
Guaranteeing Coevolutionary Objective Measures

Sean Luke

Department of Computer Science
George Mason University
Fairfax, VA 22030
sean@cs.gmu.edu
<http://www.cs.gmu.edu/~sean/>

R. Paul Wiegand

Department of Computer Science
George Mason University
Fairfax, VA 22030
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 several techniques have been proposed to use external or objective measures to help in analysis, there are clearly properties of fitness payoff, like intransitivity, for which these techniques are ineffective. We feel that a principled approach to this problem is to first establish the theoretical bounds to guarantee objective measures in one CEA model; from there one can later examine the effects of deviating from the assumptions made by these bounds. To this end, 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.

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 of the subjective nature

of fitness 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. This differs 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 our focus 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. Because it is tractable, complete mixing is a standard assumption in theoretical analyses of coevolution; however we recognize that a number of coevolutionary methods (single-elimination tournament, incomplete mixing, fitness sharing, etc.) do not necessarily fit this model in practice. However first establishing minimum bounds for guaranteeing objective measures in some valid CEA, we form a principled foundation from which to later "jump off" and explore what happens when one deviates from these bounds: for example, by gradually loosening transitivity restrictions, or by moving from a full population-wide tournament to smaller tournaments.

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 first understand 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 to include conditions which are not formally objective, but still guarantee that the CEA will behave dynamically just as some related traditional EA.

In this paper, we provide first-draft proofs of theorems.

2 TYPES OF COEVOLUTION PROGRESS MEASURES

Most attempts at competitive coevolution face the dilemma of how to chart the algorithm's progress in an objective fashion. 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 are 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. Additionally, 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 ones 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 for some

game, such that the rules are represented by binary strings. One could provide an objective measure for an individual which is a simple uniteration 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

The coevolutionary model we use will focus solely on one-population coevolution. There is some controversy over whether single-population coevolution should be called *coevolution*, and this deserves some small attention here. In a biological sense, coevolution by definition requires more than one population. What we call one-population coevolution, biologists would call *evolution*, a term unfortunately already taken to describe traditional EAs, which in reality do not have the game-theoretic aspects central to biological evolution. This leaves one-population coevolution in the confusing situation of having no good name. Thus while many researchers studying CEAs limit the term *coevolution* to situations in which there are distinct, non-interbreeding populations (Potter, 1997; Rosin and Belew, 1997), there are many which do not make this distinction (Juillé, 2001; Ficici and Pollack, 2000; Luke et al., 1998). We are tempted to refer to one-population coevolution as simply a *competitive fitness function* (Angeline and Pollack, 1993) except that this term is more commonly used to describe both one-population and multi-population competitive environments. At any rate, it is clear that both one-population and multi-population “coevolution” share certain game-theoretic properties, and as such may have more in common with one another than either does to traditional EAs. Therefore for the purposes of this paper, we define single population models using competitive fitness methods to be coevolutionary algorithms.

Since we will be focusing entirely on single population coevolutionary algorithms, our high-level framework 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 based on fitness, 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 us in terms of our questions regarding CEA dynamics. Thus 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:

$$\mathcal{G} = \mathcal{M}(\mathcal{S}(\mathcal{R}(\vec{x}), \vec{x})) \quad (1)$$

...where \vec{x} represents the population vector of genome proportions.

However, 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 values for genotypes (their *utility*) can be obtained as follows:

$$\vec{u} = A\vec{x} \tag{2}$$

Translating to the game-theoretic terminology, genotypes represent different playing strategies and A specifies the *payoff matrix* that describes what each strategy gets when it plays some other strategy. Since populations are infinite, $\vec{x} \in \Delta^n$ represents a point in the unit simplex. This means that all component values of the vector are between 0 and 1 inclusive, and that all components of the vector sum to 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.

Unlike much of the literature, our definitions of \vec{u} , A , and \vec{x} permit these vectors and matrices to be infinite in length, so long as \vec{u} consists entirely of finite-value elements. As \vec{x} is in the unit simplex, this can occur if all elements in A are less than some finite number z . It can also occur if the number of nonzero elements in \vec{x} is finite (that is, the population is finite in size). This extends the theory to cover not only infinite populations of finite genotypes, but also, among other things, the more practical finite populations of infinite genotypes.

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




Figure 1 Valid weakly transitive cases, not including rotations. These correspond to the four clauses stated in Definition 3. The leftmost figure (and its rotations) is the only valid strongly transitive case.




Figure 2 Invalid weakly transitive cases, not including rotations.

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 \mathcal{R} . That is, the coevolutionary procedure is at least guaranteed to follow the same trajectory as *some* traditional evolutionary procedure which uses a nonparametric selection mechanism.

If the number of genotypes were finite, then the second definition might be construed to imply the first: if there exists a ranking function, then it must always be discoverable. However if there are an infinite number of genotypes, we might not be able to say for certain what “rank” an individual holds in the infinite genotype set, but still could specify which of two individuals is ranked higher. Indeed, this is a common situation in practice.

4.1 RELATIONSHIPS TO PAYOFF

Although it is not obvious how to objectively monitor single-population coevolutionary algorithms in general, we will show that under certain minimum bounds on the payoff mechanism, there exist specific coevolutionary algorithms which can be objectively monitored.

Before we can begin our discussion in more detail, 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:

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

Figure 1 shows valid weakly transitive cases, and Figure 2 shows invalid cases.

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*. Loosely speaking, 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 population with the zero-sum 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 can be *no objective fitness function* that is rank-equivalent to the subjective fitness function described in this system. The reason for this is simple: 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

So 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 such a system. First let us define such a relationship more formally.

Definition 6 *Given a linear transformation, $A_{ij} = \alpha f_i + \beta f_j + \gamma$, 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,

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

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

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

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

As we've already established there are many points of study regarding monitoring the progress of coevolutionary algorithms. Two among these are comparing algorithms and diagnosing algorithm dynamics. In both cases, having a valid objective measure will be of use.

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 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*) or not. Moreover, 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 have an external measure you are using, 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

Of course the most obvious question at this point is: if there exists an objective function that can be used to chart the true progress of the algorithm, is the algorithm essentially non-coevolutionary? This is a fair and reasonable argument. But 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 state certain necessary conditions:

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 \bar{u} . Then the following must be true:

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

Proof Note that we are only showing necessary conditions, not sufficient ones. To do this, it is required only to establish necessity for *some* value of \vec{x} . We begin by establishing necessary bounds using the proportion vector family \vec{x}_{ij^*} which is defined as those vectors with non-zero proportions of the genotypes i and j , but zero proportions for all other genotypes. In this situation, the subjective fitness of i is just $A_{ij}x_j + A_{ii}x_i$. Similarly, the subjective fitness of j is $A_{jj}x_j + A_{ji}x_i$. Let \mathcal{R} be the existentially objective ranking function, if it exists. To be rank equivalent with the subjective measure, the following must then be true:

$$\begin{aligned} (\mathcal{R}_i > \mathcal{R}_j &\iff A_{ij}x_j + A_{ii}x_i > A_{jj}x_j + A_{ji}x_i) \wedge \\ (\mathcal{R}_i = \mathcal{R}_j &\iff A_{ij}x_j + A_{ii}x_i = A_{jj}x_j + A_{ji}x_i) \end{aligned}$$

Since x_i and x_j are the only nonzero values, then $x_j = 1 - x_i$. Hence these clauses may be rewritten as:

$$\begin{aligned} (\mathcal{R}_i > \mathcal{R}_j &\iff (A_{ij} - A_{jj})(1 - x_i) > (A_{ji} - A_{ii})x_i) \wedge \\ (\mathcal{R}_i = \mathcal{R}_j &\iff (A_{ij} - A_{jj})(1 - x_i) = (A_{ji} - A_{ii})x_i) \end{aligned}$$

Collecting x_i , we have:

$$\begin{aligned} (\mathcal{R}_i > \mathcal{R}_j &\iff (A_{ij} - A_{jj} > (A_{ji} + A_{ij} - A_{ii} - A_{jj})x_i) \wedge \\ (\mathcal{R}_i = \mathcal{R}_j &\iff (A_{ij} - A_{jj} = (A_{ji} + A_{ij} - A_{ii} - A_{jj})x_i) \end{aligned}$$

x_i may range between, but not be equal to, 0 and 1. The second clause can be true over all such x_i only when $A_{ij} = A_{jj}$ and $A_{ji} = A_{ii}$, because ultimately the second clause boils down to $a = bx_i$, and this equation can only be true for any x_i when $a = b = 0$. A similar argument may be made for the first clause: as x_i approaches but does not equal 1, it must be the case that either $A_{ji} > A_{ii}$ or that $A_{ji} = A_{ii}$ and $A_{ij} > A_{jj}$. As x_i approaches 0 (but not quite 0), it likewise must be the case that either $A_{ij} > A_{jj}$ or that $A_{ij} = A_{jj}$ and $A_{ji} > A_{ii}$. Together, we have the requirement that either $A_{ji} > A_{ii}$ and $A_{ij} \geq A_{jj}$ or that $A_{ji} \geq A_{ii}$ and $A_{ij} > A_{jj}$.

We have established the first clause in our theorem, and most of the second clause. For the remainder of the second clause, note that for rank equivalence, the following is true in general:

$$(\mathcal{R}_i = \mathcal{R}_j \iff A_{ij}x_j + A_{ii}x_i + \sum_{k \neq i, j} A_{ik}x_k = A_{jj}x_j + A_{ji}x_i + \sum_{k \neq i, j} A_{jk}x_k)$$

Rearranging the right half of this clause, we get

$$(A_{ij} - A_{jj})x_j + (A_{ii} - A_{ji})x_i = \sum_{k \neq i, j} (A_{jk} - A_{ik})x_k$$

We have established that if $\mathcal{R}_i = \mathcal{R}_j$ then it is necessary for $A_{ij} = A_{jj}$ and for $A_{ii} = A_{ji}$, so this last rearrangement further simplifies to just $\sum_{k \neq i, j} (A_{jk} - A_{ik})x_k = 0$. Over all values of x_k ranging from 0 to 1 not inclusive, the only way this can be true is if $\forall k \neq i, \neq j : A_{ik} = A_{jk}$. Since $A_{ij} = A_{jj}$ and $A_{ji} = A_{ii}$, we have $\forall k : A_{ik} = A_{jk}$. ■

7.1 A PRACTICAL EXAMPLE

Previously we gave necessary conditions to guarantee that the CEA behaves like an EA even in the absence of a true objective function. Here we demonstrate a practical example which is *sufficient* to make this guarantee. The features of the CEA 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 the result.

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

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:*

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

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

Lemma 1 *If $a \geq b$ and $a + b = c$ then $a \geq \frac{1}{2}c$ and $b \leq \frac{1}{2}c$.*

Proof $a + a \geq c$ so $a \geq \frac{1}{2}c$; and $b + b \leq c$ so $b \leq \frac{1}{2}c$. ■

Lemma 2 *A constant sum matrix is a constant sum plus matrix.*

Proof Whenever $A_{ij} = A_{ji} = \frac{1}{2}C$ then $A_{ii} = A_{jj} = \frac{1}{2}C$. Also, if $A_{ij} > A_{ji}$ and $A_{ij} + A_{ji} = C$ then from Lemma 1, $A_{ij} \geq \frac{1}{2}C$ and $A_{ji} \leq \frac{1}{2}C$. Since $A_{ii} = \frac{1}{2}C$, then $A_{ii} \geq A_{ji}$. Since $A_{jj} = \frac{1}{2}C$, then $A_{ij} \geq A_{jj}$. *Note:* it follows that a zero-sum matrix is also a constant sum plus matrix. ■

Lemma 3 *If $a = \max(b, c)$ and $b = \max(a, c)$ then $c \leq a = b$.*

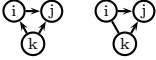
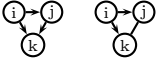
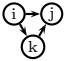
Proof Suppose that $a > b$. In this case, since $a = \max(b, c)$, therefore $a = c$. But since $b = \max(a, c)$, then $b = a$, which would be a contradiction. Similar results follow from $a < b$. With $a = b$, there is no contradiction, and further if $a = b$ and $a = \max(b, c)$, then $c \leq a$. ■

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 $u_i = u_j$.*

Proof By definition, $f_i = \sum_k A_{ik}x_k$ and $f_j = \sum_k A_{jk}x_k$. We divide the proof into two parts. First, we show that for any distinct genotype k not i or j , if $A_{ij} > A_{ji}$ then $A_{ik} \geq A_{jk}$ hence $A_{ik}x_k \geq A_{jk}x_k$. Since i and j are both in the population, then x_i and x_j are both nonzero, and so because A is a constant-sum-plus matrix, we know that

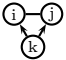
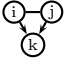
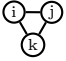
$A_{ij}x_j > A_{jj}x_j$ and $A_{ii}x_i \geq A_{ji}x_i$ or $A_{ij}x_j \geq A_{jj}x_j$ and $A_{ii}x_i > A_{ji}x_i$. At any rate, for each and every k , when $A_{ij} > A_{ji}$, then $A_{ik}x_k \geq A_{jk}x_k$, and further there exists a k such that $A_{ik}x_k > A_{jk}x_k$, so $f_i = \sum_k A_{ik}x_k > f_j = \sum_k A_{jk}x_k$.

To show this, note that for a given k not i or j , there are three general cases that do not violate weakly transitive constraints.

1.  $A_{ij} > A_{ji} \wedge A_{ik} \leq A_{ki} \wedge A_{jk} < A_{kj}$.
Then $A_{kj} \geq \max(A_{ki}, A_{ij})$, so $A_{kj} \geq A_{ki}$, so $C - A_{kj} \geq C - A_{ki}$, thus $A_{jk} \leq A_{ik}$.
2.  $A_{ij} > A_{ji} \wedge A_{ik} > A_{ki} \wedge A_{jk} \geq A_{kj}$.
Then $A_{ik} \geq \max(A_{ij}, A_{jk})$, so $A_{ik} \geq A_{jk}$.
3.  $A_{ij} > A_{ji} \wedge A_{ik} > A_{ki} \wedge A_{jk} < A_{kj}$.
Let C be the constant sum of matrix A . By Lemma 1, $A_{jk} < \frac{1}{2}C < A_{ik}$.

Second, we will show that any distinct genotype k not i or j , if $A_{ij} = A_{ji}$ then $A_{ik} = A_{jk}$ hence $A_{ik}x_k = A_{jk}x_k$. Then by the definition of a constant sum matrix, $A_{ii} = A_{jj} = A_{ij} = A_{ji}$. Thus $A_{ii}x_i = A_{ji}x_i$ and $A_{ij}x_j = A_{jj}x_j$ so for each and every k , when $A_{ij} = A_{ji}$, then $A_{ik}x_k = A_{jk}x_k$ so $f_i = \sum_k A_{ik}x_k = f_j = \sum_k A_{jk}x_k$.

There are three cases that do not violate weakly transitive constraints:

4.  $A_{ij} = A_{ji} \wedge A_{ik} < A_{ki} \wedge A_{jk} < A_{kj}$.
Then $A_{ki} = \max(A_{kj}, A_{ij} = A_{ji})$ and $A_{kj} = \max(A_{ki}, A_{ij} = A_{ji})$, so from Lemma 3, $A_{ki} = A_{kj}$ so $A_{ik} = A_{jk}$.
5.  $A_{ij} = A_{ji} \wedge A_{ik} > A_{ki} \wedge A_{jk} > A_{kj}$.
Then $A_{ik} = \max(A_{jk}, A_{ij} = A_{ji})$ and $A_{jk} = \max(A_{ik}, A_{ij} = A_{ji})$, so from Lemma 3, $A_{ik} = A_{jk}$.
6.  $A_{ij} = A_{ji} \wedge A_{ik} = A_{ki} \wedge A_{jk} = A_{kj}$.
Then from the definition of a constant plus matrix it necessarily follows that $A_{ik} = A_{jk}$.

This proves that if $A_{ij} > A_{ji}$ then $f_i > f_j$; and if $A_{ij} = A_{ji}$ then $f_i = f_j$. To show that the implication is also an equivalence, note that these are the only possible cases: either A_{ij} is equal to A_{ji} , or one is greater than the other, and either f_i is equal to f_j or one is greater than the other. Further this is a one-to-one, onto mapping: the A and f relations each use a unique i and j pair. ■

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, the existentially objective ranking function is always *possible* to determine, with enough work, 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 is applied 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

- Angeline, P. and J. Pollack (1993). Competitive environments evolve better solutions for complex tasks. In S. Forest (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms (ICGA)*, San Mateo, CA, pp. 264–270. Morgan Kaufmann.
- Cliff, D. and G. F. Miller (1995). Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *Proceedings of the Third European Conference on Artificial Life*, pp. 200–218. Springer-Verlag.
- Ficici, S. and J. Pollack (1998). Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In A. et al (Ed.), *Proceedings of the Sixth International Conference on Artificial Life*, Cambridge, MA, pp. 238–247. MIT Press.
- Ficici, S. and J. Pollack (2000). 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 (Eds.), *Proceedings of the Sixth International Conference on Parallel Problem Solving from Nature (PPSN VI)*, pp. 467–476. Springer-Verlag.
- Juillé, H. (2001). Basic concepts in coevolution. Presentation at GECCO-01 Coevolutionary Workshop.
- Luke, S., C. Hohn, J. Farris, G. Jackson, and J. Hendler (1998). Co-evolving soccer softbot team coordination with genetic programming. In *RoboCup-97: Robot Soccer World Cup I (Lecture Notes in AI No. 1395)*, Berlin, pp. 398–411. Springer-Verlag.
- Panait, L. and S. Luke (2002). A comparison of two competitive fitness functions. Submitted to GECCO 2002.
- Potter, M. (1997). *The Design and Analysis of a Computational Model of Cooperative CoEvolution*. Ph. D. thesis, George Mason University, Fairfax, Virginia.
- Rosin, C. and R. Belew (1997). New methods for competitive coevolution. *Evolutionary Computation* 5(1), 1–29.
- Vose, M. (1999). *The Simple Genetic Algorithm*. MIT Press.
- Watson, R. and J. Pollack (2001). Coevolutionary dynamics in a minimal substrate. In L. Spector (Ed.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2001*, pp. 702–709. Morgan Kaufmann.
- Wiegand, R. P., W. Liles, and K. De Jong (2002). Analyzing cooperative coevolution with evolutionary game theory. In D. Fogel (Ed.), *Proceedings of CEC 2002*. IEEE Press. (To appear).