

Coevolutionary & Coadaptive Systems Symposium Program

Friday, November 4

09.00–
10.30 INTRODUCTION AND COMMENTS

Solution Concepts in Coevolutionary Algorithms

Sevan Ficici

Coevolutionary algorithms are stochastic population-based search methods that are often applied to games of strategy; the nature of these games can run the gamut from cooperative to competitive. Since evolution entails the "survival of the fittest," we are enticed to assume that the co-evolution of game-playing agents should necessarily lead to increasingly competent strategic play. In practice, however, this expectation is often disappointed. In this talk, I will discuss how the consideration of solution concepts for coevolutionary search can, first, help explicate the gap between our expectations and the actual behavior of coevolutionary algorithms, and, second, help suggest new algorithms that behave more in line with our goals.

11.00–
12.30 **Coevolution of Neural Networks Using a Layered Pareto Archive**

German Monroy, Kenneth O. Stanley, & Risto Miikkulainen

The Layered Pareto Coevolution Archive (LAPCA; De Jong, 2004b) was recently proposed as an effective Coevolutionary Memory (CM) which, under certain assumptions, guarantees monotonic progress in coevolution. In this paper, a technique is developed that interfaces the LAPCA algorithm with NeuroEvolution of Augmenting Topologies (NEAT), a method to evolve neural networks with demonstrated efficiency in game playing domains. In addition, the behavior of LAPCA is analyzed for the first time in a complex gameplaying domain: evolving neural network controllers for the game Pong. The technique is shown to keep the total number of evaluations in the order of those required by NEAT, making it applicable to complex domains. Pong players evolved with a LAPCA and with the Hall of Fame (HOF) perform equally well, but the LAPCA is shown to require significantly less space than the HOF. Therefore, combining NEAT and LAPCA is found to be an effective approach to coevolution.

Evaluation Structures and Innovative Dynamics

Anthony Bucci

One characteristic of coevolutionary algorithms is their lack of an explicit measure of value. Unlike typical optimization or evolutionary algorithms, which rely on the existence of an objective function, coevolutionary algorithms must develop metrics through time while simultaneously discovering entities which optimize them. With a few notable exceptions, generally the means for developing such metrics has been a black art. My focus here is therefore on the problem of evaluation. With this problem in mind, I pose the question: what should a coevolutionary algorithm be doing? My simple answer is that an algorithm should be discovering and maintaining as many innovations as possible. In order to pose the question and answer precisely, in this talk I will develop a theoretical tool called an evaluation structure. Such is a structured mathematical object built from a population and representing salient aspects, or measurements, of entities' capabilities. Examples include the Pareto covering order and Ficici's notion of measurement table. I use this tool to distinguish two types of algorithm transitions: natural ones, which preserve but elaborate existing structure; and unnatural ones, which upset structure. I prove that, under certain conditions, unnatural transitions are necessary in the sense that an algorithm which does not make them will never reach a solution. Under these conditions I dub such transitions "innovative," as they are more or less radical changes in how we perceive the population to be structured which are required to reach solution. I conclude by discussing algorithm heuristics derivable from these observations as well as implications for the significance of selection in coevolutionary algorithms.

14.00–
15.30 **Complexification and the Arms Race**
Kenneth Stanley

Much recent research in competitive coevolution has focused on the problem of selecting appropriate opponents to establish a valid gradient for improvement. By carefully selecting opponents the hope is that problems like running in circles (i.e. the Red Queen effect) or disengagement can be avoided, thereby facilitating a genuine arms race of increasing sophistication. However, an arms race requires more than measuring against appropriate opponents; it also requires that the *representation* of the strategies can continue to acquire new elaborations without losing what is already known. For such a process to be possible, the genome must be able to expand, i.e. *complexify*, by adding new genes. That way, the representational space of strategies is virtually unlimited. In this talk, I describe such a system, NeuroEvolution of Augmenting Topologies (NEAT), and how it addresses the challenges that arise when genomes are allowed to grow.

Towards Metrics and Visualizations Sensitive to Coevolutionary Failures

Ari Bader-Natal & Jordan B. Pollack

The task of monitoring success and failure in coevolution is inherently difficult, as domains need not have any external metric to measure performance. Past metrics and visualizations for coevolution have been limited to identification and measurement of success but not failure. We suggest circumventing this limitation by switching from best-of-generation-based techniques to all-of-generation-based techniques. Using all-of-generation data, we demonstrate one such technique—a population-differential technique that allows us to profile and distinguish an assortment of coevolutionary successes and failures, including arms-race dynamics, disengagement, cycling, forgetting, and relativism.

16.00–
17.30 **On the Coevolutionary Construction of Learnable Gradients**

Shivakumar Viswanathan & Jordan B. Pollack

The best way for adaptive agents to learn is to be exposed to problems that are just a little more difficult than those they already know how to solve. While this has been a guiding concept in developing algorithms for gradient construction in coevolution, this has remained largely an intuition rather than a formal concept. In this paper, we build on the order-theoretic formulation of coevolution to develop some preliminary formal concepts towards clarifying the nature of the relation between the variational structure imposed by the representation and coevolutionary learning. By explicating marrying the learnability problem to the variational structure of the learner space, we describe a basic idealization of how coevolution with an Ideal Teacher could inherently address the problem of appropriate gradient creation with the intent that this could serve as a basis to developing practical algorithmic mechanisms that approximate this idealized behavior.

The Effects of Synthetic Social Structures on Teams of Autonomous Vehicles

Annie Wu

This work explores the effects of synthetic social structures on the behavior of teams of autonomous vehicles (AVs) that are working together to achieve a mutual goal. We examine the performance of teams of AVs on the "opera problem" and find that the inclusion of social structures and social interactions can improve the efficiency of the team as a whole.

09.00– **Multi-Objective Multi-Agent Planning**10.30 *Abdel-Ilhah Mouaddib*

Multi-Objective Multi-Agent Planning addresses the problem of resolving conflicts between individual agent interests and the group interests. In this paper, we address this problem using a Vector-Valued Decentralized Markov Decision Process (2V-DEC-MDP). The formal framework of a Vector-Valued MDP considered uses the value function which returns a vector representing the individual and the group interests. To define an optimal policy, each agent has to perform a multicriteria optimization process. To do that, the approach we present uses Egalitarian Social Welfare orderings that allow to an agent to consider during its local optimization the satisfaction of all criteria and reducing their differences. The obtained result is an equilibrium of individual and group satisfactions where the local policies can lead to more global satisfying behaviors in some settings. This result is illustrated in an example and compared to alternate local policies.

Response Regret*Martin Zinkevich*

The concept of regret is designed for the long-term interaction of multiple agents. However, most concepts of regret do not consider even the short-term consequences of an agent's actions: e.g., how other agents may be nice to you tomorrow if you are "nice" to them today. For instance, an agent that always defects while playing the Prisoners Dilemma will never have any internal or external regret. In this paper, we introduce a new concept of regret, called response regret, that allows one to consider both the immediate and short-term consequences of one's actions. Thus, instead of measuring how an action affected the utility on the time step it was played, we also consider the consequences of the action on the next few time steps, subject to the dynamic nature of the other agents' responses: e.g. if the other agent always is nice to us after we are nice to it, then we should always be nice; however, if the other agent sometimes returns favors and sometimes doesn't, we will not penalize our algorithm for knowing when these times are. We develop algorithms for both external response regret and internal response regret, and show how if two agents minimize internal response regret, then they converge to a correlated equilibrium in repeated bimatrix games, stochastic games, and partially observable stochastic games.

11.00– **Asynchronous Chess**12.30 *Nathaniel Gemelli, Robert Wright, & Roger Mailler*

We present adversarial agent work being done in a real-time asynchronous environment, Asynchronous Chess (AChess). AChess is a real-time competitive experiment platform for developing new agent technologies using single-agent reasoning and/or multi-agent cooperative strategies. We aim to provide a simplified asynchronous environment that is stochastic in its nature, but easily described from its foundations as a synchronized game. The goal is to design agent technologies that perform better in these domains than existing single- and multi-agent methods. This research applies to non-deterministic agent-based search and reasoning in a simplified, yet still very complex environment.

Coevolution in Strategic Game AI*Chris Miles & Sushil Lewis*14.00– **Time-Dependent Collaboration Schemes for Cooperative Coevolutionary Algorithms**15.30 *Liviu Panait & Sean Luke*

Cooperative coevolutionary algorithms are a popular approach to learning via problem decomposition. One important aspect of cooperative coevolutionary algorithms concerns how to select collaborators for computing the fitness of individuals in different populations. We argue that using a fixed number of collaborators during the entire search may be suboptimal. We experiment with a simple ad-hoc scheme that varies the numbers of collaborators over time. Empirical comparisons in a series of problem domains indicate that decreasing the numbers of collaborators over time fares better than keeping the number fixed. We conclude with a brief discussion of our findings and suggest directions for future research.

Recent Advances in Cooperative Multiagent Learning*Liviu Panait*

Concurrent learning is a form of cooperative multiagent learning in which each agent has an independent learning process and little or no control over its teammates' actions. An agent's action selection directly influences the rewards received by all the agents; this results in a co-adaptation among the concurrent learning processes. Co-adaptation can drive the team towards suboptimal solutions because agents tend to select those actions that are rewarded better, without any consideration for how such actions may affect the search of their teammates. We propose two mechanisms to reduce such undesirable effects. First, we argue that agents may mutually benefit if they show lenience to one another during the learning process: ignore many of the low rewards initially, and fewer rewards as learning progresses. Second, we suggest that agents should also prefer actions that inform their teammates about the structure of the joint search space in order to help them choose from among various action options. We introduce several novel multiagent learning algorithms to illustrate the advantages of lenience and informativeness.

16.00– **A Dynamical Systems Analysis of Collaboration Methods in Cooperative Coevolution**

17.30

Elena Popovici & Kenneth De Jong

Cooperative co-evolution is often used to solve difficult optimization problems by means of problem decomposition. In order to do so efficiently, one must have a good understanding of co-evolutions dynamical behavior. To build such understanding, we have constructed a methodology for analyzing co-evolution based on dynamical systems theory. In this paper we show how it can be applied to investigate the effects that collaboration methods have on performance and to identify a problem property relevant in this respect.

Robustness in Compositional Coevolution*R. Paul Wiegand*

Recent advances in coevolutionary theory have improved not only our understanding of the nature of coevolutionary problems and coevolutionary algorithm dynamics, but have also provided us with a better grasp of the relationship between the two. Centered around Sevan Ficici's idea of *solution concept*, a philosophy of application of coevolution is crystallizing: Understand what solution concept is induced by the algorithm, understand what solution concept the engineer has in mind, and attempt to match the two. This philosophy has been successful in applying coevolutionary algorithms to test-based problems, and I believe it is a useful philosophy for compositional problems as well.

This talk presents a general discussion of compositional coevolutionary algorithms from the solution concept perspective. I will begin with a high-level discussion of terms and concepts, explaining what I mean by "compositional problems/algorithms", and why I am abandoning terms like "cooperative coevolution". I will also review a common solution concept for compositional coevolutionary algorithms, *ideal partnership*, which relates to static function optimization—and explain why it is not particularly useful. Finally, the talk will consider a new solution concept, one that centers around the idea of *robustness*. I will outline a framework for defining robustness in a variety of contexts, offer a plausible specific definition useful to problem domains such as multiagent learning, and outline why such a definition may be a useful solution concept for composition coevolutionary algorithms. The goal of the talk is primarily to stimulate discussion to help guide research of the robustness solution concept for compositional coevolution. Secondly, my intention is to provide examples of a "solution concept view" in the compositional coevolutionary setting.