
When Coevolutionary Algorithms Exhibit Evolutionary Dynamics

Sean Luke

Department of Computer Science
George Mason University
<http://www.cs.gmu.edu/~sean/>

R. Paul Wiegand

Department of Computer Science
George Mason University
paul@tesseract.org

Abstract

The task of understanding the dynamics of coevolutionary algorithms or comparing performance between such algorithms is complicated by the fact the internal fitness measures are subjective. Though a variety of techniques have been proposed to use external or objective measures to help in analysis, there are clearly properties of fitness payoff (e.g., intransitivity) which call such methods into question in certain contexts. We present a model of competitive fitness assessment with a single population and non-parametric selection (such as tournament selection), and show minimum conditions and examples under which an objective measure exists, and when the dynamics of the coevolutionary algorithm are identical to those of a traditional EA. We also discuss terminological difficulties in the coevolution literature, and present a detailed description of external measures presently in use in the literature.

1 Introduction

Traditional evolutionary algorithms (EAs) assess the fitness of an individual objectively, that is, independent of the population context in which the individual is placed. Coevolutionary algorithms (CEAs) operate much like traditional EAs except that fitness assessment is not objective but subjective: an individual is evaluated through its interaction with other individuals in the evolutionary system. Because fitness is subjective in CEAs, it is not clear under what conditions a CEA would be expected to optimize in a fashion like a traditional EA would solving a static problem. Without a firmer understanding of these conditions, the usefulness of CEAs as optimization procedures is unclear.

We consider two common reasons why knowing these conditions would be helpful. First, one might want to feel safe

in the knowledge that the designed CEA is following the same trajectory as a traditional EA; even if it is not clear what the traditional EA's objective fitness would be. Second, one might want to know an actual equivalent objective fitness in order to measure the progress of the coevolutionary algorithm, either to analyze how the algorithm is improving during a given run, or to compare performance between algorithms. Indeed, understanding how the algorithm is progressing during a run is of particular importance to coevolutionary algorithms, since a variety of pathological dynamics such as cycling, mediocre stability, and loss of an adaptive gradient can occur to spoil search efforts. As it turns out, these two reasons do not result in the same set of conditions: the second condition is more specific than and implies the first.

In this paper we will shed some light on these issues with respect to a specific, common form of coevolutionary procedure: coevolution with a single population, using a non-parametric selection procedure such as tournament selection, rank selection, or truncation selection. We realize this is somewhat different from much of the coevolution theory, which has focused on two-population coevolution and has typically applied a fitness-based (parametric) selection method. A non-parametric selection procedure permits us to shift from a fitness function to a *ranking function*, allowing us to make guarantees even with unusual fitness values as long as the ranks among individuals remain consistent. For purposes of this paper, we will also operate under the standard EGT assumption that selection is done with full mixing, that is, an individual's fitness is determined through competition with every other individual in the population, including the individual itself. In addition, the mathematical models presented here assume infinite population sizes.

This paper will first clarify terminology and make use of a basic game-theoretic framework. We will then use this model to leverage properties of the internal subjective payoff of CEAs, beyond that of simple transitivity, to help understand first under what conditions we may establish a true

objective measure for a given CEA, in order to chart evolutionary progress or compare techniques. Then we relax the framework enough to include conditions which are not formally objective, but still guarantee that the CEA will behave dynamically just as some related traditional EA.

2 Types of Coevolution Progress Measures

Most attempts at competitive coevolution have met with the problem of attempting to chart the algorithm’s progress due to the subjectivity of the fitness measure. To address this problem, there is a growing body of research which is focusing on the questions surrounding assessment of coevolutionary progress using *external* and *objective* measures. However the terminology has been used relatively loosely until this point, even though there are clear differences between the notions of externality and objectivity.

Historically, getting a handle on how to measure the true progress of coevolutionary algorithms almost always involves some kind of measurement *external* to the algorithm’s dynamics. For instance, Cliff and Miller (1995) uses several methods (frequently involving ancestral history of individuals) as an external method for attempting to detect when pathological dynamics such as cycling occurs.

This method is still *subjective* in the sense that the progress measurement of individuals is contextually dependent on other individuals. They are different than the actual fitness used by the algorithm, however, because they are *external* to the algorithm’s dynamics. There are other external progress measures which have been proposed which are *objective*. For instance, many optimization oriented problems have been charted using a static (and external) representative sampling set of the strategy space (Rosin and Belew, 1997). Watson and Pollack (2001) also explicitly describe an analysis substrate which involves using a known external and objective measure. Also, order statistics have been employed in order to determine the conditions under which arms races can occur (Ficici and Pollack, 1998).

We feel that a clear terminological distinction should be made between the concepts of externality and objectivity. External progress measures are measures which do not affect the dynamics of a running algorithm, while internal progress measures are used directly or indirectly to affect the course of an algorithm’s run. Objective measures are those in which a given individual receives a measurement value irrespective of other individuals, while subjective fitness measures depend on individuals which either currently exist in some population participating in the evolution, or existed at some point during the evolutionary run.

The problem with many subjective external measures is that, while they may help one understand the dynamics of a given run of the algorithm, it isn’t clear how they can be

used to compare performances between algorithms. On the other hand, it is often true that we know little about whether a given objective progress measure being used measures behavior which is reasonable to expect from a given algorithm. Moreover, objective measures based on sampled sets depend on how representative those sets are of true teaching sets (and assume such exist) (Rosin and Belew, 1997).

Although these different progress measurement methods for CEAs serve different purposes, there are at least two important characteristics one would like out of such a measure in cases where some kind of optimization is being attempted. First, the measure should give us some indication of an algorithm’s performance in terms of the optimization problem we want to solve. Second, there should be some reason to believe the measure is somehow *connected* to the problem in terms of the real dynamics of the algorithm. For example, suppose the problem space being searched is that of competing behavior rule sets represented as bit strings. An objective measure for an individual could be a simple unitation of the bit string; however, we cannot expect this measure to be helpful for understanding the progress of the algorithm, since it has little or nothing to do with it.

This underscores one of the fundamental problems with coevolution: it is not always easy to know what optimization problem is being solved.

3 Our Coevolutionary Model

We will be focusing entirely on single population coevolutionary algorithms. As a result, the high-level model is similar to existing abstract descriptions of evolutionary algorithms (Vose, 1999). However, we are further restricting ourselves to the class of algorithms which use non-parametric selection methods. In these cases, selection only considers the rank of individuals rather than their actual raw fitness values. As such, the model in equation 1 is slightly modified to illustrate that it is some *ranking* function which is of interest to. A single population, coevolutionary algorithm can be described by a dynamical system in which the map function is defined as a composition of some variation operation (\mathcal{M}), selection operation (\mathcal{S}), and some ranking function (\mathcal{R}) as follows:

$$\mathcal{G} = \mathcal{M} \circ \mathcal{S} \circ \mathcal{R} \tag{1}$$

In a coevolutionary algorithm internal fitness assessment is subjective by its very nature. When comparisons to internal subjective fitness are important, we model CEA fitness assessment in an evolutionary game theoretic way (Ficici and Pollack, 2000; Wiegand et al., 2002). This means that, assuming an infinite population size and complete mixing (i.e.: each individual is paired with every other individual in the population, including itself), aggregate subjective val-

ues for genotypes (their *utility*) can be obtained as follows:

$$\vec{u} = A\vec{x} \quad (2)$$

Translating to the game-theoretic terminology, genotypes represent playing strategies and A specifies the *payoff matrix* that describes what each strategy gets when it plays some other strategy. $\vec{x} \in \Delta^n$ represents a point in the unit simplex. This means that $\forall i : x_i \in [0, 1]$ and $\sum_i x_i = 1$. In other words, a strategy receives a utility which is a weighted sum of the payoffs it receives from playing the other strategies in the population.

4 Preliminaries

We begin our preliminary definitions with a more formal treatment of the term *objective*. An objective measure is commonly thought of as a fitness function f which takes a genotype and assigns it a unique value. But because our model employs ranking functions \mathcal{R} rather than absolute fitness functions, we must establish what objectivity means with regard to ranking.

Definition 1 A ranking function \mathcal{R} is **objective** if, for any population vectors \vec{x} , it always establishes the same ordering among all individuals in an infinite population regardless of the proportions in \vec{x} .

This is not really a redefinition of objectivity: note that if there is an objective ranking function \mathcal{R} , then there trivially exists an objective fitness measure f : simply order all genotypes by rank, then assign them their rank as fitness.

Definition 2 A ranking function \mathcal{R} is **existentially-objective** if, for any population vectors \vec{x} , it always establishes the same ordering among individuals in an infinite population which have non-zero proportions in \vec{x} . An objective function is *existentially-objective*.

Note the crucial difference between these two definitions. Because a ranking function can be mapped to an absolute fitness function f , the first definition implies an objective fitness of an individual that is invariant over any population context — including populations that do not actually contain the individual. A truly objective fitness measure allows us to directly compare the results between techniques, or chart how fitness progresses over the course of the run.

The second definition does not imply this: instead, it only guarantees that individuals existing in a population will always be ranked the same way relative to one another. If the evolutionary system uses a non-parametric selection procedure, then this second definition is sufficient to guarantee that individuals with a coevolved subjective fitness measure will be selected in the exact same way that they would

if they were selected with the objective ranking function \mathcal{R} . That is, the coevolutionary procedure is at least guaranteed to follow the same trajectory as *some* traditional evolutionary procedure which uses nonparametric selection.

4.1 Relationships to Payoff

It will be helpful to establish some basic definitions of properties of payoff matrices, as well as properties of some kinds of relationships between measures and payoffs. To begin with, in our model payoff values for contests between different strategies are described by a *payoff matrix* (denoted A_{ij} , where i and j indicate specific genotypes).

One payoff property on which the coevolution community has focused has been the issue of *intransitivity*. Loosely speaking, a matrix is transitive if and only if there are no cycles in its payoff matrix. It turns out that intransitivity is not the only important attribute of payoff matrices which make eliciting objective measures difficult, but it is a good place to start. Below we formally define *weak transitivity* and its stricter form, *strong transitivity*. Strong transitivity will be used immediately following in the paper; weak transitivity will be used later on.

Definition 3 A payoff matrix A is **weakly transitive** iff for any distinct i, j, k the following holds:

$$\begin{aligned} (A_{ij} > A_{ji} \wedge A_{jk} > A_{kj} &\longrightarrow A_{ik} > A_{ki}) \wedge \\ (A_{ij} > A_{ji} \wedge A_{jk} = A_{kj} &\longrightarrow A_{ik} > A_{ki}) \wedge \\ (A_{ij} = A_{ji} \wedge A_{jk} > A_{kj} &\longrightarrow A_{ik} > A_{ki}) \wedge \\ (A_{ij} = A_{ji} \wedge A_{jk} = A_{kj} &\longrightarrow A_{ik} = A_{ki}) \end{aligned}$$

Definition 4 A payoff matrix A is **strongly transitive** iff A is weakly transitive and for any distinct $i, j : A_{ij} \neq A_{ji}$.

Since much of our discussion centers around the notion of rank, another important concept is one of *rank equivalence*. Informally, two measures are rank-equivalent if they order genotypes in the same way.

Definition 5 Two fitness measures, f and g are **rank equivalent**, $f \equiv_R g$ if and only if there exists an ordering of f and of g over the possible genotypes such that $\forall j, k$ genotypes: $(f_j > f_k \iff g_j > g_k)$

5 Transitivity is Insufficient

Even if there is a completely transitive ordering among individuals (and thus an objective measure), it turns out that such an ordering is not sufficient to guarantee that the CEA will follow a traditional EA trajectory with respect to this objective measure. Watson and Pollack (2001) demonstrated that loss of adaptive gradient can occur in

two-population models even when they are fully transitive. However, this phenomenon may be symptomatic of the use of a multi-population model. Even so, it still turns out that neither weak nor strong transitivity is a sufficient condition to guarantee that a coevolutionary system will have the same dynamics as an evolutionary system, much less that an objective measure exists which is rank-equivalent to the CEA's subjective measure. For example, consider a zero-sum population with the following payoff matrix A and proportion vector \vec{x} :

$$A = \begin{bmatrix} 0 & 6 & 3 \\ -6 & 0 & 18 \\ -3 & -18 & 0 \end{bmatrix} \quad \vec{x} = \begin{bmatrix} x_i \\ x_j \\ x_k \end{bmatrix}$$

This system is clearly transitive. i beats j and j beats k and i beats k . Now, imagine that $x_i = \frac{1}{2}, x_j = \frac{1}{2}, x_k = 0$. In this situation, the subjective fitness (utility) of genotype i is 3, and the utility of j is -3. But, now imagine that $x_i = \frac{1}{3}, x_j = \frac{1}{3}, x_k = \frac{1}{3}$. Here, the utility of i is 3, but the utility of j is 4! Since the ranks of i and j depend on their proportions relative to k , there is *no objective fitness function* that is rank-equivalent to the subjective fitness function described in this system. Just saying that in a transitive matrix i beats j is not sufficient to argue that the coevolutionary subjective fitness has a rank equivalence to some fitness function where $f_i > f_j$, because what matters is not that i beat j but by *how much* i beat j . The spread is important.

6 Guaranteeing an Objective Measure

How might we go about guaranteeing that an objective measure exists which is rank-equivalent to our CEA's subjective measure? One of the complications in making this guarantee is that the measure for genotype i must be independent of population context: that is, it must be the same regardless of whether or not i is actually in presently *in the population*. Because of this issue, we can establish the hard minimum requirement for a payoff matrix A necessary for there to exist an objective measure f that is rank equivalent to the subjective utility measure \vec{u} .

Theorem 1 *Given a payoff matrix A with the corresponding subjective utility measure $\vec{u} = A\vec{x}$, if there exists an objective measure f that is rank equivalent to \vec{u} , the following condition must be true:*

$$\forall i, j : (f_i > f_j \iff \forall k : A_{ik} > A_{jk})$$

Proof First we establish that this is a necessary condition for the corners of the simplex, that is, those vectors \vec{x}_{k^*} representing a population consisting entirely of a sole genotype k . In this situation, $u_i = A_{ik}$ and $u_j = A_{jk}$. Thus for \vec{u} to be rank equivalent with f , we must have: $(f_i > f_j \iff u_i = A_{ik} > u_j = A_{jk})$

Second, we establish that this is a sufficient condition for all remaining possible values of the simplex \vec{x}_{k^*} . Since the simplex corners form basis vectors for the simplex, \vec{x} is simply a linear function of various simplex corners \vec{x}_{k^*} . As all elements in \vec{x} are non-negative, and at least one is non-zero, then it must be that if the corners of the simplex obey the condition above, then for any \vec{x} : $(f_i > f_j \iff u_i = \sum_k A_{ik}x_k > u_j = \sum_k A_{jk}x_k)$ ■

A practical example which clearly meets this property is the *greater than game* (Watson and Pollack, 2001; Panait and Luke, 2002). In the latter study, competitions were scored by scaling the difference between two objective functions, while in the former competitions are scored with 0 or 1 depending on whether a strategy scores higher than another. Although not strictly the same, these relationships can be more or less generalized to a linear relationship without affecting the fact that an external objective measure exists for the system. First we define this relationship more formally.

Definition 6 *Given a linear transformation, $A_{ij} = \alpha f_i + \beta f_j$, the internal subjective utility u is **linearly related** to an objective function f , $u \sim_L f$, if the transitive payoff matrix A is produced using this transformation.*

It's useful to know that such a relationship exists, and thus that there exists some objective measure that is rank-equivalent with the internal subjective utility measure. We will learn more about this in the next couple of sections. First we prove that this generalized form of the *greater than game* does indeed have this rank-equivalence property.

Theorem 2 *Given a transitive payoff matrix A produced by a linear transformation of some arbitrary objective function, $u \sim_L f \rightarrow u \equiv_R f$ (that is, u is rank equivalent to f) as long as $\alpha > 0$.*

Proof First note that the utility for a given genotype, i , can be obtained through algebraic expansion,

$$\begin{aligned} u_i &= (A\vec{x})_i = \sum_{j=1}^n A_{ij}x_j \\ &= \sum_{j=1}^n (\alpha f_i + \beta f_j)x_j = \sum_{j=1}^n \alpha f_i x_j + \sum_{j=1}^n \beta f_j x_j \\ &= \alpha f_i + \beta \sum_{j=1}^n f_j x_j \end{aligned}$$

It suffices now to prove that $u_k > u_i \rightarrow f_k > f_i$. Starting with the first part, $u_k > u_i$, so

$$\begin{aligned} \alpha f_k + \beta \sum_{j=1}^n f_j x_j &> \alpha f_i + \beta \sum_{j=1}^n f_j x_j \\ \alpha f_k &> \alpha f_i \end{aligned}$$

Since $\alpha > 0$, we can divide through on both sides without reversing the inequality, such that $f_k > f_i$. ■

6.1 Analyzing Coevolutionary Algorithms

It is clear that having an objective measure of genotypes being coevolved will allow researchers to confirm which of two algorithms has performed better than the other. More specifically, we will have a reasonable external measurement with which to judge the quality of solutions found in general. Perhaps more importantly, however, having a valid objective measure will help us detect when apparently uninteresting subjective dynamics correspond to real improvement (so-called *arms race conditions*). Further, pathological dynamics such as cycling or mediocre stability can be more easily assessed (Watson and Pollack, 2001).

Another use of this knowledge is the reverse question: if you are using an external measure, does it really measure the objective on which the algorithm is working? If it does not, then the objective measure is inappropriate and potentially harmful, since it is likely to mislead its user regarding the true nature of the problem being solved by the system.

6.2 Dynamical Equivalence to Traditional EAs

An obvious question at this point is: if an objective function exists that can be used to chart the true progress of the algorithm, isn't the algorithm really non-coevolutionary? This is a reasonable question, and depending on one's point of view and the circumstances involved, the answer might well be that it is not a coevolutionary algorithm at all.

Since the term "coevolutionary algorithm" has more to do with the concepts involved in constructing the evolutionary system, it is perhaps more constructive to continue to think of algorithms with such properties as CEAs, but ask the question of whether or not their dynamics are equivalent to some non-coevolutionary algorithm. In fact, when true rank-equivalence of the internal subjective fitness measure exists, the coevolutionary algorithms of the type we have been discussing are dynamically equivalent to an EA solving some problem with the related objective measure.

Theorem 3 *A single population CEA under complete mixing, employing a non-parametric selection method using the internal subjective utility $\vec{u} = A\vec{x}$ is dynamically equivalent to an evolutionary algorithm with the same selection method, using the objective function f , if $u \equiv_R f$.*

Proof Given the general model of a coevolutionary or evolutionary algorithm shown in equation 1, it suffices to prove that the result of selection will be the same, $S(A\vec{x}, \vec{x}) = S(f(\vec{x}), \vec{x})$. However, since the selection method is non-parametric, we can see more specifically that $S(R(A\vec{x}), \vec{x}) = S(R(f(\vec{x})), \vec{x})$, where R is a function which assigns an ordering to the genotypes based on the ranks of their measures. Because $u \equiv_R f$, we know that $R(A\vec{x}) = R(f(\vec{x}))$ by the definition of *rank equivalence*. ■

7 Beyond Objectivity

It turns out that one can guarantee that a coevolutionary system follows the same dynamical trajectory as some traditional EA, even if one does not know the relevant objective measure. All that is necessary is to be able to demonstrate that the coevolutionary system will order members of a population in the same way that a traditional EA would. This requires that we find an *existentially objective* ranking function which is rank-equivalent to the subjective utility measure. We have not yet discovered minimum necessary and sufficient bounds for this situation: but we can and will state certain necessary conditions. We omit the proof for space reasons, though it will be provided in the future.

Theorem 4 *Given a payoff matrix A with the corresponding subjective utility measure $\vec{u} = A\vec{x}$, imagine that there exists an existentially-objective ranking function \mathcal{R} which is rank-equivalent to \vec{u} . Then the following must be true:*

$$\forall i, j, k : (\mathcal{R}_i > \mathcal{R}_j \longrightarrow (A_{ij} > A_{jj} \wedge A_{ii} \geq A_{ji}) \vee (A_{ij} \geq A_{jj} \wedge A_{ii} > A_{ji}))$$

7.1 A Practical Example

As it so happens, there exist certain common CEA features which, if all true for a given CEA, guarantee that the system behaves like an EA even in the absence of a true objective function. The features are as follows: first, the payoff matrix must be at least weakly transitive; second, the matrix must be constant-sum (in fact, the requirement is slightly looser than this); third, the matrix must be *monotone*, that is, if i is better than j , and j is better than k , then i must beat k by at least as much as i beats j and by at least as much as j beats k . We begin with some definitions of matrix features, then prove this fact.

Definition 7 *Payoff matrix A is **monotone** if and only if for any i, j, k : $A_{ij} \geq A_{ji} \wedge A_{jk} \geq A_{kj} \longrightarrow A_{ik} \geq \max(A_{ij}, A_{jk})$. Note that i, j, k need not be distinct.*

Definition 8 *Payoff matrix A is a **constant sum plus matrix** if it is constant sum (let the constant be C) on all elements except possibly the diagonal. On the diagonal, the following hold:*

$$A_{ij} = A_{ji} = \frac{1}{2}C \longrightarrow A_{ii} = A_{jj} = \frac{1}{2}C$$

$$A_{ij} > A_{ji} \longrightarrow (A_{ij} > A_{jj} \wedge A_{ii} \geq A_{ji}) \vee (A_{ij} \geq A_{jj} \wedge A_{ii} > A_{ji})$$

Theorem 5 *Let A be a monotone, weakly transitive, constant-sum-plus payoff matrix. Then given any genotypes i and j , where proportions x_i and x_j are non-zero, then using a monotone, weakly transitive, constant-sum-plus payoff matrix A , iff $A_{ij} > A_{ji}$ then $u_i > u_j$. Further, iff $A_{ij} = A_{ji}$ then $u_i = u_j$.*

Again the proof is omitted due to space restrictions and will be included in a publication in the near future.

In the preceding theorem, because A is weakly transitive, there is clearly an ordering among all the genotypes. But because the proportions of the two genotypes i and j must be nonzero in the population in order to be ranked properly, we cannot stipulate that A 's transitive ordering acts as a truly objective measure: it is dependent in some sense on population context. Instead the ordering is an existentially objective measure. One consequence of this is the fact that we do not actually need to do full mixing in order to evolve the population, if we use tournament selection: instead, we can perform lazy evaluation at selection-time to determine which member of the tournament beats the others, and select that member.

7.2 Dynamical Equivalence Without Objectivity

In most evolutionary game theory, the assumption is made that A and \vec{x} and \vec{u} are finite in size, that is, that there is a finite number of possible genotypes. This is primarily because it is problematic to establish a distribution \vec{x} over an infinite population with an infinite set of genotypes. However, if we assume a finite population size, we may relax this assumption, at least for the previous two theorems, because although \vec{x} will be infinitely long, we know it can have only a finite number of nonzero elements.

Thus we can apply the previous two theorems in two different ways. First, they apply to the traditional game theory approach with infinite populations of finite genotypes. Here, with enough work, it is always *possible* to discover the existentially objective ranking function, because the number of genotypes is fixed. But more interestingly, the theorems also apply to an approach with finite populations of (countably) infinite genotypes. This is important because it gives us practical conditions under which we may make guarantees of EA-like dynamics in a real-world, finite population, even with countably infinite representations such as graph or tree structures. With these representations, we may never be able to discover the actual existentially objective ranking function in full (because the genotypes are infinite), but we know how this function would rank any two individuals, and we also know that the CEA is equivalent to an EA in applying this function.

8 Conclusions and Future Work

In this paper we discussed the history behind objective and external ranking functions, and the difficulty which coevolutionary's subjective function presents in obtaining useful external measures or in guaranteeing that the CEA will have dynamics similar to an EA's. We then established minimum conditions for a CEA to have a truly objective

measure, then gave a practical example of a CEA with an objective measure. Last, we relaxed the conditions of objectivity, and still demonstrated situations where the CEA would have an EA dynamics, even in the absence of an obvious objective measure.

Our technique as described here applies to single-population coevolution, with full mixing, and a non-parametric selection function such as truncation selection or tournament selection. As future work, we are examining three questions: first, what additional conditions can be placed upon *intransitive* situations such that they exhibit reasonable evolutionary dynamics? Second, what may be said in the partial-mixing situations common in real-world CEA work? Third, how can we extend these results to apply to multiple-population models?

References

- D. Cliff and G. F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *Proceedings of the Third European Conference on Artificial Life*, pages 200–218. Springer-Verlag, 1995.
- S. Ficici and J. Pollack. Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In A. et al, editor, *Proceedings of the Sixth International Conference on Artificial Life*, Cambridge, MA, 1998. MIT Press.
- S. Ficici and J. Pollack. A game-theoretic approach to the simple coevolutionary algorithm. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. Merelo, and H.-P. Schwefel, editors, *Proceedings of the Sixth International Conference on Parallel Problem Solving from Nature (PPSN VI)*, pages 467–476. Springer-Verlag, 2000.
- L. Panait and S. Luke. A comparison of two competitive fitness functions, 2002. Submitted to GECCO 2002.
- C. Rosin and R. Belew. New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29, 1997.
- M. Vose. *The Simple Genetic Algorithm*. MIT Press, 1999.
- R. Watson and J. Pollack. Coevolutionary dynamics in a minimal substrate. In L. Spector, editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2001*, pages 702–709. Morgan Kaufmann, 2001.
- R. P. Wiegand, W. Liles, and K. De Jong. Analyzing cooperative coevolution with evolutionary game theory. In D. Fogel, editor, *Proceedings of CEC 2002*. IEEE Press, 2002. (To appear).