in High-Dimensional Dynamical Systems”, Ph.D. Dissertation Defense, Physics Department, University of Wisconsin, Madison, 18 August 2004.
95. Carl McTague (SFI), “Automated Pattern Discovery—Constructing Multi-Regular Language Filters”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 17 August 2004.
96. James P. Crutchfield, “Pattern and Pattern Discovery—A Review of Computation Mechanics”, Redwood Neurosciences Institute, Menlo Park, CA, 2 September 2004.
97. James P. Crutchfield, “Cellular Automata Computational Mechanics—Patterns and Discovery”, Keynote speech, ACRI 2004, Amsterdam, Netherlands, 26 October 2004.
98. James P. Crutchfield, “The Evolution of Structural Complexity”, Center for Living Technology and Statistics Department, Universite de Ca Forsca, Venice, Italy, 2 November 2004.

99. James P. Crutchfield, “Chaos and Complexity”, Graduate Seminar, Venice International Program of the University of Virginia, Venice, Italy, 4 November 2004.
100. James P. Crutchfield, “Objects that Make Objects: The Population Dynamics of Structural Complexity”, Mathematical Biology Seminar, UC Davis, 15 February 2005.
101. James P. Crutchfield, “Pattern Discovery for Spatial Processes”, Sun Microsystems Laboratory 2005 Open House, Computer Museum, Palo Alto, California, 27 April 2005.
102. James P. Crutchfield, “Multiagent Dynamical Systems”, Seminar Series in Complex Systems, Environmental Sciences Department, UC Davis, 2 June 2005.
103. James P. Crutchfield, “Multiagent Dynamical Systems”, Sun Microsystems Research Laboratories, Mountain View, California, 17 August 2005.
104. James P. Crutchfield, “Multiagent Dynamical Systems”, Evolution and Economics Seminar, University of California, Davis, 25 October 2005.
105. James P. Crutchfield, Keynote speech on “Frontiers in Complex Systems”, Hackers 2005, Santa Cruz, California, 13 November 2005.
106. James P. Crutchfield, “Objects that Make Objects: The Population Dynamics of Structural Complexity”, Workshop on the Evolution of Biological Complexity, Ohio State University, Columbus, Ohio, 17 November 2005.
107. James P. Crutchfield, “Structure, Meaning, and Function: A Dynamical Systems Perspective”, International School on Semiotic Dynamics, Language, and Complexity, Ettore Majorana Foundation and Centre for Scientific Culture Erice, Italy, 14 December 2005.

APPENDIX D: PRESENTATIONS HOSTED BY THE PROJECT

1. Martin Nilsson (LANL), “Time Series Modeling and Sensor Fusion”, Dynamics of Learning Group seminar, 24 September 2002.
2. Lukasz Debowski, “Infinitary Processes and Linguistics”, Dynamics of Learning Group seminar, 6 August 2002.
3. Scott and Jena Page (Economics, University of Michigan), “Problem Solving in Groups”, Dynamics of Learning Group seminar, 16 July 2002.
4. Eric Smith, “Self-Organization as Structural Refrigeration”, Dynamics of Learning group seminar, 19 and 26 November 2002.
5. John Miller (Carnegie-Mellon), “Coordination and Competition”, Dynamics of Learning Group seminar, 18 March 2003.

6. Sam Bowles (Amherst/SFI) and Jessica Flack “Social Cooperation in Chimps and Economies), Dynamics of Learning Group seminar, 21 January 2003.
7. Koichi Fujimoto, “How Fast Elements can Affect Slow Dynamics”, Dynamics of Learning Seminar, 8 July 2003.
8. Matt Tanner and Eric Smith, “How do Smart Agents Use Their Internal Knowledge for Effective Group Interactions?”, Dynamics of Learning Seminar, 5 August 2003.
9. Geoff Canright and Kenth Engo, “Biologically Inspired Self-Organization in Dynamical Networks”, Dynamics of Learning Seminar, 22 August 2003.
10. Chi-Chi May, “Biological Coding Theory”, Dynamics of Learning Seminar, 30 September 2003.
11. Emily Stone (Utah State), “Dynamic models of PCR”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 5 March 2004.
12. Eizo Akiyama (Tokyo Institute of Technology), “Game Dynamical Systems”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 9 March 2004.
13. Garrett Kenyon (LANL), “The Neurodynamics of Pattern Recognition in Early Vision”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 13 April 2004.
14. David Peak (Utah State), “Distributed Computation in Plant Respiration”, Colloquium, Santa Fe Institute, 16 April 2004.
15. Tony Bell (Redwood Neuroscience Institute), “Topics in the Dynamics of Neurolearning”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 27 April 2004.
16. James Theiler (LANL), “Machine Learning Approaches to Modeling Dynamics Systems”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 5 May 2004.
17. Luis Bettencourt (LANL), “In vitro Neurodynamics and Distributed Biosensors”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 11 May 2004.
18. Frantisek Matus (University of Prague), “Information Geometry and Inference”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 25 May 2004.
19. Ken Burns (Northwestern), “Dynamical Systems Tools in High Dimensions”, Dynamics of Learning Group, Seminar, Santa Fe Institute, 24 August 2004.

REFERENCES

- [1] J. Hendler. Is there an intelligent agent in your future? <http://www.nature.com/nature/webmatters/agents/agents.html>, 2000.

- [2] G. Weiss. *Distributed Artificial Intelligence meets Machine Learning*. Springer-Verlag, New York, 1997.
- [3] P. Stone and M. Velos. Multiagent systems: A survey from a machine learning perspective. Technical Report CMU-CS-97-193, Carnegie Mellon University, 1997.
- [4] N. R. Jennings, K. Sycara, and M. Wooldridge. A roadmap of agent research and development. *Autonomous Agents and Multi-Agent Systems*, 1:7–38, 1998.
- [5] S. Sen. *Multi-Agent Learning: Papers from the 1997 AAAI Workshop*. AAAI Press, Menlo Park, California, 1997. Technical Report WS-97-03.
- [6] K. Sycara. Multiagent systems. *AI Magazine*, 19:79–92, 1998.
- [7] D. H. Wolpert. An overview of collective intelligence. In J. M. Bradshaw, editor, *Handbook of Agent Technology*, Menlo Park, California, 1988. AAAI Press.
- [8] D. H. Wolpert and K. Tumer. An introduction to collective intelligence. Technical Report cs.MA/9908014, arXiv.org, 1999.
- [9] R. Axelrod. *The Complexity of Cooperation: Agent-Based Models of Competition and Cooperation*. Princeton University Press, Princeton, New Jersey, 1997.
- [10] J. M. Epstein and R. Axtell. *Growing Artificial Societies: Social Sciences from the Ground Up*. MIT Press, Cambridge, Massachusetts, 1996.
- [11] M. E. Huhn. *Distributed Artificial Intelligence*. Pittman, London, 1987.
- [12] L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- [13] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, Massachusetts, 1998.
- [14] H. Peyton Young. *Individual strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, Princeton, New Jersey, 1998.
- [15] J. M. Vidal and E. H. Durfee. Predicting the expected behavior of agents that learn about agents: The clri framework. Technical Report cs.MA/0001008, arXiv.org, 2000.
- [16] W. B. Arthur. Complexity in economic theory: Inductive reasoning and bounded rationality. *Economic Review*, 84:406–411, 1994.
- [17] W. B. Arthur, S. Durlauf, and D. Lane. *The Economy as an Evolving Complex System II*. Santa Fe Institute Series in the Sciences of Complexity. Addison-Wesley, Reading, Massachusetts, 1997.
- [18] E. Bonabeau, F. Henaux, S. Guerin, D. Synders, P. Kuntz, and G. Theraulaz. Routing in telecommunications networks with “smart” ant-like agents. Technical Report 98-01-003, Santa Fe Institute, 1998.
- [19] R. A. Brooks. Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on AI*, pages 569–595, 1991.
- [20] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
- [21] D. Carmel and S. Markovich. Exploration and adaptation in multiagents systems: A model-based approach. In *Proceedings of the International Joint Conference on AI*, 1997. [http://www.cs.technion.ac.il/users/wwwb/cgi-bin/tr-info.cgi?1997/CIS/CI%\\$9704](http://www.cs.technion.ac.il/users/wwwb/cgi-bin/tr-info.cgi?1997/CIS/CI%$9704).
- [22] A. R. Cassandra, L. P. Kaelbling, and M. L. Littman. Acting optimally in partially observable stochastic domains. In *Proceedings of the Twelfth International Joint Conference on AI*, 1994.

- [23] A. Chavez and P. Maes. An agent marketplace for buying and selling goods. In *First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, pages 75–90, 1996. <http://ecommerce.media.mit.edu>.
- [24] C. Claus and C. Boutilier. The dynamics of reinforcement learning in cooperative multiagent systems. In *Proceedings of the Workshop on Multiagent Learning*, pages 746–752, Menlo Park, California, 1998. AAAI Press. <http://www.cs.ubc.ca/spider/cebly/Papers/multirl.ps>.
- [25] M. Dorigo and L. M. Gambardella. Ant colonies for the travelling salesman problem. *Biosystems*, 39, 1997.
- [26] M. Dorigo and L. M. Gambardella. Ant colony systems: A cooperative learning approach to the travelling salesman problem. *IEEE Trans. Evolutionary Computation*, 1:53–66, 1997.
- [27] J. H. Holland and J. H. Miller. Artificial adaptive agents in economic theory. *American Economic Review*, 81:365–370, 1991.
- [28] J. Hu and M. P. Wellman. Multiagent reinforcement learning: Theoretical framework and an algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning*, pages 242–250, 1998.
- [29] J. Hu and M. P. Wellman. Online learning about other agents in a dynamic multiagent system. In *Proceedings of the Second International Conference on Autonomous Agents*, pages 239–246, 1998.
- [30] A. Hubler and D. Pines. Prediction and adaptation in an evolving chaotic environment. In G. Cowan, D. Pines, and D. Meltzer, editors, *Complexity: Metaphors, Models and Reality*, Santa Fe Institute Series in the Sciences of Complexity, pages 343–379, Reading, Massachusetts, 1994. Addison-Wesley.
- [31] J. O. Kephart, T. Hogg, and B. Huberman. Collective behavior of predictive agents. *Physica D*, 42:48–65, 1990.
- [32] V. R. Lesser and D. D. Corkill. Functionally accurate, cooperative distributed systems. *IEEE Trans. Systems, Man, and Cybernetics*, 11:81–96, 1981.
- [33] M. L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the Eleventh International Conference on Machine Learning*, pages 157–163. Morgan Kaufman, 1994. <http://www.cs.duke.edu/~mlittman/docs/ml94-final.ps>.
- [34] D. H. Wolpert, K. R. Wheeler, and K. Tumer. General principles of learning-based multi-agent systems. In *Proceedings of the Third International Conference of Autonomous Agents*, pages 77–83, 1999.
- [35] C. Boutilier, Y. Shoham, and M. P. Wellman. Editorial: Economic principles of multi-agent systems. *Artificial Intelligence Journal*, 94:1–6, 1997.
- [36] J. M. Bradshaw. *Software Agents*. MIT Press, Cambridge, Massachusetts, 1997.
- [37] D. Challet and Y. C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407, 1997.
- [38] D. Challet and Y. C. Zhang. On the minority game: Analytical and numerical studies. *Physica A*, 256:514, 1998.
- [39] B. A. Huberman and T. Hogg. The behavior of computational ecologies. In *The Ecology of Computation*, pages 77–115, Amsterdam, 1988. North Holland.
- [40] J. O. Kephart, J. E. Hanson, and J. Sairamesh. Price and niche wars in a free market economy

- of software agents. *Artificial Life*, 4:1–13, 1998.
- [41] D. McFarland. Toward robotic cooperation. In *Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 440–443, Cambridge, Massachusetts, 1994. MIT Press.
 - [42] J. H. Miller. The coevolution of automata in the repeated prisoner’s dilemma. *Journal of Economic Behavior and Organization*, 29:87–112, 1996.
 - [43] A. Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoner’s dilemma. *Economics Letters*, 19:227–230, 1985.
 - [44] S. I. Nishimura and T. Ikegami. Emergence of collective strategies in a prey-predator game model. *Artificial Life*, 3:243–260, 1997.
 - [45] S. Olafsson. Games on networks. *Proceedings of the IEEE*, 85:1556–1562, 1997.
 - [46] A. Schaerf, Y. Shoham, and M. Tennenholtz. Adaptive load balancing: A study in multi-agent learning. *Journal of Artificial Intelligence Research*, 94:99–137, 1997.
 - [47] T. Sandholm and V. R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94:99–137, 1997.
 - [48] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Let.*, 63:105–108, 1989.
 - [49] J. P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54, 1994.
 - [50] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, Cambridge, Massachusetts, 1998.
 - [51] M. J. Kearns and U. V. Vazirani. *An Introduction to Computational Learning Theory*. MIT Press, Cambridge, Massachusetts, 1994.
 - [52] S. Russell. Learning agents for uncertain environments. In *Proceedings of Computational Learning Theory ’98*. ACM Press, 1998.
 - [53] M. J. Kearns. *The Computational Complexity of Machine Learning*. MIT Press, Cambridge, Mass., 1990.
 - [54] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
 - [55] E. Baum. Manifesto for an evolutionary economics of intelligence. In C. M. Bishop, editor, *Neural Networks and Machine Learning*, New York, 1998. Springer-Verlag.
 - [56] A. Hastings. *Population Biology: Concepts and Models*. Springer-Verlag, New York, 1997.
 - [57] J. H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, Massachusetts, 1993.
 - [58] J. Maynard-Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, England, 1987.
 - [59] P. Berge, Y. Pomeau, and C. Vidal. *Order within chaos: towards a deterministic approach to turbulence*. Wiley, New York, 1986.
 - [60] S. H. Strogatz. *Nonlinear Dynamics and Chaos: with applications to physics, biology, chemistry, and engineering*. Addison-Wesley, Reading, Massachusetts, 1994.
 - [61] C. Beck and F. Schlögl. *Thermodynamics of Chaotic Systems*. Cambridge University Press, 1993.
 - [62] T. Shinbrot, C. Grebogi, E. Ott, and J. A. Yorke. Using small perturbations to control chaos.

- Nature*, 363:411–417, 1993.
- [63] K. Huang. *Statistical Mechanics*. J. Wiley and Sons, New York, 1963.
 - [64] J. J. Binney, N. J. Dowrick, A. J. Fisher, and M. E. J. Newman. *The Theory of Critical Phenomena*. Oxford University Press, Oxford, 1992.
 - [65] L. E. Reichl. *A Modern Course in Statistical Physics*. University of Texas Press, Austin, Texas, 1980.
 - [66] J. M. Yeomans. *Statistical Mechanics of Phase Transitions*. Clarendon Press, Oxford, 1992.
 - [67] J. Rissanen. *Stochastic Complexity in Statistical Inquiry*. World Scientific, Singapore, 1989.
 - [68] Vladimir Naumovich Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, Berlin, 1995.
 - [69] P. Manneville. *Dissipative structures and weak turbulence*. Academic Press, Boston, 1990.
 - [70] C. Godreche and P. Manneville, editors. *Hydrodynamics and nonlinear instabilities*. Cambridge University Press, New York, 1998.
 - [71] E. Baum. Toward a model of mind as an economy of agents. *Machine Learning*, page in press, 1999.
 - [72] L. G. Valiant. A theory of the learnable. *Comm. ACM*, 27:1134–1142, 1984.
 - [73] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, 1991.
 - [74] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, Reading, 1979.
 - [75] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, Reading, Massachusetts, 1994.
 - [76] J. P. Crutchfield and C. R. Shalizi. Thermodynamic depth of causal states: Objective complexity via minimal representations. *Physical Review E*, 59(1):275–283, 1999.
 - [77] N. Ay and J. P. Crutchfield. Reductions of hidden information sources. *J. Stat. Phys.*, 210(3-4):659–684, 2005.
 - [78] C. R. Shalizi, K. L. Shalizi, and J. P. Crutchfield. Pattern discovery in time series, Part I: Theory, algorithm, analysis, and convergence. *Journal of Machine Learning Research*, submitted, 2002. Santa Fe Institute Working Paper 02-10-060; arXiv.org/abs/cs.LG/0210025.
 - [79] K. L. Shalizi, C. R. Shalizi, and J. P. Crutchfield. Pattern discovery in time series, Part II: Implementation, evaluation, and comparison. 2002. in preparation.
 - [80] D. P. Varn and J. P. Crutchfield. From finite to infinite range order via annealing: The causal architecture of deformation faulting in annealed close-packed crystals. *Phys. Lett. A*, 234(4):299–307, 2004.
 - [81] D. P. Varn, G. S. Canright, and J. P. Crutchfield. Inferring pattern and disorder in close-packed structures from X-ray diffraction studies, Part II: Structure and intrinsic computation in zinc sulphide. *Phys. Rev. B*, 2002. submitted; Santa Fe Institute Working Paper 03-03-022; arXiv.org/abs/cond-mat/0302587.
 - [82] D. P. Varn, G. S. Canright, and J. P. Crutchfield. Inferring pattern and disorder in close-packed structures from X-ray diffraction studies, Part I: ϵ -Machine spectral reconstruction theory. *Phys. Rev. B*, 2002. submitted; Santa Fe Institute Working Paper 03-03-021; arXiv.org/abs/cond-mat/0302585.

- [83] D. P. Varn, G. S. Canright, and J. P. Crutchfield. Discovering planar disorder in close-packed structures from x-ray diffraction: Beyond the fault model. *Phys. Rev. B*, 66(17):174110–3, 2002.
- [84] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence. *CHAOS*, 13(1):25–54, 2003.
- [85] J. P. Crutchfield and D. P. Feldman. Synchronizing to the environment: Information theoretic limits on agent learning. *Adv. in Complex Systems*, 4(2):251–264, 2001.
- [86] D. P. Feldman and J. P. Crutchfield. Structural information in two-dimensional patterns: Entropy convergence and excess entropy. *Phys. Rev. E*, 67(5):051103, 2003.
- [87] C. S. McTague and J. P. Crutchfield. Automated pattern discovery—an algorithm for constructing optimally synchronizing multi-regular language filters. *Theoretical Computer Science*, in press, 2005. Santa Fe Institute Working Paper 04-09-027; arxiv.org:adap-org/0410017.
- [88] H. E. Stanley. *Introduction to Phase Transitions and Critical Phenomena*. Oxford University Press, Oxford, 1971.
- [89] J. Rissanen. Universal coding, information, prediction, and estimation. *IEEE Trans. Info. Th.*, IT-30:629, 1984.
- [90] C. S. Wallace and D. M. Boulton. An information measure for classification. *Comput. J.*, 11:185, 1968.
- [91] J. P. Crutchfield and B. S. McNamara. Equations of motion from a data series. *Complex Systems*, 1:417 – 452, 1987.
- [92] V. Vapnik. Three fundamental concepts of the capacity of learning machines. *Physica A*, 200:538–544, 1993.
- [93] J. Ziv and A. Lempel. Compression of individual sequences via variable-rate encoding. *IEEE Trans. Info. Th.*, IT-24:530, 1978.
- [94] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Champaign-Urbana, 1962.
- [95] E. van Nimwegen, J. P. Crutchfield, and M. Mitchell. Finite populations induce metastability in evolutionary search. *Phys. Lett. A*, 229:144–150, 1997.
- [96] E. van Nimwegen, J. P. Crutchfield, and M. Mitchell. Statistical dynamics of the royal road genetic algorithm. *Theoret. Comp. Sci.*, 229:41–102, 1999.
- [97] E. van Nimwegen and J. P. Crutchfield. Optimizing epochal evolutionary search: Population-size independent theory. *Computer Methods in Applied Mechanics and Engineering, special issue on Evolutionary and Genetic Algorithms in Computational Mechanics and Engineering*, D. Goldberg and K. Deb, editors, 186:171–194, 2000.
- [98] E. van Nimwegen and J. P. Crutchfield. Optimizing epochal evolutionary search: Population-size dependent theory. *Machine Learning*, 45(1):77–114, 2001.
- [99] F. Jimenez-Morales, M. Mitchell, and J. P. Crutchfield. Evolving one dimensional cellular automata to perform a non-trivial collective behavior task: One case study. In *Computational Science–ICCS 2002, Part I, Proceedings 2329*, pages 793–802, 2002.
- [100] J. P. Crutchfield, M. Mitchell, and R. Das. The evolutionary design of collective computation in cellular automata. In *Evolutionary Dynamics—Exploring the Interplay of Selection*,

- Neutrality, Accident, and Function*, Santa Fe Institute Series in the Sciences of Complexity, pages 361–411. Oxford University Press, 2001.
- [101] J. P. Crutchfield and M. Mitchell. The evolution of emergent computation. *Proc. Natl. Acad. Sci.*, 92:10742–10746, 1995.
 - [102] L. J. Fogel, A. J. Owens, and M. J. Walsh, editors. *Artificial Intelligence through Simulated Evolution*. Wiley, New York, 1966.
 - [103] H. Zhou and J. J. Grefenstette. Induction of finite automata by genetic algorithms. In *Proc. of the 1986 IEEE Int’l Conf. on Systems, Man, and Cybernetics*, Atlanta, GA, 1986. IEEE Press.
 - [104] M. M. Lankhorst. Grammatical inference with a genetic algorithm. In *Proc. of the 1994 EUROSIM Conference on Massively Parallel Processing Applications and Development*, Amsterdam, 1994. Elsevier.
 - [105] A. Paz and J. W. Carlyle. Chaotic evolution via generalized probabilistic automata (probabilistic arrays). In S. Bozapalidis, editor, *Proc. of the 3rd International Conference Developments in Language Theory*, pages 9–35, Thessaloniki, Greece, 1997. Aristotle University.
 - [106] M. Schroder. Automatic design of multidimensional controllers by means of an evolutionary algorithm. *Fuzzy Sets and Systems*, 89(3):321–333, 1997.
 - [107] L. J. Fogel, P. J. Angeline, and D. B. Fogel. An evolutionary programming approach to self-adaptation on finite state machines. In J. R. McDonnell, R. G. Reynolds, and D. B. Fogel, editors, *Evolutionary Programming IV*, pages 355–365, Cambridge, MA, 1995. MIT Press.
 - [108] J. P. Crutchfield and P. K. Schuster. *Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function*. Santa Fe Institute Series in the Sciences of Complexity. Oxford University Press, 2002.
 - [109] J. P. Crutchfield and O. Gernerup. Objects that make objects: The population dynamics of structural complexity. *J. Roy. Soc. Interface*, 3:345–349, 2006. doi:10.1098/rsif.2006.0114.
 - [110] O. Gernerup and J. P. Crutchfield. Hierarchical self-organization in the finitary process soup. In *Proceedings of Artificial Life X*, volume in press, 2006.
 - [111] D. Albers, J. C. Sprott, and J. P. Crutchfield. Persistent chaos in high dimensions. *Physical Review Letters*, submitted, 2005. Santa Fe Institute Working Paper 05-04-011; arXiv.org/abs/nlin/0504040.
 - [112] Y. Sato and J. P. Crutchfield. Coupled replicator equations for the dynamics of learning in multiagent systems. *Physical Review E*, 67(1):40–43, 2003. Santa Fe Institute Working Paper 02-04-017; arXiv.org/abs/nlin/0204057.
 - [113] Y. Sato, E. Akiyama, and J. P. Crutchfield. Stability and diversity in collective adaptation. *Physica D*, 210(1-2):21–57, 2005. Santa Fe Institute Working Paper 04-08-025; arXiv.org/abs/nlin/0408039.