

On the coevolutionary construction of learnable gradients

Shivakumar Viswanathan and Jordan B. Pollack

DEMO lab, Dept of Computer Science, Brandeis University
Waltham, MA-02139
{shiva,pollack}@cs.brandeis.edu

Abstract

“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, it 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 explicitly 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.

Introduction

Coevolutionary problem-solving involves the simultaneous processes of gradient creation and gradient following. This dynamic has long been framed in terms of an arms-race between species competing for the same niche (Hillis 1991; Sims 1994; Cliff & Miller 1995; Rosin 1997; Floreano, Nolfi, & Mondada 1998). However, another metaphor that has been increasingly used to frame this dynamic treats the coevolving populations as playing the role of a *Learner* and a *Teacher/Trainer* respectively (Epstein 1994; Pollack & Blair 1998; Juillé 1999; Ficici 2004; Bucci & Pollack 2003; De Jong & Pollack 2004). Here, the teacher poses different problems for the learner (i.e. creates gradient) and the learner attempts to acquire the capability to solve these problems by repeated interactions with these problems (i.e. follows gradient).

An important guiding idealization with this metaphor is that of an *Ideal Trainer* (Juillé 1999) (for the sake of consistency, we will use the term Ideal Teacher instead of Ideal Trainer from this point on). An Ideal Teacher consistently poses problems that are not too “difficult” or too “easy” but having a level of difficulty that provides a learning gradient that is just appropriate to promote the adaptation of the learner based on its current capabilities. This is based on the

notion that “*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*” [emphasis in original] (Juillé 1999). By being able to pose such appropriate problems as the learner dynamically adapts, the Ideal Teacher is envisioned as directing learning in a way that enables a continuous open-ended improvement of the learner’s capability.

This idealization leads to the fundamental design question of how the Ideal Teacher can be operationalized by low inductive bias coevolutionary algorithms. In this regard, two differing interpretations of the Ideal Teacher, which we will refer to as the *challenge-centered* and *evaluation-centered* perspectives, have recently emerged. From a challenge-centered perspective, the intended role of the Ideal Teacher is one of *constructing* a learning gradient (or challenge) that can elicit a desired *dynamic transformation* of the learner population. This has been the classic approach of choice, especially within the arm-race framework. However, the intended role of an Ideal Teacher in the more recent evaluation-centered approach is one of *revealing* the gradient in capabilities that exists among the learners in the population, by effectively treating the current population as a more *static* entity (Ficici & Pollack 2001).

The motivating rationale of the evaluation-centered approach is that when the Teacher can consistently provide an accurate assessment of the current population of learners it can provide a fine gradient for selection to act on. Furthermore, it would rarely be the case that a “good” learner appearing in the population would be prematurely lost due to inaccurate evaluation. In this way, such an approach could, in principle, ensure that no regress occurs in the evolution of the learners. The key operational difficulty here is of consistently identifying the problems or tests that provide an *accurate* evaluation of the current learner population. Addressing this difficulty and related aspects of such an evaluation-centered approach has so far been the dominant focus of recent work in developing coevolutionary algorithms within the Teacher-Learner framework (Bucci & Pollack 2003; Bucci, Pollack, & De Jong 2004; De Jong & Pollack 2004).

While accurate evaluation is indeed an important aspect of coevolutionary problem solving, an evaluation-centered approach is silent on the question of how a succession of learners of increasing capabilities can be *actively* generated. So far, the variational properties of the learners are treated

as a blackbox and this approach is more geared to be able to identify suitable learners when they happen to appear in the population. This is in sharp contrast to the challenged-centered approach which is explicitly centered on promoting the active adaptation of the learners. Noting this difference, the question motivating the work here is of how a more active view of gradient construction could be incorporated into the dominantly evaluation-centered concepts of Pareto co-evolution.

In this paper, we translate the intuition of the Ideal Teacher into concrete formal concepts using the order-theoretic formulation of coevolution proposed by Bucci and Pollack (Bucci & Pollack 2003). The key proposal emerging from this theoretical exercise is the notion of a *Complete Learnable Test* set that describes the gradient construction properties of an Ideal Teacher.

Background

Juillé (1999) noted that coevolutionary problems can be considered to be a form of *multi-objective optimization*. Rather than a single objective defined by the environment, every problem posed by the teacher and encountered by the learning population presents an independent objective to be solved. Based on this conception of multi-objective problems, Ficici and Pollack (Ficici & Pollack 2001) and Noble and Watson (Noble & Watson 2001) explicitly brought the techniques of Multi-Objective Optimization (MOO) to bear on coevolution with the use of Pareto dominance as the corresponding solution concept. Here we adopt the order-theoretic framing of this Pareto coevolution problem proposed by Bucci and Pollack (Bucci & Pollack 2003), the relevant basic elements of which are described below.

Order-theoretic framework

Let the finite set of learners or students be S and the finite set of problems or tests (posed by the teacher) be T . The interaction between the learners and the tests is defined by the function $p : S \times T \rightarrow R$, where R is the set of ordered outcomes of the interactions between the learners and the tests. We restrict the focus to the case where $R = \{0 < 1\}$. Though represented by 0 and 1, these outcomes are tokens representing “not solved”(lose) and “solved”(win) respectively, rather than numerical payoffs and the relation between 0 and 1 is ordinal.

The function $p : S \times T \rightarrow R$ can be curried on T to obtain $p : S \rightarrow (T \rightarrow R)$, which assigns to each $s \in S$ the function $p_s : T \rightarrow R$. This function p_s describing the behavior of s for all tests in T will be referred to as the *interaction profile* for s . Each learner in S has such an interaction profile for all tests in T . Similarly, the function can be curried on S to obtain $p : T \rightarrow (S \rightarrow R)$. The corresponding function p_t for each test $t \in T$ is referred to as the *response profile* of t .

The function p can be represented as shown in Table 1. The value in position (i, j) is the outcome of the interaction between s_i and t_j , where $|S| = N$ and $|T| = M$. The entire row corresponding to learner s_i represents its interaction profile, and the entire column corresponding to test t_j represents its response profile. In keeping with the represen-

tion of the function p as a matrix, we will use the shorthand (s_i, t_j) to refer to the interaction $p(s_i, t_j)$.

	t_1	t_2	t_3	t_4	\dots	t_M
s_1	1	1	0	1	\dots	1
s_2	0	0	0	1	\dots	1
s_3	0	1	0	1	\dots	0
\dots	\dots	\dots	\dots	\dots	\dots	\dots
s_N	0	0	1	1	\dots	1

Table 1: Matrix representation of p

As the outcomes are ordered, a preference relation \preceq between each $s, s' \in S$ is defined by a pairwise comparison \leq_{pw} of the outcomes of each test in p_s with $p_{s'}$. The relation \preceq is treated here as being the *Pareto dominance* relation. So $s \preceq s'$ implies that s' Pareto-dominates s . This ordering on (S, \preceq) obtained from the complete set of interaction profiles defines a pre-order on S . So, the maximal elements of (S, \preceq) are the desired solutions to the problem.

Accordingly, the coevolutionary search problem is to find the maximal elements of the pre-order (S, \preceq) defined by the Pareto dominance relation over the set of all interaction outcomes as defined by p . Since there may be several maximal elements, it is important to note here that a solution to the problem may be a subset of S rather than necessarily being a single individual learner.

The class of problems of interest here are those where the size of T is extremely large, i.e. where $|T|$ is of the same order as $|S|$. For such problems evaluating each individual from S in the population against *all* the elements in T at every step of the search process (as with EMOO algorithms) is highly impractical. So, a motivating premise in using a coevolutionary approach is to solve such search problems by the use of only a limited population of tests at any given time in order to evaluate the learner population, where the test population can itself evolve over the search process.

From this point on, we will use the term “teacher” to refer to the search algorithm operating on the population of tests. The algorithm operating on the learner population, however, will simply be referred to as the learner algorithm. In the next section, we use these basic concepts to identify the test properties required to serve as a learnable gradient.

Test difficulty

As noted earlier, an Ideal Teacher can consistently provide tests that are neither too “difficult” nor too “easy” but that are at a level of difficulty that provides a learning gradient that is “just appropriate” to the promote the adaptation of the learner based on its current capabilities. In order to address how this notion can be operationalized, we need to define what it means for a test to be “difficult” or “easy” for a particular learner. We interpret this as a difference in terms of the *learnability* of a test.

Learnability and test difficulty

Ficici and Pollack (2001) define the *learnability* of a test with respect to a particular learner as “the probability that

the learner can be transformed, *over some number of variation steps*, to become competent (or more competent) at the task posed by the teacher” [emphasis added].

From this definition, we can see that given a learner $s \in S$ and a test $t \in T$ such that $(s, t) = 0$, the ability of s to learn to solve t is dependent on the variational structure of the learner space S . This space is essentially the set S augmented by the topological structure induced by the variational operators particular to the encoding of the members of S .

For simplicity, we restrict our attention to variation with mutation operators. With mutational operators, the topology induced on S can be assumed to take the form of an undirected graph $S = (S, E)$, where S is the vertex set and E is the set of edges. An edge $e \in E$ exists between s_i and s_j ($s_i, s_j \in S$) if and only if s_i can be obtained by a single application of the mutational operator μ to s_j . We assume here that the effect of the mutation operator is reversible, i.e. if s_i can be obtained from s_j by a single application of the operator, then the reverse is also possible.

Given this space S , if a learner $s' = \mu^n(s)$ can be obtained by n applications of the mutation operator to s such that $(s', t) = 1$, then it would follow that t is learnable by s ¹. Critical to this interpretation is the value of n . Here we focus on the case where $n = 1$. Therefore, the learnability of a test t by the learner s is the likelihood that there exists a learner $s' = \mu(s)$ such that $(s', t) = 1$.

Based on this notion of learnability, the “difficulty” of a test t for a learner s , can be interpreted as follows. A test t is said to be “too difficult” for a learner s , if $(s, t) = 0$ and $n > 1$ mutations of s are required to produce s' such that $(s', t) = 1$. However, if $(\mu(s), t) = 1$ then the problem is “appropriate”, in being just beyond the present capability of s . On the other hand, if $(s, t) = 1$ then no variation on s is required to solve the test. Such a test can be considered to be “too easy”. The “fitness landscape” corresponding to these three cases is shown in Figure 1.

Learnability and Improvement

The above definition of learnability with respect to a single test however requires a critical amendment in the context of the global search problem. Suppose $s' = \mu(s)$ was such that it indeed solves the problem t . This by itself is insufficient to determine whether s' Pareto dominates s , i.e. s' is better or no worse than s on *all* the tests in T .

In order for a teacher to ascertain whether s' is indeed a true improvement over s , the relative performance of the two learners would, in principle, need to be evaluated across *all* the tests in T . Indeed if the teacher could present all the tests to the learner at each instance then there would no demands on the teacher to provide graded challenges to the learner and there would no need for coevolution. Therefore, rather than posing a gradient defined by a single learnable test for s , we would ideally like the teacher to pose a *small and sufficient* collection of tests $\Delta \subset T$ such that if learnable by s would indicate an improvement with respect to the global solution concept.

In this regard, Δ would need to contain tests that s *can solve*, rather than only containing tests that s cannot solve. This is to avoid a situation where a variant s' that solves tests that s cannot solve also “forgets” how to solve the tests that s can solve. For example, consider the scenario in Table 2. Let $\Delta = \{t_1, t_2, t_3\}$ be the set of tests that s cannot solve. The perceived learnability of Δ due to the existence of a variant s' that solves all the tests in Δ is deceptive. Even though s' solves all the tests in Δ , it has “forgotten” how to solve t_5 . So, an evaluation of $s \preceq s'$ based on Δ alone would be inaccurate in this case.

	t_1	t_2	t_3	t_4	t_5	t_6
s	0	0	0	1	1	1
s'	1	1	1	1	0	1

Table 2: Deceptive evaluation due to “forgetting”

It is learnability in this stronger sense that is of particular relevance to the overall goals of coevolutionary search. This brings us to the question – given a learner s in S what is the set of sufficient tests Δ_s that are learnable by s ? The answer to this question follows from the above definitions, as discussed next.

The Complete Learnable Test set

As S can be treated as a graph, each learner s is associated with a set $\epsilon_s \subset S$ of all 1-neighbors obtained by a single application of the mutation operator to s . Corresponding to the edge between s and each member of ϵ_s is a unique test set as described below.

Consider the interaction profile of a learner s given by p_s . Applying a mutation to s produces another learner, say $s' \in \epsilon_s$. If the interaction profiles are such that $p_{s'} \neq p_s$, it implies that s and s' have different behaviors. Let $\Delta_{s,s'}$ be the set of all tests in T such that $\Delta_{s,s'} = \{t | (s, t) \neq (s', t), s' \in \epsilon_s, t \in T\}$.

The properties of the tests in $\Delta_{s,s'}$ can be interpreted in a dynamic way. The tests $\Delta_{s,s'} \subseteq T$ are *sensitive* to the variation of s by responding to this change by a change in their outcomes. Since each test $t \in \Delta_{s,s'}$ produces different values corresponding to s and s' , each t distinguishes between s and s' . Similarly, the tests in $T - \Delta_{s,s'}$ are *insensitive* as the change of s to s' does not result in a change in their values.

From this perspective, if the change in s to s' is such that only the tests in $\Delta_{s,s'}$ corresponding to interaction outcome of 0 with s change their values to being 1 with s' then it implies that $s \preceq s'$. Similarly if this change in the test outcomes is from 1 to 0 then it implies that $s' \preceq s$. And finally, if there exist at least two tests in $\Delta_{s,s'}$ such that one changes its outcome from being 0 to 1 and the other from 1 to 0, then s and s' are mutually non-dominated or incomparable by \preceq .

Such a set Δ_{s,s_i} of “sensitive” tests, with respect to s and s_i , exists for each $s_i \in \epsilon_s$. This set Δ_{s,s_i} can be considered to be an attribute associated with the edge joining s and each s_i as shown in Figure 2. The *complete* set of tests which are learnable (and possibly improvable) by s can therefore be

¹ $\mu^n(s)$ is used to indicate n applications of μ to s




Figure 1: Effect of variation on the “difficulty” of learning to solve test t by a learner s




Figure 2: Subgraph of \mathcal{S} corresponding to $\epsilon_s \cup \{s\}$

obtained as $\Delta_s = \bigcup \Delta_{s,s_i}$. Δ_s is *complete* in that s cannot learn to or forget how to solve any further test from T , for the given variational structure \mathcal{S} . From this point on we will refer to Δ_s as being the *Complete Learnable Test* (CLT) set for s .

At the outset, we can see that the Complete Learnable Test set has the following characteristics:

- If $s \preceq s'$ with respect to Δ_s ($s' \in \epsilon_s$) then $s \preceq s'$ with respect to T .
- A test t , where $(s, t) = 0$, is learnable by s if and only if t is a member of Δ_s .
- Similarly, for every test $t' \in \Delta_s$ where $(s, t') = 1$, there exists some variant in ϵ_s that can forget how to solve t' .

It is important to note that Δ_s is not necessarily the minimal set of tests required to *accurately* evaluate the relation between s and its neighbors, if they were simultaneously present. From the perspective of the *underlying dimensions* of the problem (Bucci, Pollack, & De Jong 2004), the set $\Delta_{s,s'}$ may contain a number of tests that are redundant in the information that they provide. Furthermore, there may be different non-minimal proper subsets of $\Delta_{s,s'}$ that can perform the same role, i.e. where the relation between s and every $s' \in \epsilon_s$ as evaluated using these test sets is identical to that obtained with Δ_s . It is in this sense that Δ_s is the set

of *sufficient* tests for evaluating learning though all the tests are not *necessary* for this purpose.

So, to summarize what we have achieved: Starting from the general intuition about the dynamic behavior of an Ideal Teacher, we have arrived at a definition of a specific concept describing the exact properties of the tests generated by the Ideal Teacher to achieve this dynamic process of continuous learning. So, when we speak of an Ideal Teacher that constructs a learnable gradient for an individual learner, the gradient it provides to the learner takes the form the tests of a Complete Learnable Test set. In the next section, we describe how such an Ideal Teacher can produce the dynamic of continuous improvement of a gradient following learner by generating the CLT test for a succession of learners.

Idealized coevolution

Let $\gamma = \{\Delta_s | s \in \mathcal{S}\}$. This is the set of all the CLT sets corresponding to each of the elements s in \mathcal{S} . We can define a topological structure as $\Gamma = (\gamma, E)$, where an edge e exists between Δ_s and $\Delta_{s'}$ if and only if there exists an edge between s and s' in \mathcal{S} , i.e. Γ and \mathcal{S} are isomorphic. The key idea that we propose here is the conception of the Ideal Teacher as operating on the structured state space defined by Γ rather than on the test space T . Whenever presented with a learner from \mathcal{S} , the meta-problem that the teacher poses to this learner is not a single test but a collection of tests corresponding to a particular member of γ .

This process can be conceived as taking the form shown in Figure 3. Given a learner s , the meta-problem posed by the teacher is the corresponding collection of tests Δ_s . Given the gradient posed by Δ_s , the learner performs a local hill-climbing operation. All the variants of s are generated, and if a variant $s' \in \epsilon_s$ dominates s with respect to the tests in Δ_s then it is selected.

Rather than a synchronous adaptation, at the next iteration when presented with s' , the teacher correspondingly performs a local search in Γ using s' as the basis to find the corresponding CLT set for s' , i.e. ideally “moving” along the edge from Δ_s to $\Delta_{s'}$. The test set $\Delta_{s'}$ is in turn presented as the learning gradient for s' , and so on. In this idealization, the tests posed by the teacher are *always* learnable and the learning that occurs corresponds to progress with respect to the global learning problem.




Figure 3: Idealized (asynchronous) coevolution with CLT sets

If this idealized coevolutionary process were realizable in this form, the pathologies of “disengagement” (i.e. loss of gradient), “forgetting” and cycling (Watson & Pollack 2001) would be impossible. Even so, one pathology typical of hill-climbers would however be present, namely, of the learners getting stuck on local optima when the learner dominates or is incomparable to all its neighbors. As all the relevant tests are intrinsically contained in each CLT set, there would be no possibility of making a locally inaccurate evaluation to escape the local optimum and the notion of variation opening up a new dimension along which learning could continue (Ficici & Pollack 2001; Stanley & Miikkulainen 2004) would be meaningless.

Discussion

Given the formulation developed in the previous sections and the notion of the CLT set, the question that follows is how such an idealized teaching-learning process could be approximated by an actual algorithm.

Central to addressing this issue is that of how a CLT set Δ_s of a learner s could be used to construct the CLT set of its neighbors in ϵ_s , to continually provide relevant gradient to the iteratively improving learner. At the outset, we can see that if s and s' are neighbors in \mathcal{S} , then $\Delta_{s,s'} \subset \Delta_s \cap \Delta_{s'}$. So the question is one of how the tests in $\Delta_{s'} - \Delta_{s,s'}$ can be obtained by search through T using $\epsilon_s \cup s$ as the gradient and without generating any further variants of s' .

A starting point towards this end could be to consider the teachability properties of the test space \mathcal{T} . Due to the dual nature of learners and tests as defined by p , exactly the same rationale as described in Section 3.2 applies to the test space as well. Much like the Complete Learnable Test set, there exists a *Complete Teachable Learner* set, ∇_t , associated with each $t \in \mathcal{T}$. This is the set of all learners that exhibit differences in behavior for all 1-mutations of the test t . This is equivalent to the distinctions observable when moving columnwise in the interaction matrix of p .

This formulation also brings to light another critical conceptual issue that we believe has not been widely appreciated before, namely the amount of bias required to discover useful variation. At the minimum, given a single learner s , the teacher is posed with the search problem of finding an appropriate learnable test t from the set of possible tests T . This test t requires to satisfy two criteria, namely, (a) $(s, t) = 0$ and (b) there exists some $s' = \mu(s)$ such that

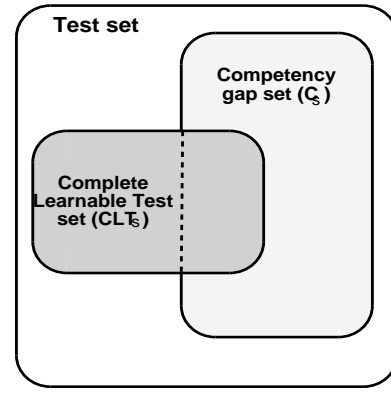


Figure 4: Venn diagram of relation between T , C_s and Δ_s for a learner s

$(s', t) = 1$. These requirements present a formidable search problem for the teacher.

If a student s is randomly selected from \mathcal{S} , then the teacher “knows” little about the student beyond the fact that the outcome would either be 0 or 1 for whatever test t may be picked from T . The teacher’s lack of knowledge about the student’s identity has another broader implication.

Let the *competence gap set* $C_s \subseteq T$ be the set of all tests such that $t \in C$ if and only if $(s, t) = 0$. The relation between T , C_s and Δ_s is shown in Venn-diagram form in Figure 4. Given this relation, suppose the teacher randomly picks two tests t_a and t_b and the resultant outcome of interactions are $(s, t_a) = (s, t_b) = 0$. While t_a and t_b reveal gaps in the competence of the learner, this raises the question of how the teacher can assess which (if any) of these two tests (or both) is in L_s and can provide an appropriate learning gradient for s , given that learnability involves variational behavior that has *yet to occur*.

One obvious way to make this assessment is by actually evaluating the tests with each of the 1-variants of s . However, such an after-the-fact assessment is effectively an exhaustive trial and error process, and reduces learnability to an irrelevant concept. So, we can see that a fundamental issue in addressing the teacher’s search problem is in being able to assess the learnability of tests with respect to s *indirectly* without exhaustively generating the variants of s . This raises the fundamental issue of whether such an indirect evaluation could be achieved without some external heuristic that provides additional information about where tests belonging to the CLT may be more readily found in the test-space. Our conjecture, based on preliminary investigations, suggest that such meta-heuristics could play a considerable role in achieving this end.

Conclusion

While the impact of the dispositional properties of the variational properties on evolution has been receiving increasing attention in the Evolutionary Computation literature (Wagner & Altenberg 1996; Wolpert & Macready 1997), their avatars in coevolution taking the form of *learnability* and *teachability* have remained largely unstudied despite the

recognition of their importance. Towards engaging this complex issue, this paper presents a preliminary attempt to develop the formal concepts and machinery to make explicit the relation between the variational structure imposed by the representation and coevolutionary learning. While the approach adopted here is clearly simplistic and focussed on an idealization, it presents a framework that could act as a source of intuitions and analysis in developing algorithmic operationalizations to address the issue of gradient construction that is central to coevolution.

Acknowledgements

The authors would like to thank Sevan Ficici, Anthony Bucci and Ari Bader-Natal for discussions leading up to this paper.

References

- Bucci, A., and Pollack, J. B. 2003. A mathematical framework for the study of coevolution. In De Jong, K. A.; Poli, R.; and Rowe, J. E., eds., *Foundations of Genetic Algorithms 7*. San Francisco: Morgan Kaufmann. 221–236.
- Bucci, A.; Pollack, J. B.; and De Jong, E. D. 2004. Automated Extraction of Problem Structure. In et al., K. D., ed., *Proceedings of the Genetic and Evolutionary Computation Conference 2004 (GECCO 2004)*, 501–512. Seattle, Washington, USA: Springer-Verlag, Lecture Notes in Computer Science Vol. 3102.
- Cliff, D., and Miller, G. 1995. Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Morán, F.; Moreno, A.; Merelo, J. J.; and Chacón, P., eds., *Advances in Artificial Life: Third European Conference on Artificial Life*, number 929 in Lecture Notes in Computer Science, 200–218. Berlin, New York: Springer.
- De Jong, E., and Pollack, J. B. 2004. Ideal evaluation from coevolution. *Evolutionary Computation* 12(2):159 – 192.
- Epstein, S. L. 1994. Toward an ideal trainer. *Machine Learning* 15(3):251–277.
- Ficici, S. G., and Pollack, J. B. 2001. Pareto optimality in coevolutionary learning. In Kelemen, J., and Sosik, P., eds., *Advances in Artificial Life: 6th European Conference (ECAL 2001)*. Springer.
- Ficici, S. 2004. *Solution Concepts in Coevolutionary Algorithms*. Ph.D. Dissertation, Brandeis University.
- Floreano, D.; Nolfi, S.; and Mondada, F. 1998. Competitive co-evolutionary robotics: From theory to practice. In *From Animals to Animats 4*. MIT Press.
- Hillis, D. 1991. Co-evolving parasites improves simulated evolution as an optimization procedure. In C. Langton, C. Taylor, J. F., and Rasmussen, S., eds., *Artificial Life II*. Reading, MA: Addison-Wesley.
- Juillé, H. 1999. *Methods for Statistical Inference: Extending the Evolutionary Computation Paradigm*. Ph.D. Dissertation, Brandeis University.
- Noble, J., and Watson, R. A. 2001. Pareto coevolution: Using performance against coevolved opponents in a game as dimensions for Pareto selection. In et al., L. S., ed., *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2001)*, 493–500. San Francisco, California: Morgan Kaufmann Publishers.
- Pollack, J. B., and Blair, A. D. 1998. Coevolution in the successful learning of backgammon strategy. *Machine Learning* 32:225–240.
- Rosin, C. D. 1997. *Coevolutionary Search Among Adversaries*. Ph.D. Dissertation, University of California, San Diego.
- Sims, K. 1994. Evolving 3D morphology and behavior by competition. In Brooks, R., and Maes, P., eds., *Artificial Life IV*, 28–39. MIT Press.
- Stanley, K. O., and Miikkulainen, R. 2004. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research* 21:63–100.
- Wagner, G., and Altenberg, L. 1996. Complex adaptations and the evolution of evolvability. *Evolution* 50:967–976.
- Watson, R. A., and Pollack, J. B. 2001. Coevolutionary dynamics in a minimal substrate. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2001)*, 702–709.
- Wolpert, D. H., and Macready, W. G. 1997. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* 1(1):67–82.