Conditions for the Robustness of Compositional Coevolution

Dagstuhl 2008: Theory of Evolutionary Algorithms

R. Paul Wiegand

Institute for Simulation & Training
University of Central Florida

Coevolutionary Algorithms

- Very similar to traditional EA methods:
  - Individuals encode aspect of potential solutions
  - They are altered during search by genetic operators
  - Search is directed by selection based on fitness

- But differ in fundamental ways:
  - Evaluation requires interaction between multiple individuals
  - Interacting individuals may reside in same population or in different populations
  - Evokes notions of cooperation and competition in new ways
  - Some representation issues are unique to coevolution
  - Methods of evaluation are particularly important

Compositional Coevolutionary Algorithm

Consider a simple framework for algorithms designed to approach different kinds of compositional problems:

- Partition problem space into components
- Apply an EA to each component
- Assemble components for a complete solution (collaboration)
- Evaluate solution with objective function

CCEA Example

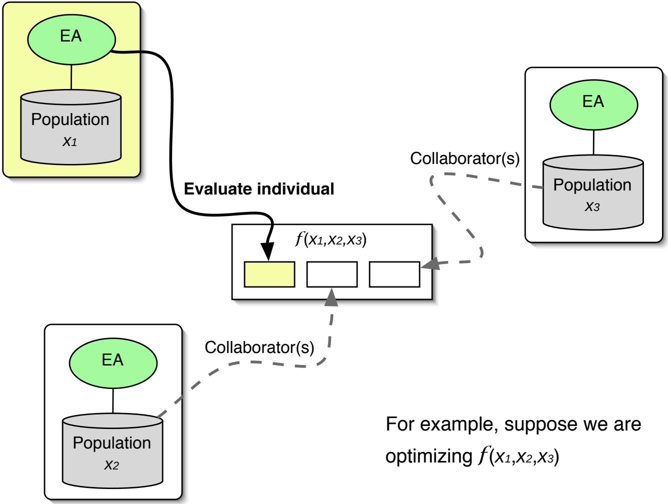

For example, suppose we are optimizing $f(x_1,x_2,x_3)$

Modeling the CCEA for Analysis

- Evolutionary Game Theoretic (EGT):
  - Discrete time dynamical systems model
  - Tracking *population state* as generations progress
  - Model fitness as the payoff of a game
  - Interested in properties of limit behaviors of system

- Model assumptions & properties
  - Two populations, each infinitely large
  - *Role symmetry* in the payoff ("cooperative")
  - Individuals interact in a variety of ways
  - No genetic search operators
  - Parents are selected via proportionate
  - All children survive
  - Populations evaluated & updated synchronously, in parallel

EGT Model Details

- Populations are ratios of genotypes
  - If there are n distinct genotypes in each population:
  - $\vec{x} \in \mathbb{R}^n,\ x_i \in [0,1]\ \wedge\ \sum_{i=1}^{n} x_i = 1$
  - $\vec{y} \in \mathbb{R}^n,\ y_i \in [0,1]\ \wedge\ \sum_{i=1}^{n} y_i = 1$
- Fitness modeled using a *payoff matrix*
  - Given a matrix A and...
  - Genotype i from x population and j from y:
  - $a_{ij} = f(i,j)$, whether scoring for x or for y

$$x_i' = \frac{(A\vec{y})_i}{\vec{x}^T A\vec{y}} \cdot x_i$$

$$y_j' = \frac{\left(A^T\vec{x}\right)_j}{\vec{x}^T A\vec{y}} \cdot y_j$$

EGT & Ideal Partnership

- Can use evolutionary game theory (EGT) to model CCEAs
  - Dynamical systems theory (population state trajectories)
  - Game theory (fitness as payoff matrix)
- What is "best" in compositional coevolution?
  - Perhaps: *Strategy that does the best when partnered with ideal partner strategy* ("ideal partnership")
- What does it mean to "optimize" a compositional problem?
  - Find the global optima in composite space
  - Global optima $\equiv$ ideal partnership

$$
\begin{aligned}
a^* &= \max\{A\} \\
f(x_i, y_j) &= a^*
\end{aligned}
$$

Ideal Partnership: (i, j)

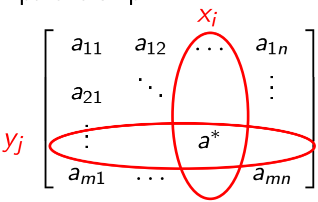

Motivations & Perspectives

- From analysis, we know:
  - Properties affecting CCEAs are complex
  - Can be suboptima that correspond with stable attracting f.p.
  - Model studies suggest suboptima may be *more* attractive
- But, experience suggests ...
  - CCEAs can be "fixed" to improve their optimization potential
  - CCEAs work well on many multiagent learning problems
  - Often we are less interested in strictly optimal team performance, but in high performing teams that are also robust
- Questions:
  - What do we mean by "fixed"?
  - What do we mean by "robust"?
  - **How can robustness property be shown more formally?**

What Does This Mean?

Global, static optimization of a compositional problem is *not* a good solution concept for a traditional compositional coevolutionary algorithm.

- We can "fix" the algorithm....
  - Bias toward ideal partnership (Panait *et al.*, 2004)
  - Apply Pareto & memory methods (Bucci & Pollack, 2005)

- But the traditional algorithm *seems* to do something useful
  - Performs well on many coadaptive multiagent learning tasks
  - Here, we want teams to be "relatively good" but "robust" to changes in member behaviors
  - May want to consider some kind of *robustness* solution concept
  - At first blush, this would be consistent with what is known about CCEAs (i.e., cumulative rewards ...)

What is “Robustness”?

- Term “robustness” is often ill-defined
- There are different, equally valid interpretations
- In general, how does one know what is meant?
- How does one know whether a given interpretation is *useful*?

Framework for Defining Robustness

Stolen liberally from the RA community

- Need two things:
  - Partitioning into two components (design & control)
  - Robustness criterion
- Ask: How do changes in control variables affect fixed design variables wrt the robustness criterion?

Notice: This describes a very general framework for instantiating different definitions of robustness

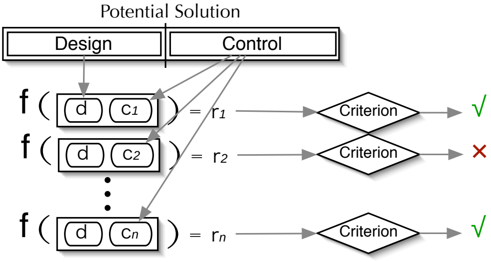

where,
$\{c_1, c_2, \ldots, c_n\} = C$
$d \in D$

A Framework for CCEAs

- We consider only a two-population CCEA, generalizing is straightforward
- We take the perspective of the first component, the second follows by symmetry

Definition

Let $S = \{s_1, s_2, \ldots, s_n\}$ and $T = \{t_1, t_2, \ldots, t_n\}$ represent the set of possible component values for the first and second population, respectively.

We call $\mathcal{C} : S \times T \mapsto \{0, 1\}$ a **robustness criterion**. We say that a particular $\bar{s}$ is **perfectly robust** over T with respect to the robustness criterion $\mathcal{C}_s$ if $\forall t_j \in T, \mathcal{C}_s\left(\bar{s}, t_j\right) = 1$.

We call $\mathcal{D}_s\left(\bar{s}, T\right) = \frac{\sum_j \mathcal{C}_s\left(\bar{s}, t_j\right)}{|T|}$ the **(non-parametric) degree of robustness** of $\bar{s}$ over T with respect to the robustness criterion $\mathcal{C}_s$.

- We can reverse the two components and apply the same logic to get $\mathcal{D}_t\left(t_j, S\right)$
- So, component value s_i is more robust than s_j if $\mathcal{D}_s\left(s_i, T\right) > \mathcal{D}_s\left(s_j, T\right)$
- In practice one hasn't access to the entire space of potential components, so some sampling method will be required

Translating to Coadaptive Multiagent Learning

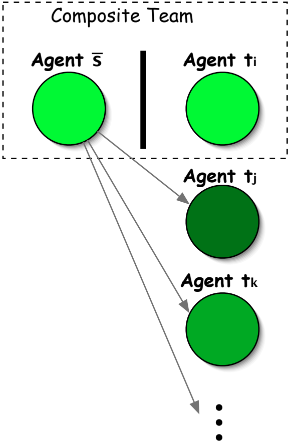

- Tend to coevolve teams by representing behaviors for each agent in a separate population
- Natural to partition a "team" of agents along agent boundaries
- We want the team to perform well, even when one member does something unexpected
- By our definition, agent $\bar{s}$ is *robust* with respect to the rest of the team if the team meets the criterion in spite of changes in the behaviors of the T agent
- The key is to select a meaningful criterion

Some Robustness Criteria

Given an objective function $f : S \times T \mapsto \mathbb{R}$,

Mediocre Criterion

$$\mathcal{C}_s\left(\bar{s}, t_j\right) = \begin{cases} 1 & \text{if } f\left(\bar{s}, t_j\right) \geq \sum_i \frac{f(s_i, t_j)}{|S|} \\ 0 & \text{otherwise} \end{cases}$$

Paul's Crazy Min-Max Criterion

$A = T - \{t_j\},\ B = S - \{\bar{s}\}$

$$\mathcal{C}_s\left(\bar{s}, t_j\right) = \begin{cases} 1 & \text{if } \frac{\min f(\bar{s}, A)}{\max f(B, t_j)} > c_1 \\ 0 & \text{otherwise} \end{cases}$$

Sufficient Value Criterion

$$\mathcal{C}_s\left(\bar{s}, s_j\right) = \begin{cases} 1 & \text{if } f\left(\bar{s}, s_j\right) \geq s^* - \delta \\ 0 & \text{otherwise} \end{cases}$$

	s_1	s_2	s_3	
t_1	**5**	4	2	$\mathcal{D}_s(s_1, T) = \frac{2}{3}$
t_2	4	3	1	
t_3	2	1	**6**	$\mathcal{D}_s(s_3, T) = \frac{1}{3}$

	s_1	s_2	s_3	
t_1	**5**	4	2	$\mathcal{D}_s(s_1, T) = \frac{2}{3}$
t_2	4	3	1	(for $c_1 = \frac{1}{4}$)
t_3	2	1	**6**	$\mathcal{D}_s(s_3, T) = 0$

	s_1	s_2	s_3	
t_1	**5**	4	2	$\mathcal{D}_s(s_1, T) = \frac{2}{3}$
t_2	4	3	1	(for $\delta = 2$)
t_3	2	1	**6**	$\mathcal{D}_s(s_3, T) = \frac{1}{3}$

Observations & Musings

Observations:

- Early proofs involving trapping regions, but ...
- Hard to connect trapping region concept (rigorously) to robustness definitions
- What I could prove was much weaker than what I believed ...

Musings:

- Might posit parametric degree of robustness definitions: $\mathcal{D}_s(m, T) := \frac{\sum_j a_{mj}}{\sum_i \sum_j a_{ij}}$
- Can compare two strategies in terms of their robustness without an explicit robustness criterion

What do I believe?

- Basis vectors associated with strategy sets that have maximal cumulative joint rewards will attract more trajectories than those that do not
- The crucial "robustness" definition that captures this notion is this parametric degree of robustness measure just mentioned
- Key lies in relating the size of the basins of attraction to robustness definition (but not with trapping regions)
- Useful to think in terms of *progress*: When are components of the population state guaranteed *increase*

A Simple Condition for Progress

Theorem

Given a population vector $\vec{x}$ such that $x_m > x_k \, \forall k \neq m$ and then $x'_m > x_m$.

Proof Sketch

Consider two cases:

i) $x_m \leq \frac{1}{n}$:

$$\sum_j (a_{mj} - a_{kj}) y_j > 0 \quad \rightsquigarrow \quad \frac{1}{n} \sum_j a_{mj} y_j > x_k \sum_j a_{kj} y_j$$

$$n \cdot \left(\frac{1}{n} \sum_j a_{mj} y_j \right) > \sum_i \sum_j a_{ij} x_i y_j \quad \rightsquigarrow \quad \frac{\sum_j a_{mj} y_j}{\sum_i \sum_j a_{ij} x_i y_j} > 1$$

$$\Rightarrow \quad x'_m > x_m$$

ii) $x_m > \frac{1}{n}$: *assume $x'_m < x_m$, show that when true, $x_m \not> \frac{1}{n}$, a contradiction*

Other Obvious Observations

- $\vec{y}$ is a vital part of the condition $\boxed{\sum_{j=1}^{n} \left(a_{mj} - a_{kj}\right) y_j > 0}$
- If we have $x_m > x_k$ and $\sum_j a_{mj} y_j > \sum_j a_{kj} y_j$ then $x'_m > x'_k$
- And again, $\vec{y}$ affects the order of $\mathcal{D}_s$ $\boxed{\frac{\sum_j a_{mj}}{\sum_i \sum_j a_{ij}}}$
- For *doubly symmetric* payoff matrices ($A = A^T$), if $\mathcal{D}_x(m) > \mathcal{D}_x(k)$ then once x_m exceeds x_k, it remains larger

Less Obvious, More Interesting

(and incomplete and a bit speculative ...)

- So $\vec{y}$ specifies the region of the search space in which progress on the x population can be ensured
- The size of that space relates to the relative differences in $\mathcal{D}_s$
- The bigger those differences, the bigger the region of the $\vec{y}$ space under which the condition holds
- One should be able to characterize the relative sizes of the basins of attraction using this relationship

Suggestions, Discussion, Questions?

R. Paul Wiegand
wiegand@ist.ucf.edu
Natural Computing and Coadaptive Systems Laboratory
Institute for Simulation & Training
University of Central Florida

Acknowledgments

This work is part of the "Learning Robust Behaviors in Mixed-Agent, Heterogeneous Teams" portion of the ARL/HRED funded *Team Performance in Human-Agent Collaboration* project (W911NF-07-R-0001).