Advanced Tutorial on Coevolution

Sevan G. Ficici
Harvard University
Cambridge MA USA 02138
sevan@eecs.harvard.edu

Anthony Bucci
Brandeis University
Waltham MA USA 02454
abucci@cs.brandeis.edu

GECCO 2007

Copyright is held by the author/owner(s).
GECCO'07, July 7–11, 2007, London, England, United Kingdom.
ACM 978-1-59593-697-4/07/0007.

© Copyright 2007 by Sevan G. Ficici & Anthony Bucci

Conventional Coevolution

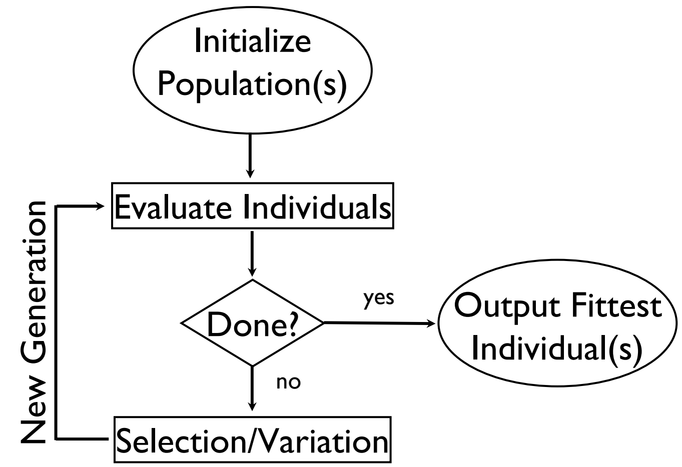

Conventional Coevolution

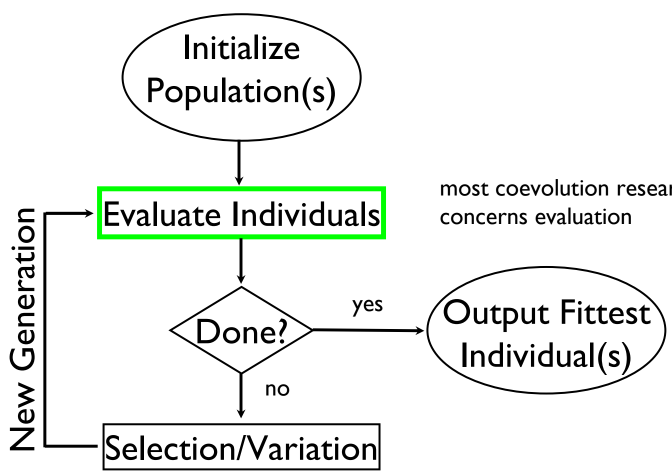

most coevolution research concerns evaluation

Conventional Coevolution

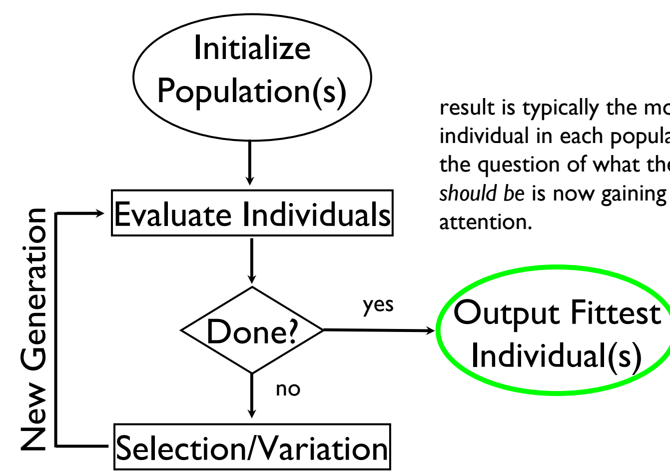

result is typically the most-fit individual in each population; the question of what the result *should be* is now gaining more attention.

Conventional Coevolution

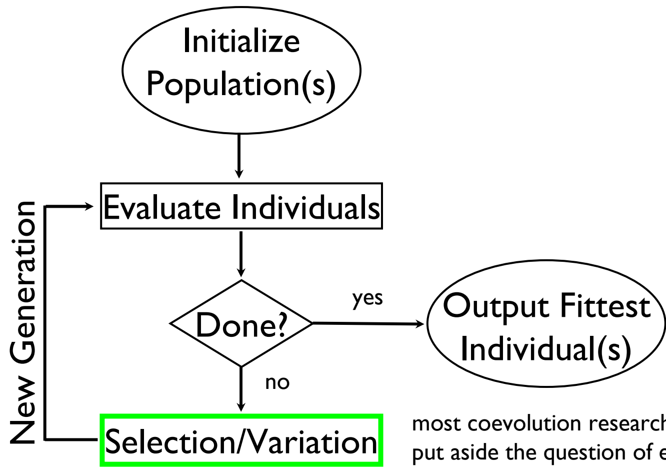

most coevolution research has put aside the question of evolutionary representation; this is an area that is gaining more attention now.

Main Themes

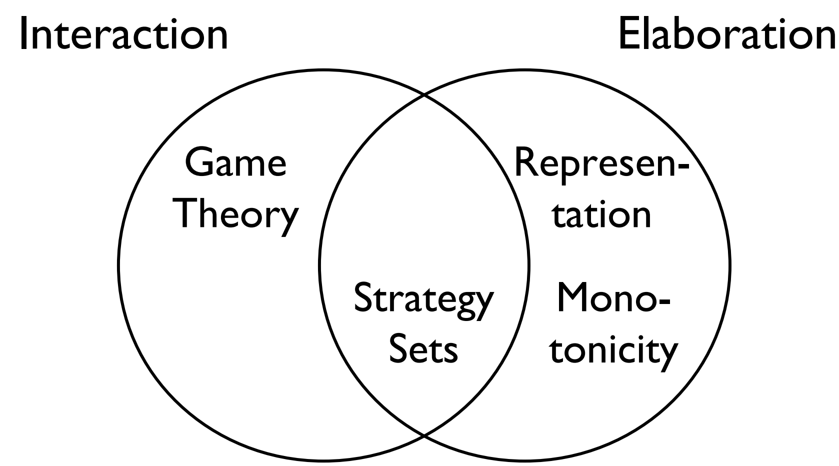

Interaction

- To evaluate an individual in coevolution, we must have it interact with others
- The outcome of evaluation is contingent upon whom the individual interacts with
- The individual may appear good in one context and poor in another context
- This context sensitivity is game theoretic in nature
- Solutions may be sets of individuals

Elaboration

- We want the evolving individuals to improve over evolutionary time
- Coevolutionary “arms race” is an example
- Improvement can be viewed as an accumulation of competences, or elaboration
- We will discuss different forms of elaboration

Main Topics

- Game theory
  - game, strategies, payoffs
  - solution concepts: implementation
- Strategy sets
  - Mixtures, Pareto front, archives, ...
- Representation
- Monotonic improvement over time

Motivation: Coevolutionary Pathologies

- <u>Cycling</u>: algorithm revisits a portion of state-space periodically—no progress
- <u>Disengagement</u>: loss of fitness gradient
- <u>Overspecialization</u>: lack of elaboration
- <u>Forgetting</u>: loss of potentially useful traits
- <u>Relative overgeneralization</u>: favoring of versatile components over those of optimal solution

Game Theory

- Mathematics of strategic reasoning
  [Fundenberg & Tirole 1998]
- If we have a number of interacting agents...
  - How <u>will</u> they behave; what will be outcome?
  - If we interact, how <u>should</u> we behave?
- Provides descriptive <u>predictions</u> of how players will behave
- Provides prescriptive (normative) <u>instructions</u> on how to behave

Game Theory

- Provides predictions and instructions about behavior
- Assumes all agents are rational, selfish
- Nash equilibrium [Nash 1951]
  - A configuration of strategic choices such that no player has incentive to deviate unilaterally from its current strategy
  - All finite games have at least one Nash equilibrium

Game Theory: Components

- Game specifies for each player…
  - strategies that are available
  - outcomes that result for each strategy when interacting with other players' strategies
- Solution concept
  - formal specification of what configuration of players' behaviors (strategies) constitutes a solution to the game

Rock Paper Scissors

Player 2

Player 1	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

Rock Paper Scissors

Player 2

Player 1	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

- Pure strategies: rock, paper, scissors
- Mixed strategy: any probability distribution over pure strategies

Rock Paper Scissors

Player 2

Player 1	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

- Payoffs (outcomes) for all possible pure-strategy interactions
- For mixed strategies, we calculated expected payoffs based on probability distributions used

Rock Paper Scissors

Player 2

Player 1	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

- Rock > Scissors > Paper > Rock
- No pure strategy is universally best
- Solving this game requires a <u>set</u> of strategies

Rock Paper Scissors

Player 2

Player 1	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

- Nash equilibrium strategy is mixed
  - R, P, S each played with probability = 1/3
  - expected payoff of Nash player is zero, regardless of what other player does
  - expected payoff of other player is also zero, regardless of what it does

Main Themes

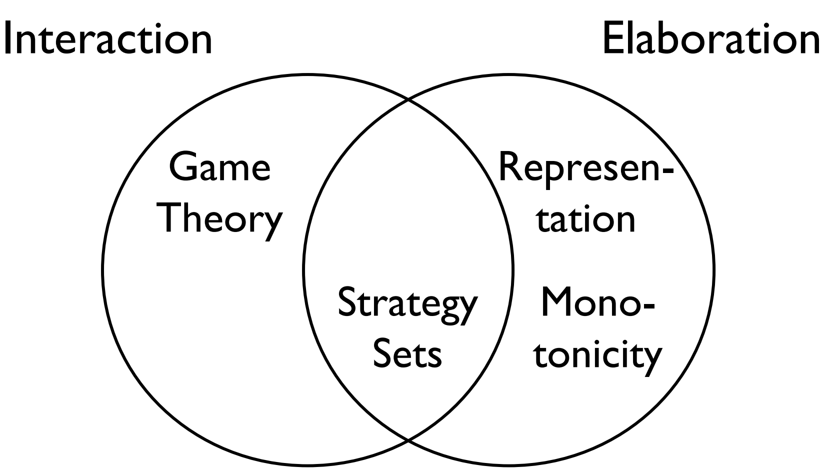

Interaction and Elaboration

- From the outcomes of pure-strategy interaction...
  - we find that no single pure strategy provides all needed competences
- The Nash mixed-strategy...
  - is a set of pure strategies...
  - and represents an elaboration of pure-strategies

Solution Concept

- Specifies <u>properties</u> of a solution
  - (not the solution itself)
- But must be implemented in search algorithm
- Incorrect implementation of solution concept will cause search algorithm to diverge from desired solution properties

Solution Concept

- Examples where algorithm fails to implement Nash equilibrium in a game
  - Proportional selection and Rock-Paper-Scissors: mixed Nash equilibrium? [Hofbauer & Sigmund 1998]
  - Alternative selection methods and Hawk-Dove game [Ficici et al. 2000, 2005]
  - Diversity maintenance methods and Hawk-Dove game [Ficici 2001]

Rock-Paper-Scissors

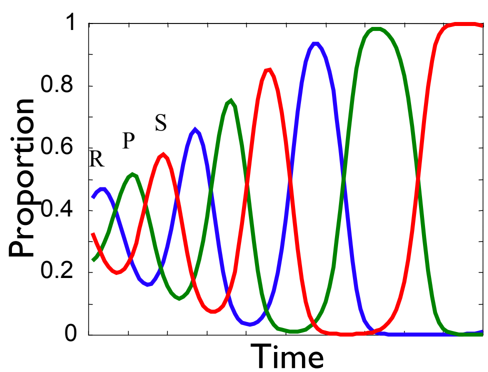

- Under fitness-proportional selection…
  - Nash equilibrium represented as polymorphic population of pure-strategists is unstable
  - Nash equilibrium also unstable for mixed strategists
- Nash concept not properly implemented here

Hawk-Dove Game

[Maynard Smith 1982]

	Hawk	Dove
Hawk	-25	50
Dove	0	15

- Nash equilibrium strategy for these payoffs:
  - 7/12 Hawk, 5/12 Dove
  - probability distribution for a mixed strategy…
  - OR proportions for polymorphic population of pure-strategists

Proportional Selection

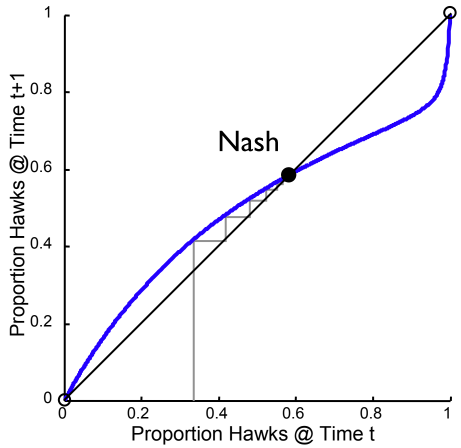

Nash equilibrium is dynamical attractor

Truncation Selection

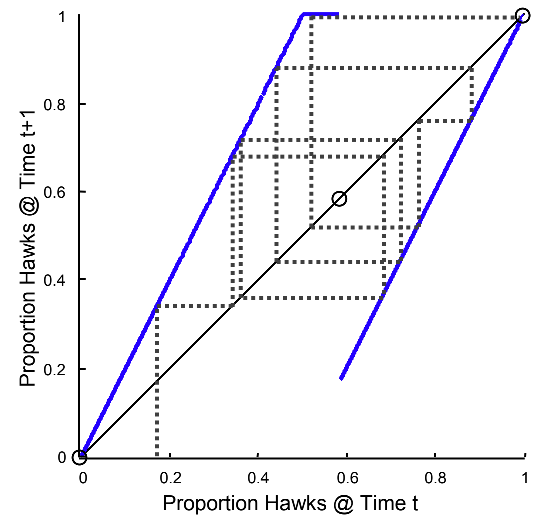

All-Hawks is now an attractor

Truncation Selection

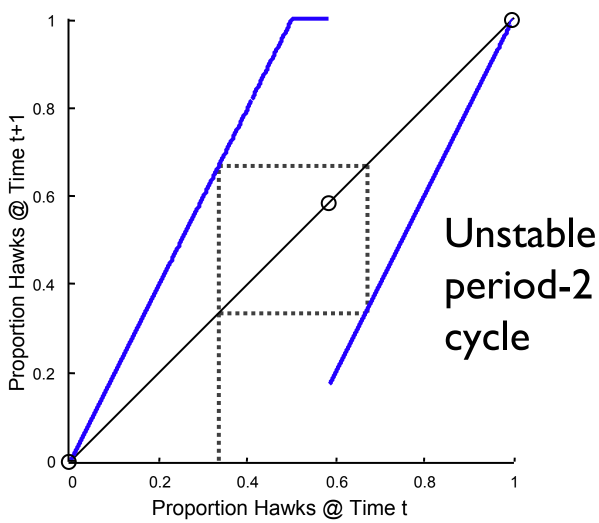

Truncation Selection

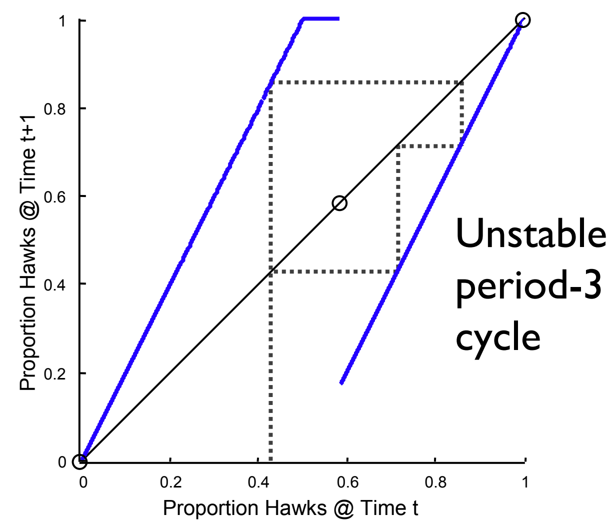

Truncation Selection

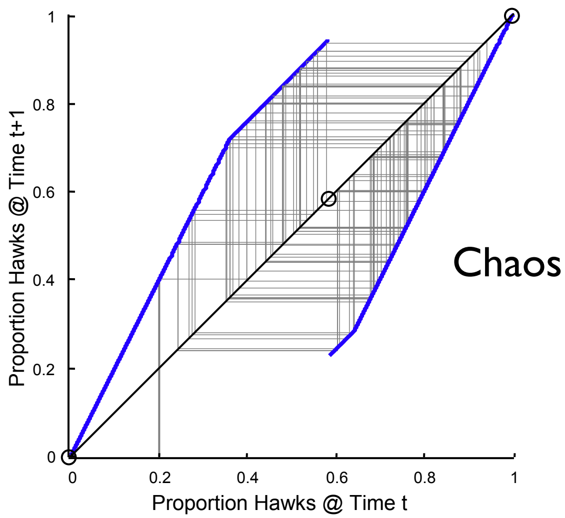

Truncation Selection

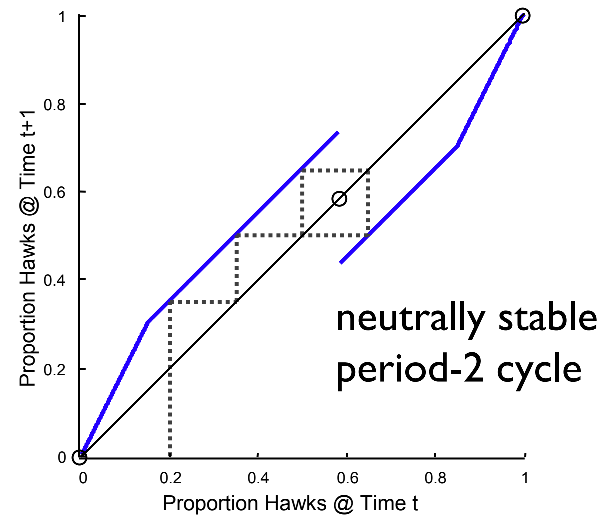

(μ, λ)-ES Selection

$\mu/\lambda = 0.6$ $\mu/\lambda = 0.5$ $\mu/\lambda = 0.3$

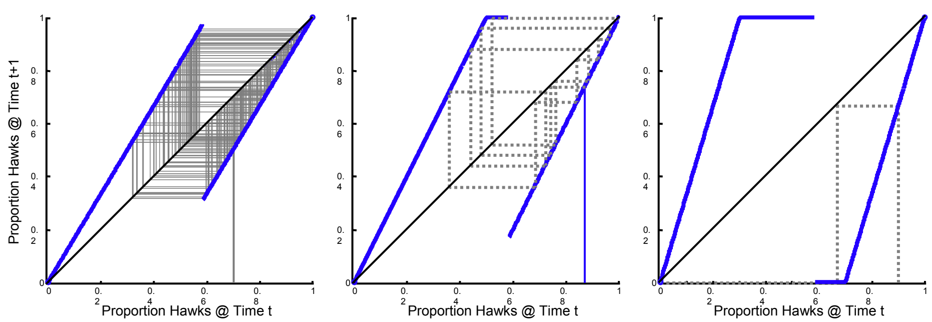

Regime 1
(0 0.41]
Chaos

Regime 2
[0.42 0.58]
All-Hawk

Regime 3
[0.59 1.0)
All-H or -D

selection pressure increase →

Rank Selection

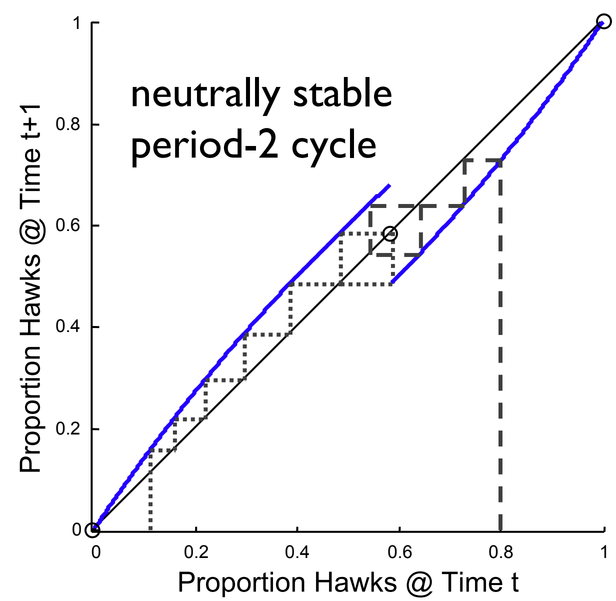

Boltzman Selection

$$f_{\text{Boltz}} = e^{\gamma f}$$

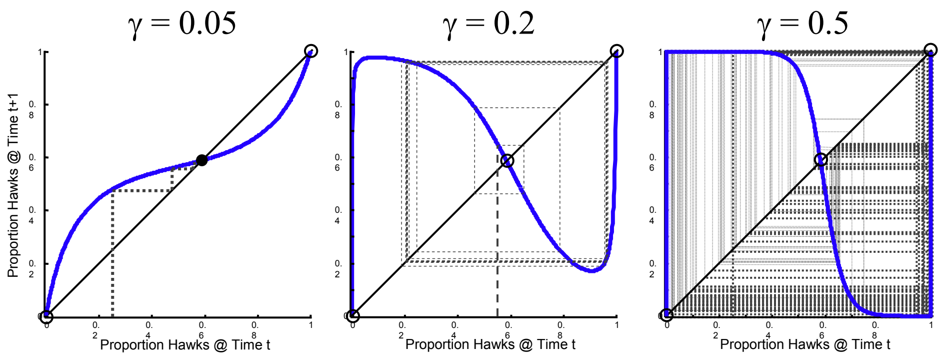

Nash Intact | Attr. Limit Cycle | Chaos

selection pressure increase →

Proportional Selection

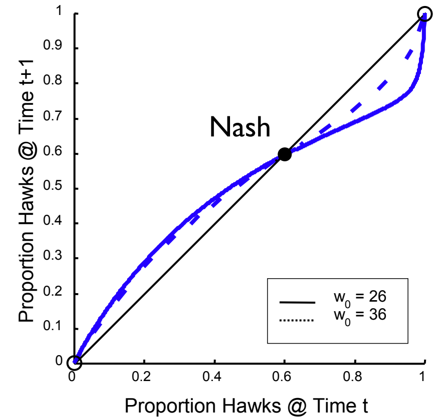

Competitive Fitness Sharing

[Rosin 1997]

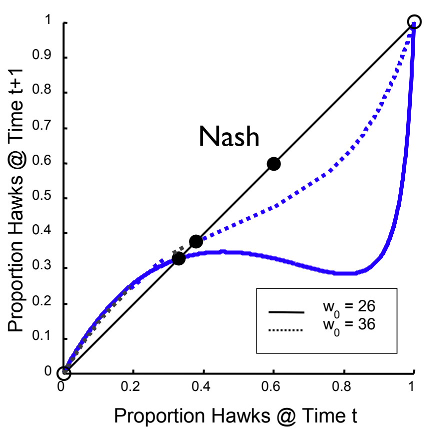

Similarity-Based Fitness Sharing

[Goldberg 1989]

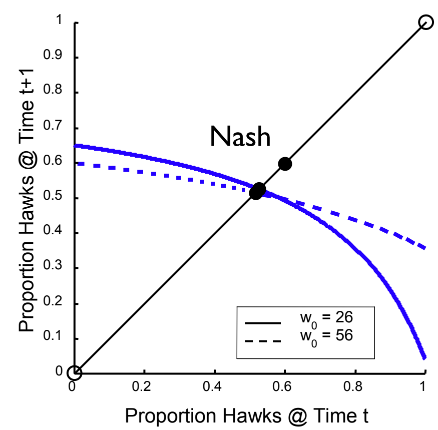

Discussion

- We use different selection methods and diversity-maintenance methods to improve search for a particular domain
- Evolving population expected to both:
  - contain sufficient genetic diversity for search
  - represent solution to search task (may be a polymorphism)
- These tasks not necessarily orthogonal
- Above illustrates pitfalls

Discussion

- Why not separate tasks?
- Let population perform search
- Let another mechanism (not population) represent best solution found so far
- Leads us to archive methods

Archive Methods

- Archives provide a way to
  - collect (according to some organizing priniple) "good" individuals over evolutionary time
  - encapsulate wider phenotypic range (than a population contains at any one moment in time)
  - broaden evaluation (and selection pressure) via augmented phenotypic diversity
  - ameliorate evolutionary forgetting
  - represent the result of the evolutionary process

Archive Methods

- Hall-of-Fame [Rosin & Belew 1997]
  - accumulate fittest of each generation
  - sample *k* members for testing current generation
  - shown to help, but weak organizing principle
- Dominance Tournament [Stanley & Miikkulainen 2002]
- Nash memory [Ficici & Pollack 2003]
- Pareto archives [de Jong 2004]

Dominance Tournament

- For zero-sum games, symmetric or not
- Organizing principle is Pareto dominance
- Add strategy X to DT archive if and only if X outperforms each member of archive
- Each new member inserted into archive has a broader demonstrated range of competence
- Avoids intransitive cycles
- Most recently added member is “solution”

Dominance Tournament

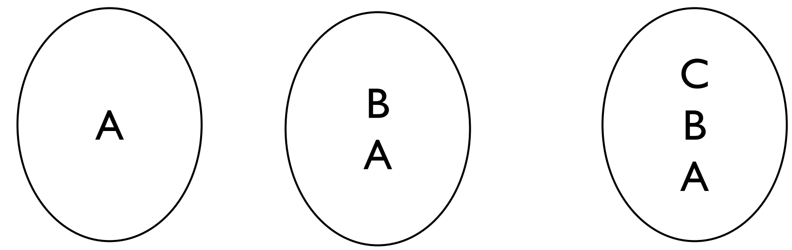

If B beats A

If C beats B and A

- If D beats C, but not B, then D excluded

Nash Memory

- For zero-sum games, symmetric or not
- Organizing principle is Nash equilibrium
- Begin with arbitrary approximation to Nash equilibrium N of game, and empty “memory” M
- If strategy S beats N, then update N and M to obtain a new Nash approximation that doesn’t lose to any strategy in S ∪ N ∪ M
- Final approximation N is “solution”

Nash Memory

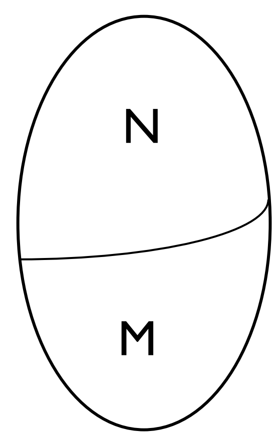

- N and M are mutually exclusive sets of pure strategies
- N is initially an arbitrary pure strategy
- M is initially the empty set

Nash Memory

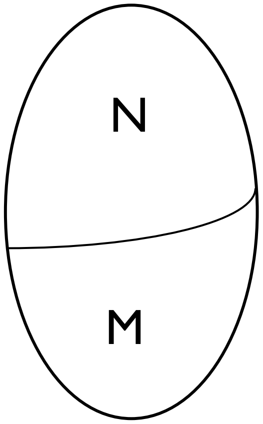

- N is Nash strategy with respect to pure strategies in N ∪ M
- M is set of strategies that used to be in N earlier and may be again in the future
- M is of bounded size

Nash Memory

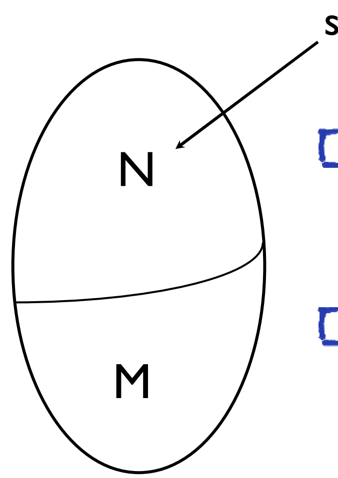

- If s does not beat N, then discard s; keep searching
- Otherwise, s indicates a weakness in N; update N

Nash Memory

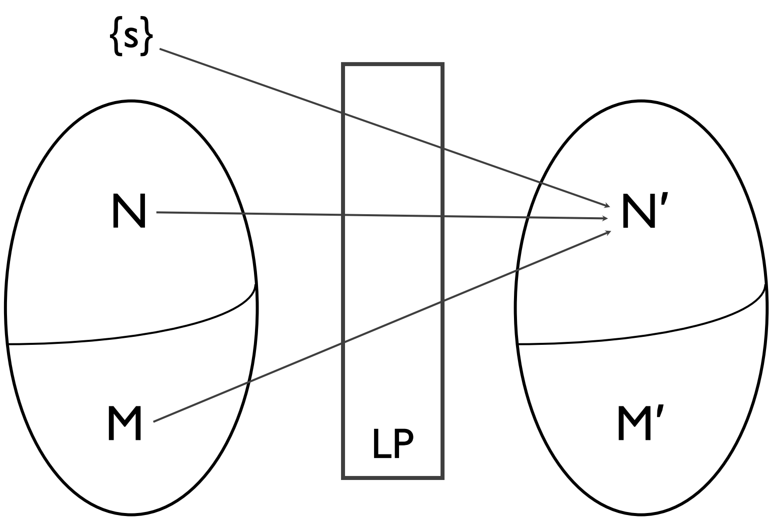

Nash Memory

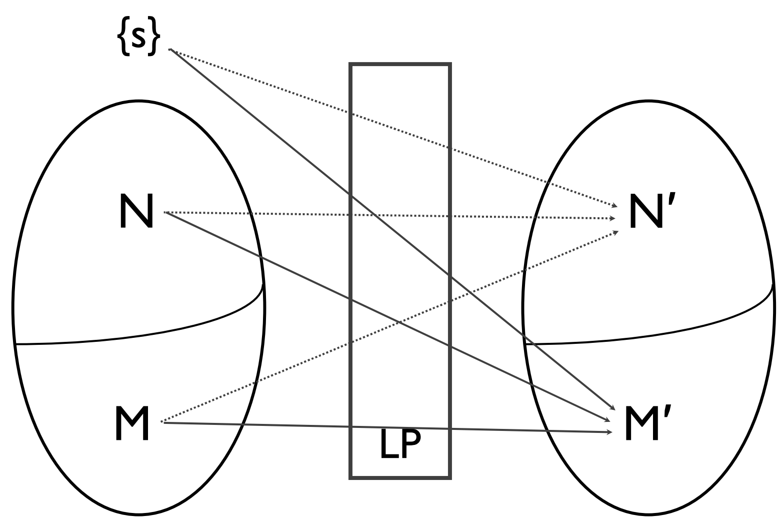

Nash Memory

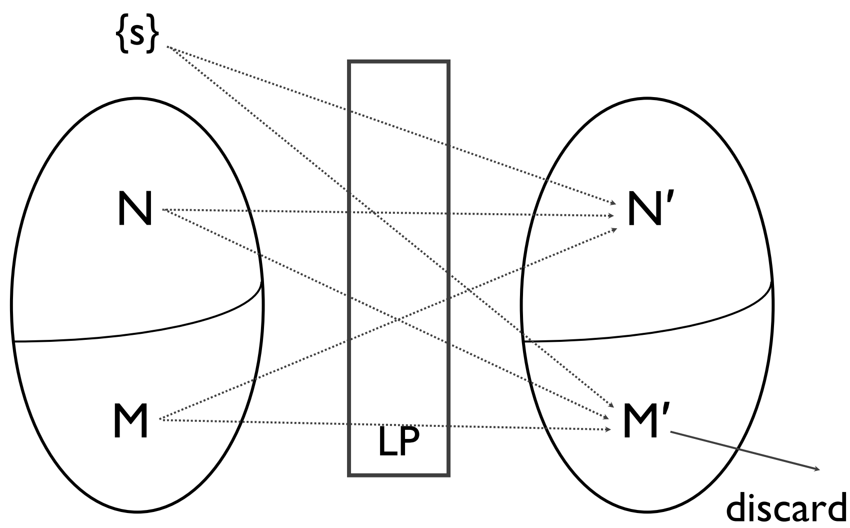

Nash Memory

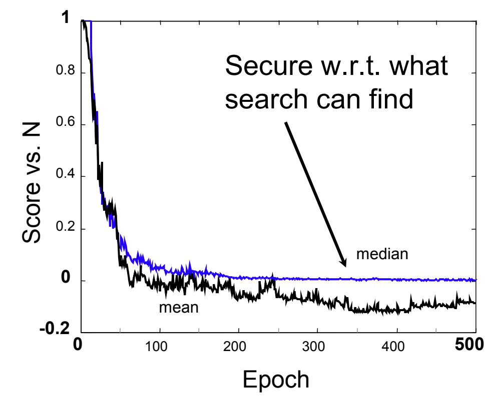

- Even though M is finite, and strategies discarded, monotonic improvement of N is approximated

Pareto Archives

- *Pareto Coevolution* [Ficici & Pollack 2001; Noble & Watson 2001] treats entities with two roles: *candidates* and *tests* (sometimes *learners* and *teachers*)
- **Candidates** are incented to *perform*
- **Tests** are incented to *inform* about candidates
- Key insight: performing ≠ informing

Performing

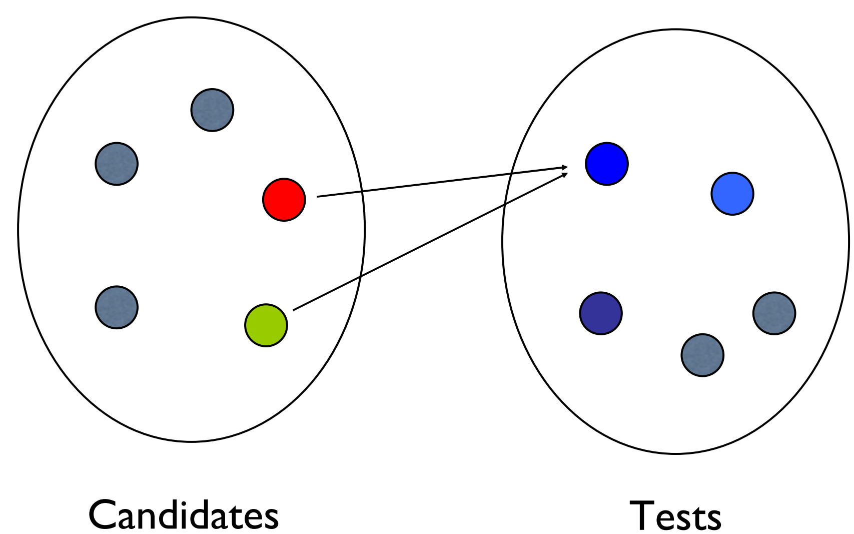

Performing

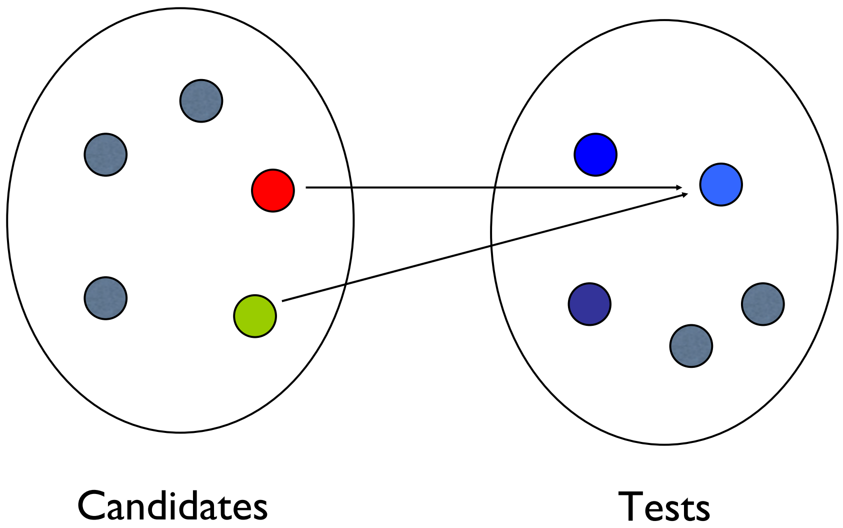

Performing

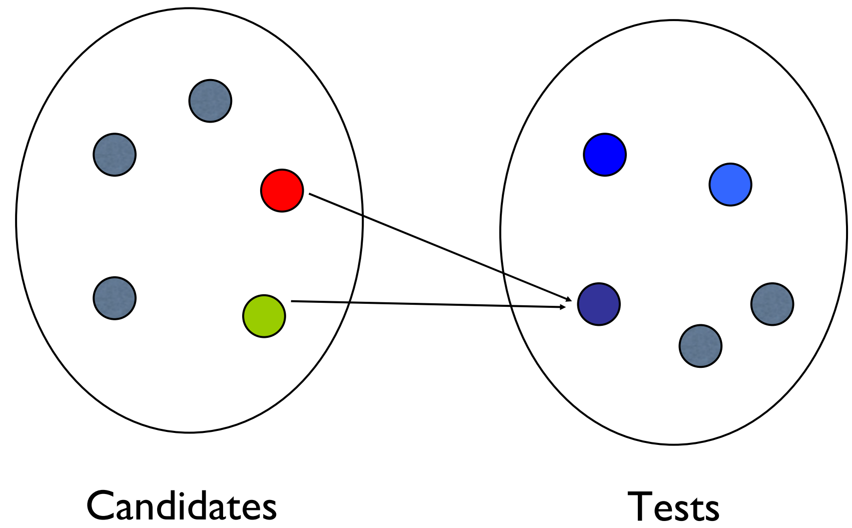

Informing

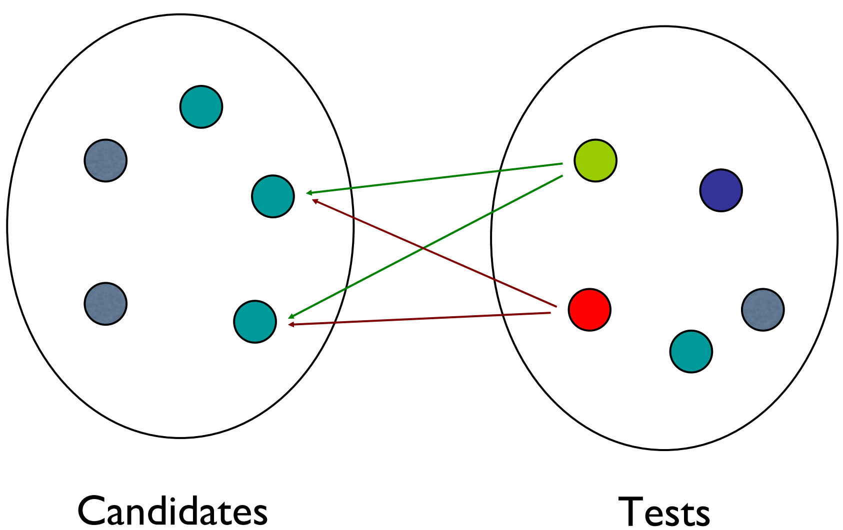

Informing

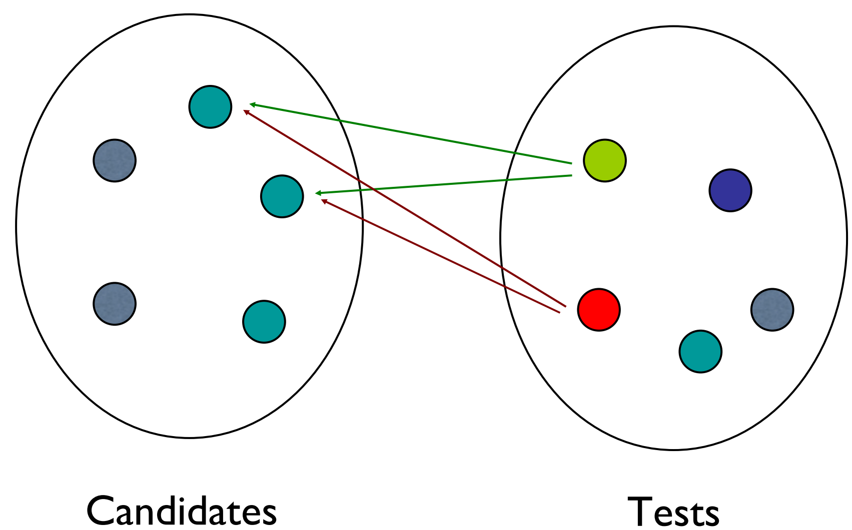

Informing

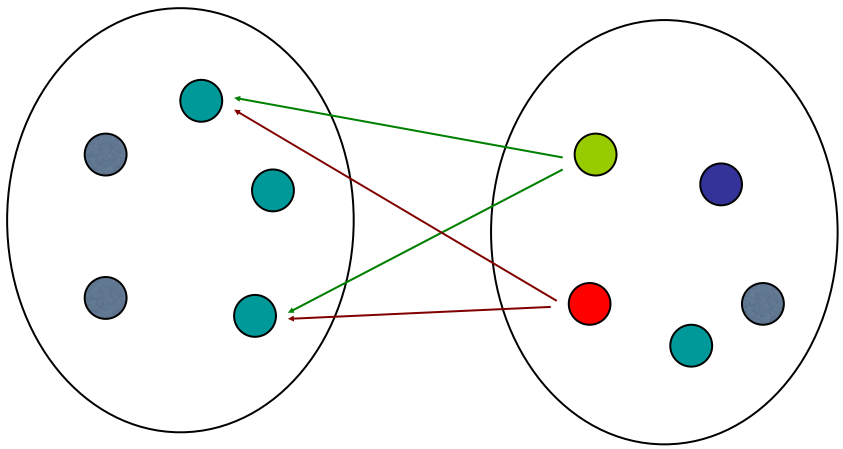

Pareto Archives: IPCA

- Incremental Pareto Coevolution Archive [de Jong 2004]
- Theoretically ensures monotonic progress for Pareto Coevolution
- Allows the candidate population to explore
- Test archive maintains candidate distinctions and can grow without bound

Pareto Archives: LAPCA

- LAyered Pareto Coevolution Archive [de Jong 2004]
- Keeps a tunable number of Pareto layers
- Approximates IPCA, but bounds the archive – loses monotonicity guarantee
- Combined with NEAT and applied to coevolve Pong players [Monroy et al. 2006]

Test-Based Problems

- Candidate solutions are *tested* by interacting with other entities, as in:

Domain	***Candidate***	***Test***
Design	Sorting network	Unsorted list
Classification	Classifier	Data point
Function/model regression	Function or model	Input
Strategy learning	First player	Second player

Pareto Coevolution

- Maintains two populations, candidate solutions and tests
- **Candidates** are compared using Pareto dominance: *A* dominates *B* if it does at least as well as *B* against all tests and better on at least one
- **Tests** are compared using *distinctions* [Ficici & Pollack 2001] or *informativeness* [Bucci & Pollack 2003]
- Solution set is non-dominated front of candidates and an informative set of tests

Shows a Distinction

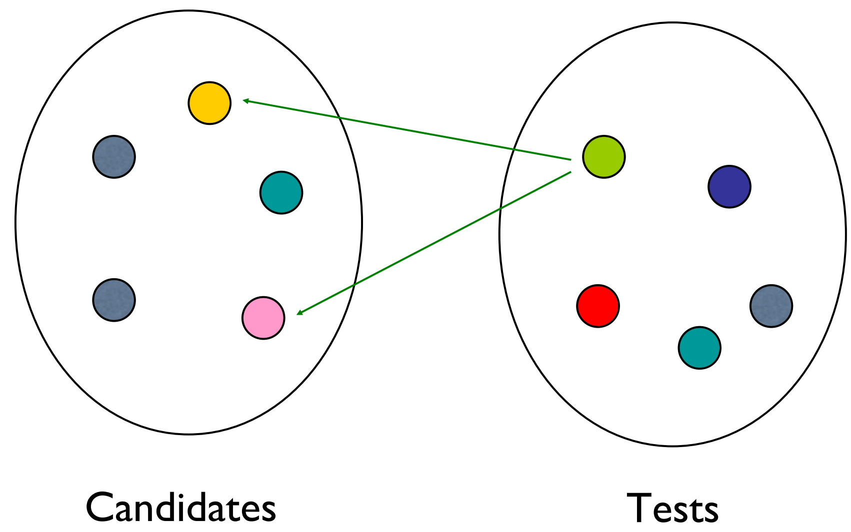

Shows No Distinction

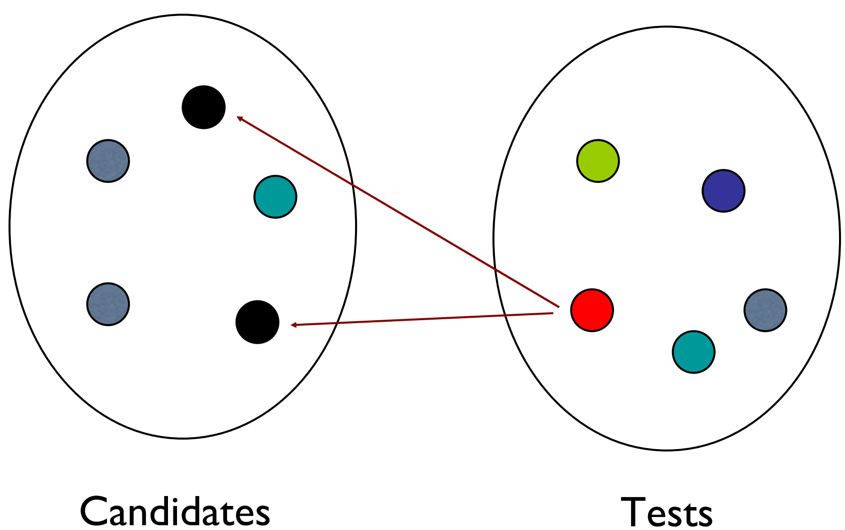

Informative

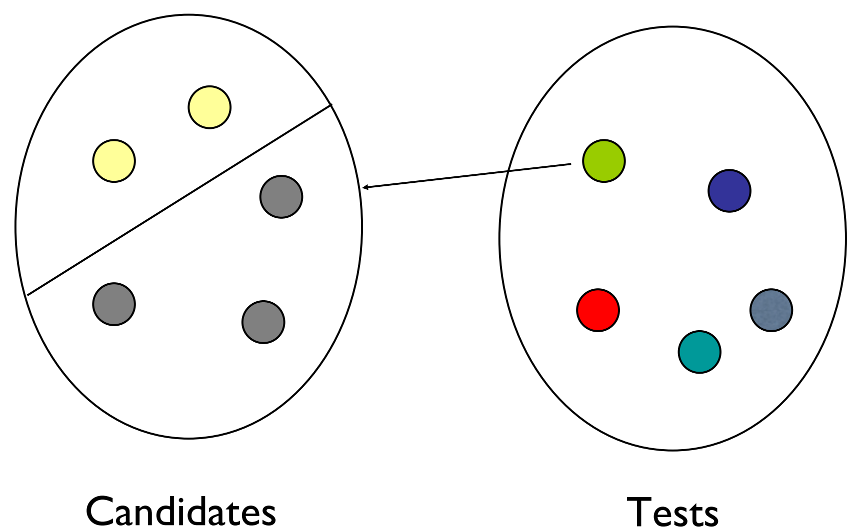

Differently Informative

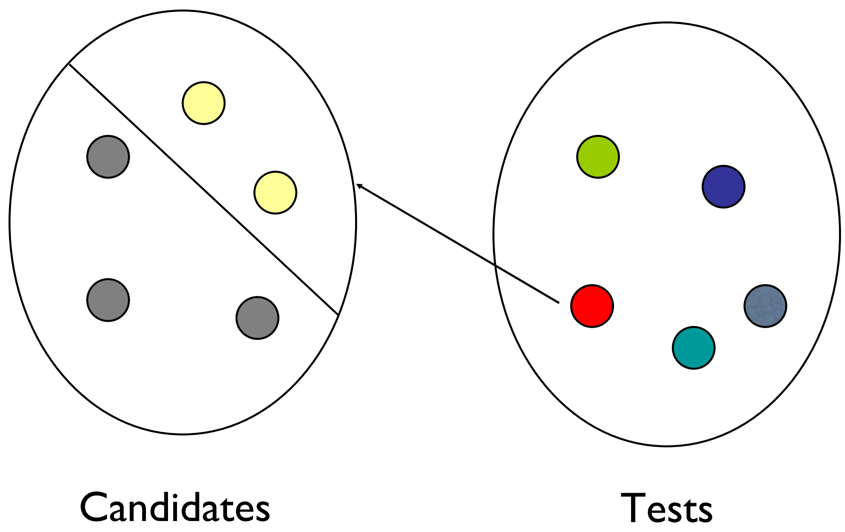

Uninformative

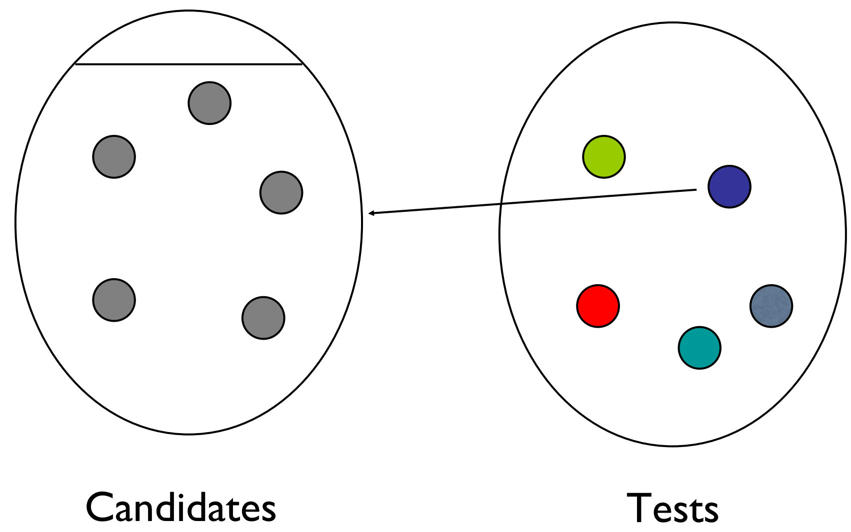

Differently Uninformative

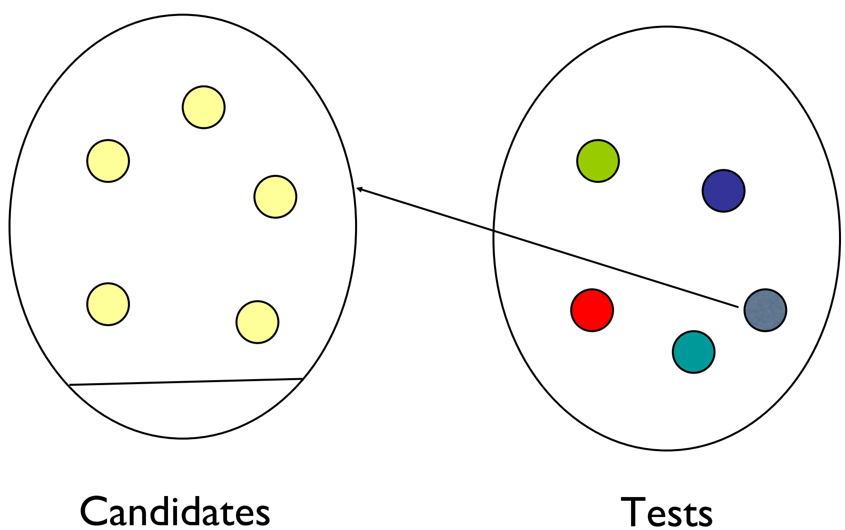

Reducing the Number of Tests

- Ideal Test Set [Bucci & Pollack 2003]
- Complete Evaluation Set [de Jong & Pollack 2004]
- [de Jong & Pollack 2004] posited the existence of multiple *underlying objectives* akin to fitness functions
- [Bucci et al. 2004] grounded the idea as *coordinate systems* (*informative dimensions*)

Dimension Extraction

- Coordinate systems collect several tests into a composite *axis*
- Set of axes forms a coordinate system analogous to a basis for a vector space
- [Bucci et al. 2004] proved coordinate systems exist and gives a polynomial-time algorithm to extract one
- [de Jong & Bucci 2006] gave a CEA, DECA, which extracts coordinate systems from populations to inform selection

Dimension Extraction

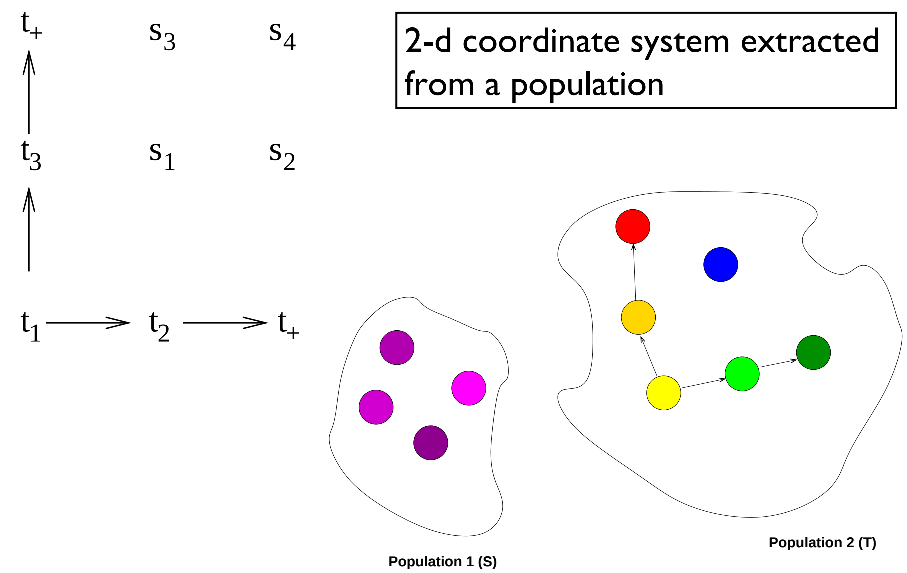

Reducing the Amount of Testing: EEA

- Estimation-Exploration Algorithm [Lipson et al. 2005]
- **Candidates** are models of a system
- **Tests** are probes of the real system (assumed to be expensive)
- Aim is to evolve a model of the real system using as few probes as possible

Main Themes

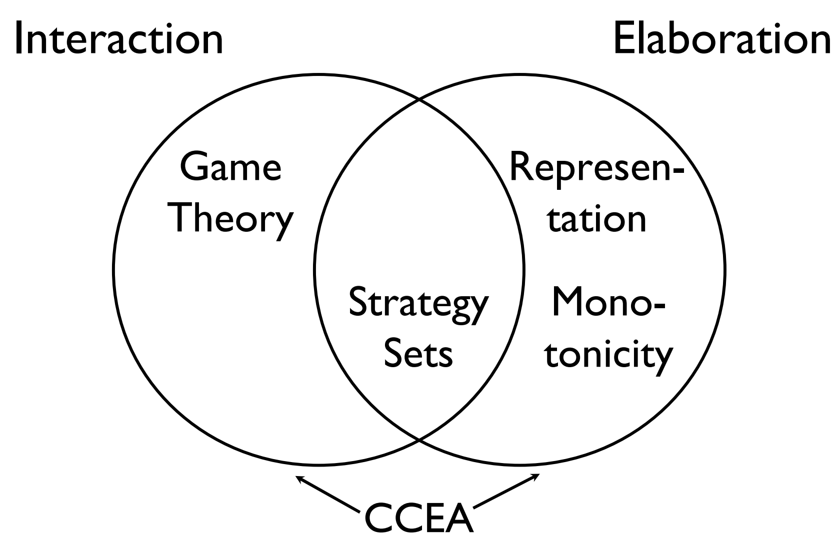

Cooperative Coevolution
[Potter & De Jong 1994]

- Monolithic problem may be too difficult

Cooperative Coevolution
[Potter & De Jong 1994]

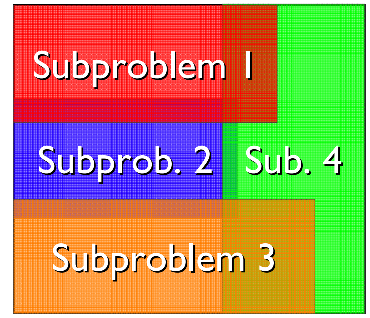

- Decompose into mosaic of semi-independent sub-problems

Cooperative Coevolution
[Potter & De Jong 1994]

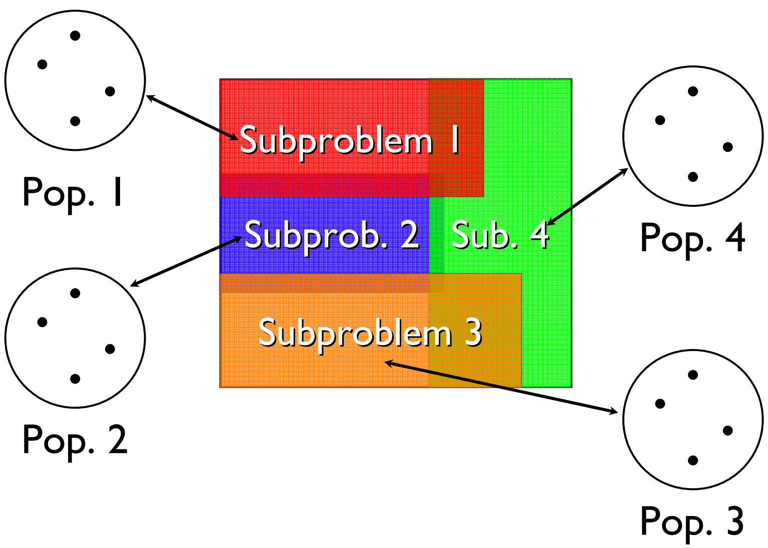

Cooperative Coevolutionary Algorithms (CCEAs)

- [Potter & De Jong 1994] argued that CCEAs optimize functions
- [Wiegand 2003] argued they do not optimize functions, but rather optimize for robustness.
- Used EGT to argue certain Nash equilibria are preferred
- Same work suggested biasing CCEA such that they do optimize

Biasing CCEA Towards Optimization

- [Panait et al. 2004] aimed to bias the CCEA by mixing evaluation with another term biasing towards its optimal evaluation
- [Bucci & Pollack 2005] used Pareto dominance comparison with no bias term; collaborators were tests
- [Panait et al. 2006] proposed an archive of good collaboration choices, iCCEA

Analyzing Collaboration Schemes

[Popovici & De Jong 2005]

- Best response curves are a property of a problem
- In CCEA, intersection points of best response curves are Nash equilibria
- Trajectories of individuals is a propety of an algorithm; e.g., the collaboration scheme
- Trajectories which land at best response curve intersection points get stuck even if they are suboptimal

Analyzing Collaboration Schemes

[Popovici & De Jong 2005]

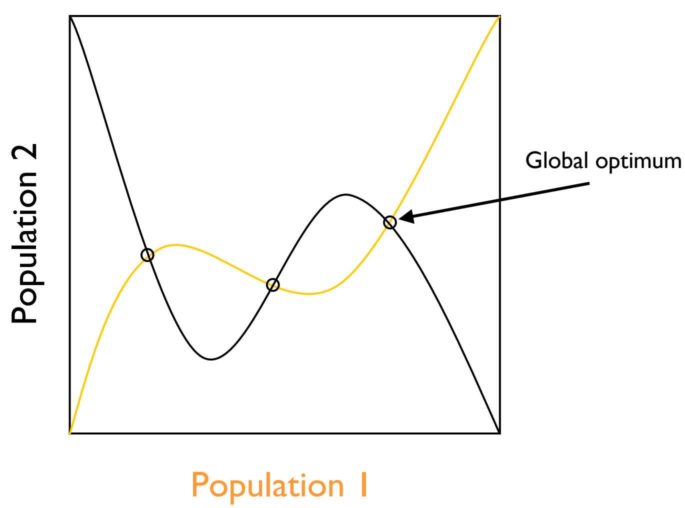

NeuroEvolution of Augmenting Topologies (NEAT)

- Evolves increasingly complex neural network topologies [Stanley & Miikkulainen 2004]
- Mutations occasionally add new structure
- Speciation protects innovative structures
- In combination, these mechanisms support elaboration

Alteration vs. Elaboration

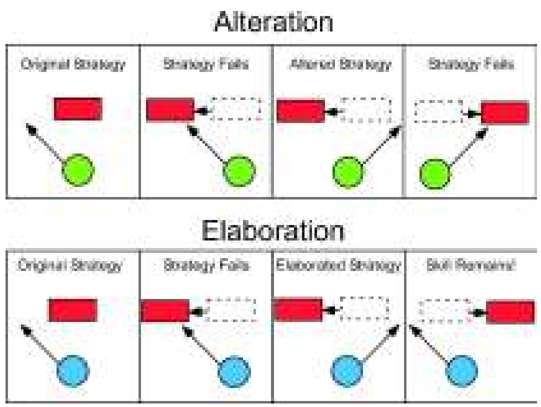

Alteration alone may damage capabilities
Elaboration accumulates capabilities
Can we abstract this idea?

Progress in Coevolution

- A core theme in coevolution research: How to ensure progress—is it possible?
- Evaluation: individual interacts with others
- Measured quality of an individual is function of which other individuals interact with it
- Constantly shifting landscape!
- Open-ended search spaces problematic

Progress in Coevolution

- Monitoring progress
  - Miller & Cliff 1994
  - Floreano & Nolfi 1997
  - Rosin 1997
  - Stanley & Miikkulainen 2002
  - Bader-Natal & Pollack 2004, 2005

Approach

- Examine the issue of progress from viewpoint of <u>solution concepts</u>
- Some solution concepts intrinsically "support" monotonic progress
- Not a value judgment—use whatever solution concept is appropriate
- But something to be aware of!

Desirable Property

- As your knowledge of a search-space increases...
- ... your estimations of a solution should improve

The longer the algorithm runs,
the better the output should be!

(Experience tells us this is not the case in coev.)

Desirable Property

$\mathcal{P}$

- As your knowledge of a search-space increases...
- ... your estimations of a solution should improve

Desirable Property

$\mathcal{P}$

- As your knowledge of a search-space increases... $\mathcal{W} \subset \mathcal{P}$
- ... your estimations of a solution should improve

Desirable Property

$\mathcal{P}$

- As your knowledge of a search-space increases... $\mathcal{W} \subset \mathcal{P}$
- ... your estimations of a solution should improve $\cancel{\mathcal{C} \in \mathcal{P}}$

$\mathcal{C} \in 2^{\mathcal{P}}$

Desirable Property

$\mathcal{P}$

- As your knowledge of a search-space increases... $\mathcal{W} \subset \mathcal{P}$
- ... your estimations of a solution should improve $\cancel{\mathcal{C} \in \mathcal{R}}$
  $\mathcal{C} \in 2^{\mathcal{P}}$

Monotonic Improvement

Knowledge

Complete Knowledge

$$\mathcal{W}_t \subset \mathcal{W}_{t+1} \subset \mathcal{W}_{t+2} \subset \ldots \subset \mathcal{W}^*$$

$$\mathcal{C}_t \preceq \mathcal{C}_{t+1} \preceq \mathcal{C}_{t+2} \preceq \ldots \preceq \mathcal{C}^*$$

Estimations

Global Solution

Monotonic Improvement

Complete Knowledge

$$\mathcal{W}_t \subset \mathcal{W}_{t+1} \subset \mathcal{W}_{t+2} \subset \ldots \subset \mathcal{W}^*$$

$$\mathcal{C}_t \preceq \mathcal{C}_{t+1} \preceq \mathcal{C}_{t+2} \preceq \ldots \preceq \mathcal{C}^*$$

Global Solution

Possible in coevolution when used to solve game of strategy?

Reasoning

$$\mathcal{W}_t \subset \mathcal{W}_{t+1} \subset \mathcal{W}_{t+2} \subset \ldots \subset \mathcal{W}^*$$

$$\mathcal{C}_t \preceq \mathcal{C}_{t+1} \preceq \mathcal{C}_{t+2} \preceq \ldots \preceq \mathcal{C}^*$$

- Should get monotonic improvement because...
  - Knowledge of strategy space strictly increasing
  - Evaluation increasingly comprehensive

Reasoning

$$\mathcal{W}_t \subset \mathcal{W}_{t+1} \subset \mathcal{W}_{t+2} \subset \cdots \subset \mathcal{W}^*$$

$$\mathcal{C}_t \preceq \mathcal{C}_{t+1} \preceq \mathcal{C}_{t+2} \preceq \cdots \preceq \mathcal{C}^*$$

- Might not get monotonic improvement because...
  - Evaluation never fully complete
  - New strategy may radically shift evaluations

Shift of Perspective

1. Begin with Rock
2. Discover Paper; Rock loses to Paper
3. Then discover Scissors; Paper loses, Rock wins

- Paper appears to lose quality; Rock gains it
- Directional improvement ill-defined concept

Depends on Solution Concept

$$\mathcal{W}_t \subset \mathcal{W}_{t+1} \subset \mathcal{W}_{t+2} \subset \cdots \subset \mathcal{W}^*$$

$$\mathcal{C}_t \preceq \mathcal{C}_{t+1} \preceq \mathcal{C}_{t+2} \preceq \cdots \preceq \mathcal{C}^*$$

- If solution concept is "monotonic"...
  - then monotonic increase in knowledge ⇒
  - monotonic improvement of estimation

Monotonicity

With a monotonic solution concept...

If $\mathcal{C}_\alpha$ is solution to X and Z, where $X \supset Z$

Then $\mathcal{C}_\alpha$ will be a solution to <u>any</u> Y where

$X \supset Y \supset Z$

Monotonicity

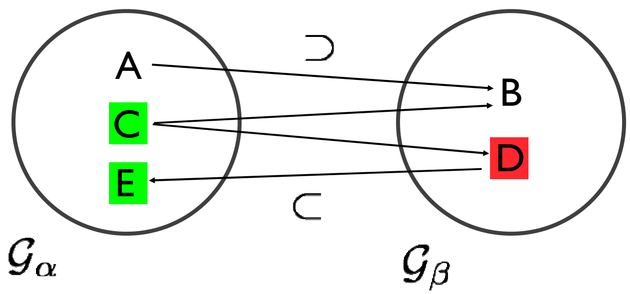

$\mathcal{C}_\alpha$ solution to games C and E, where C $\supset$ E

$\mathcal{C}_\alpha$ not solution to some game D, where C $\supset$ D $\supset$ E

Then solution concept is non-monotonic

Monotonicity

- A monotonic solution concept means:
  - once you discard an estimation in favor of another...
  - you will never return the to earlier estimation
  - ... regardless of whatever new strategies you discover in the future
- Non-monotonic solution concept means:
  - you may return to an estimation from some earlier point in time as you discover new strategies

Monotonic Solution Concepts

- Solution concepts that are monotonic
  - Nash equilibrium
  - Pareto optimality, but only if you exclude newly discovered strategies that appear identical to ones previously discovered
- Non-monotonic
  - Maximal expected payoff; best response
- This notion of monotonicity subsumes that of [de Jong 2005]

Advanced Tutorial on Coevolution—References[1]

1 Background

1.1 Game Theory

[Fudenberg and Tirole, 1998], [Nash, 1951]

1.2 Dynamical Systems

[Strogatz, 1994]

2 Solution Concepts

[Fudenberg and Tirole, 1998], [Ficici, 2004], [de Jong, 2005], [Bucci and Pollack, 2007], [Wiegand, 2003]

2.1 Solution Concept and Evolutionary Dynamics

[Maynard-Smith and Price, 1973], [Maynard-Smith, 1982], [Fogel and Fogel, 1995], [Fogel et al., 1997], [Fogel et al., 1998], [Hofbauer and Sigmund, 1998], [Liekens et al., 2004], [Ficici et al., 2005], [Ficici, 2006], [Ficici and Pollack, 2007]

3 Representation

[Moriarty and Miikkulainen, 1997], [Stanley and Miikkulainen, 2002b], [Stanley and Miikkulainen, 2004], [Ashlock et al., 2006]

4 Evaluation

[Bull, 2001], [Panait et al., 2004], [Popovici and De Jong, 2005a], [Popovici and De Jong, 2005b], [Popovici and De Jong, 2006c], [Popovici and De Jong, 2006b], [Popovici and De Jong, 2006a]

4.1 Test-Based Evaluation

[Juillé and Pollack, 1996b], [Juillé and Pollack, 1996a], [Juillé and Pollack, 1998], [Juillé, 1999], [Juillé and Pollack, 2000], [Watson and Pollack, 2000], [Ficici and Pollack, 2001], [Ashlock et al., 2004], [Bucci and Pollack, 2002], [Bucci and Pollack, 2003], [de Jong and Pollack, 2003], [Bucci et al., 2004], [de Jong, 2004a], [de Jong, 2004b], [Bongard and Lipson, 2005], [de Jong and Bucci, 2006]

[1]©2007 by Sevan G. Ficici and Anthony Bucci

5 Pareto Coevolution

[Watson and Pollack, 2000], [Ficici and Pollack, 2001], [Noble and Watson, 2001], [Bucci and Pollack, 2002], [Bucci and Pollack, 2003], [de Jong and Pollack, 2003], [Bucci et al., 2004], [de Jong, 2004a], [de Jong, 2004b], [Bongard and Lipson, 2005], [de Jong and Bucci, 2006], [Watson, 2006]

6 Archive Methods, design and use

[Rosin and Belew, 1997], [Stanley and Miikkulainen, 2002a], [Ficici and Pollack, 2003], [de Jong, 2004a], [de Jong, 2004b], [Monroy et al., 2006]

7 Progress in Coevolution

[Miller and Cliff, 1994], [Floreano and Nolfi, 1997], [Bader-Natal and Pollack, 2004], [de Jong, 2005], [Bader-Natal and Pollack, 2005], [Ficici, 2005]

8 Cooperative Coevolution

[Potter and Jong, 1994], [Potter and Jong, 2000], [Wiegand et al., 2001], [Wiegand et al., 2002b], [Wiegand et al., 2002a], [Wiegand et al., 2003], [Wiegand, 2003], [Jansen and Wiegand, 2004], [Panait et al., 2004], [Bucci and Pollack, 2005], [Popovici and De Jong, 2005a], [Popovici and De Jong, 2005b], [Popovici and De Jong, 2006c], [Popovici and De Jong, 2006b], [Popovici and De Jong, 2006a]

9 Markov Analyses

[Bull, 2001], [Schmitt, 2003a], [Schmitt, 2003b]

10 No Free Lunch

[Wolpert and Macready, 1997], [Wolpert and Macready, 2005]

References

[Ashlock et al., 2006] Ashlock, D., Kim, E.-Y., and Leahy, N. (2006). Understanding representational sensitivity in the iterated prisoners dilemma with fingerprints. *IEEE Transactions on System, Man, and Cybernetics—Part C: Applications and Reviews*, 36(4):464–475.

[Ashlock et al., 2004] Ashlock, D., Willson, S., and Leahy, N. (2004). Coevolution and tartarus. In *Proceedings of the 2004 IEEE Congress on Evolutionary Computation*, pages 1618–1624. IEEE Press.

[Bader-Natal and Pollack, 2004] Bader-Natal, A. and Pollack, J. (2004). A population-differential method of monitoring success and failure in coevolution. In *Proceedings of the 2004 Genetic and Evolutionary Computation Conference*. Springer.

[Bader-Natal and Pollack, 2005] Bader-Natal, A. and Pollack, J. (2005). Towards metrics and visualizations sensitive to coevolutionary failures. In *2005 AAAI Fall Symposium on Coevolutionary and Coadaptive Systems*. AAAI. AAAI Technical Report FS-05-03.

[Bongard and Lipson, 2005] Bongard, J. C. and Lipson, H. (2005). Nonlinear system identification using coevolution of models and tests. *IEEE Transactions on Evolutionary Computation*, 9(4):361–383.

[Bucci et al., 2004] Bucci, A., Pollack, J., and de Jong, E. (2004). Automated extraction of problem structure. In *Proceedings of the 2004 Genetic and Evolutionary Computation Conference*. Springer Verlag.

[Bucci and Pollack, 2002] Bucci, A. and Pollack, J. B. (2002). Order-theoretic analysis of coevolution problems: Coevolutionary statics. In Barry, A. M., editor, *2002 Genetic and Evolutionary Computation Conference Workshop Program*, pages 229–235.

[Bucci and Pollack, 2003] 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., editors, *Proceedings of the Foundations of Genetic Algorithms 2003 Workshop (FOGA 7)*, pages 221–235. Morgan Kaufmann Publishers.

[Bucci and Pollack, 2005] Bucci, A. and Pollack, J. B. (2005). On identifying global optima in cooperative coevolution. In *GECCO 2005: Proceedings of the 2005 conference on Genetic and evolutionary computation*, volume 1, pages 539–544. ACM Press.

[Bucci and Pollack, 2007] Bucci, A. and Pollack, J. B. (2007). Thoughts on solution concepts. In *GECCO 2007: Proceedings of the 9th annual conference on Genetic and evolutionary computation*, London. ACM Press. Forthcoming.

[Bull, 2001] Bull, L. (2001). On coevolutionary genetic algorithms. *Soft Computing*, 5(3):201–207.

[Cantú-Paz et al., 2003] Cantú-Paz et al., editors (2003). *2003 Genetic and Evolutionary Computation Conference*. Springer.

[de Jong, 2004a] de Jong, E. D. (2004a). The incremental pareto-coevolution archive. In Deb, K. et al., editors, *Proceedings of the 2004 Genetic and Evolutionary Computation Conference*, LNCS 3102, pages 525–536. Springer-Verlag.

[de Jong, 2004b] de Jong, E. D. (2004b). Towards a bounded pareto-coevolution archive. In *Proceedings of the 2004 Congress on Evolutionary Computation*, pages 2341–2348.

[de Jong, 2005] de Jong, E. D. (2005). The maxsolve algorithm for coevolution. In *Proceedings of the 2005 Genetic and Evolutionary Computation Conference*. ACM.

[de Jong and Bucci, 2006] de Jong, E. D. and Bucci, A. (2006). Deca: Dimension extracting coevolutionary algorithm. In *Proceedings of the 2006 Genetic and Evolutionary Computation Conference*.

[de Jong and Pollack, 2003] de Jong, E. D. and Pollack, J. B. (2003). Learning the ideal evaluation function. In [Cantú-Paz et al., 2003], pages 277–288.

[Ficici, 2006] Ficici, S. (2006). A game-theoretic investigation of selection methods in two-population coevolution. In *Proceedings of the 2006 Genetic and Evolutionary Computation Conference*. ACM Press.

[Ficici and Pollack, 2007] Ficici, S. and Pollack, J. (2007). Evolutionary dynamics of finite populations in games with polymorphic fitness-equilibria. *Journal of Theoretical Biology*. forthcoming.

[Ficici, 2004] Ficici, S. G. (2004). *Solution Concepts in Coevolutionary Algorithms*. PhD thesis, Brandeis University.

[Ficici, 2005] Ficici, S. G. (2005). Monotonic solution concepts in coevolution. In *Proceedings of the 2005 Genetic and Evolutionary Computation Conference*, pages 499–506.

[Ficici et al., 2005] Ficici, S. G., Melnik, O., and Pollack, J. B. (2005). A game-theoretic and dynamical-systems analysis of selection methods in coevolution. *IEEE Transactions on Evolutionary Computation*, 9(6):580–602.

[Ficici and Pollack, 2001] Ficici, S. G. and Pollack, J. B. (2001). Pareto optimality in coevolutionary learning. In Kelemen, J. and Sosík, P., editors, *Sixth European Conference on Artificial Life (ECAL 2001)*, pages 316–325. Springer.

[Ficici and Pollack, 2003] Ficici, S. G. and Pollack, J. B. (2003). A game-theoretic memory mechanism for coevolution. In [Cantú-Paz et al., 2003], pages 286–297.

[Floreano and Nolfi, 1997] Floreano, D. and Nolfi, S. (1997). God save the Red Queen! Competition in co-evolutionary robotics. In Koza, J. R., Deb, K., Dorigo, M., Fogel, D. B., Garzon, M., Iba, H., and Riolo, R. L., editors, *Proceedings of the Second Conference on Genetic Programming*, pages 398–406. Morgan Kaufmann.

[Fogel and Fogel, 1995] Fogel, D. B. and Fogel, G. B. (1995). Evolutionary stable strategies are not always stable under evolutionary dynamics. In *Evolutionary Programming IV*, pages 565–577.

[Fogel et al., 1997] Fogel, D. B., Fogel, G. B., and Andrews, P. C. (1997). On the instability of evolutionary stable states. *BioSystems*, 44:135–152.

[Fogel et al., 1998] Fogel, G. B., Andrews, P. C., and Fogel, D. B. (1998). On the instability of evolutionary stable strategies in small populations. *Ecological Modelling*, 109:283–294.

[Fudenberg and Tirole, 1998] Fudenberg, D. and Tirole, J. (1998). *Game Theory*. MIT Press.

[Hofbauer and Sigmund, 1998] Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

[Jansen and Wiegand, 2004] Jansen, T. and Wiegand, R. P. (2004). The cooperative coevolutionary (1+1) ea. *Evolutionary Computation*, 12(4):405–434.

[Juillé, 1999] Juillé, H. (1999). *Methods for Statistical Inference: Extending the Evolutionary Computation Paradigm*. PhD thesis, Brandeis University.

[Juillé and Pollack, 1996a] Juillé, H. and Pollack, J. (1996a). Co-evolving intertwined spirals. In Fogel, L. J., Angeline, P. J., and Bäck, T., editors, *Proceedings of the Fifth Annual Conference on Evolutionary Programming*, pages 461–468. MIT Press.

[Juillé and Pollack, 1996b] Juillé, H. and Pollack, J. (1996b). Dynamics of coevolutionary learning. In *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 526–534. MIT Press.

[Juillé and Pollack, 1998] Juillé, H. and Pollack, J. B. (1998). Coevolving the "ideal" trainer: Application to the discovery of cellular automata rules. In Koza, J. R. et al., editors, *Proceedings of the Third Annual Genetic Programming Conference*, pages 519–527. Morgan Kaufmann.

[Juillé and Pollack, 2000] Juillé, H. and Pollack, J. B. (2000). Coevolutionary learning and the design of complex systems. *Advances in Complex Systems*, 2(4):371–393.

[Liekens et al., 2004] Liekens, A., Eikelder, H., and Hilbers, P. (2004). Predicting genetic drift in 2×2 games. In *Proceedings from the Genetic and Evolutionary Computation Conference*, pages 549–560.

[Maynard-Smith, 1982] Maynard-Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.

[Maynard-Smith and Price, 1973] Maynard-Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.

[Miller and Cliff, 1994] Miller, G. F. and Cliff, D. (1994). Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D., Husbands, P., Meyer, J.-A., and Wilson, S. W., editors, *From Animals to Animats III*, pages 411–420. MIT Press.

[Monroy et al., 2006] Monroy, G. A., Stanley, K. O., and Miikkulainen, R. (2006). Coevolution of neural networks using a layered pareto archive. In *GECCO '06: Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pages 329–336. ACM Press.

[Moriarty and Miikkulainen, 1997] Moriarty, D. and Miikkulainen, R. (1997). Forming neural networks through efficient and adaptive coevolution. *Evolutionary Computation*, 5(4):373–399.

[Nash, 1951] Nash, J. (1951). Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295. Second Series.

[Noble and Watson, 2001] Noble, J. and Watson, R. A. (2001). Pareto coevolution: Using performance against coevolved opponents in a game as dimensions for pareto selection. In Spector, L. et al., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2001)*, pages 493–500. Morgan Kaufmann.

[Panait et al., 2004] Panait, L., Wiegand, R. P., and Luke, S. (2004). A sensitivity analysis of a cooperative coevolutionary algorithm biased for optimization. In Kalyanmoy Deb et al., editor, *Genetic and Evolutionary Computation Conference – GECCO 2004*, volume 3102 of *Lecture Notes in Computer Science*, pages 573–584. Springer.

[Popovici and De Jong, 2005a] Popovici, E. and De Jong, K. A. (2005a). A dynamical systems analysis of collaboration methods in cooperative co-evolution. In *Proceedings of the AAAI 2005 Fall Symposium on Coevolutionary and Coadaptive Systems*. AAAI Press.

[Popovici and De Jong, 2005b] Popovici, E. and De Jong, K. A. (2005b). Understanding cooperative co-evolutionary dynamics via simple fitness landscapes. In *Proceedings of the Genetic and Evolutionary Computation Conference*.

[Popovici and De Jong, 2006a] Popovici, E. and De Jong, K. A. (2006a). The dynamics of the best individuals in co-evolution. *Natural Computing*, 5(3):229–255.

[Popovici and De Jong, 2006b] Popovici, E. and De Jong, K. A. (2006b). The effects of interaction frequency on the optimization performance of cooperative coevolution. In *Proceedings of the 2006 Genetic and Evolutionary Computation Conference*. ACM Press.

[Popovici and De Jong, 2006c] Popovici, E. and De Jong, K. A. (2006c). Sequential versus parallel cooperative coevolutionary algorithms for optimization. In *Proceedings of the 2006 Congress on Evolutionary Computation*. IEEE Press.

[Potter and Jong, 1994] Potter, M. A. and Jong, K. A. D. (1994). A cooperative coevolutionary approach to function optimization. In Davidor, Y. and Schwefel, H.-P., editors, *Proceedings of the Third Conference on Parallel Problems Solving from Nature (PPSN 3)*, pages 249–257. Springer-Verlag.

[Potter and Jong, 2000] Potter, M. A. and Jong, K. A. D. (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8(1):1–29.

[Rosin and Belew, 1997] Rosin, C. D. and Belew, R. (1997). New methods for competitive co-evolution. *Evolutionary Computation*, 5(1):1–29.

[Schmitt, 2003a] Schmitt, L. M. (2003a). Coevolutionary convergence to a global optima. In [Cantú-Paz et al., 2003], pages 373–374.

[Schmitt, 2003b] Schmitt, L. M. (2003b). Theory of coevolutionary genetic algorithms. In Guo, M. and Yang, L. T., editors, *International Symposium on Parallel and Distributed Processing and Applications (ISPA)*, volume 2745 of *Lecture Notes in Computer Science*, pages 285–293. Springer.

[Schoenauer et al., 2000] Schoenauer, M. et al., editors (2000). *Parallel Problem Solving from Nature VI*. Springer-Verlag.

[Stanley and Miikkulainen, 2002a] Stanley, K. O. and Miikkulainen, R. (2002a). The dominance tournament method of monitoring progress in coevolution. In Barry, A., editor, *2002 Genetic and Evolutionary Computation Conference Workshop Program*, pages 242–248.

[Stanley and Miikkulainen, 2002b] Stanley, K. O. and Miikkulainen, R. (2002b). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.

[Stanley and Miikkulainen, 2004] Stanley, K. O. and Miikkulainen, R. (2004). Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21:63–100.

[Strogatz, 1994] Strogatz, S. H. (1994). *Nonlinear Dynamics and Chaos*. Addison-Wesley Publishing Company.

[Watson, 2006] Watson, R. A. (2006). *Compositional Evolution: The Impact of Sex, Symbiosis, and Modularity on the Gradualist Framework of Evolution*. MIT Press.

[Watson and Pollack, 2000] Watson, R. A. and Pollack, J. B. (2000). Symbiotic combination as an alternative to sexual recombination in genetic algorithms. In [Schoenauer et al., 2000], pages 425–434.

[Wiegand, 2003] Wiegand, R. P. (2003). *An Analysis of Cooperative Coevolutionary Algorithms*. PhD thesis, George Mason University.

[Wiegand et al., 2002a] Wiegand, R. P., De Jong, K. A., and Liles, W. C. (2002a). The effects of representational bias on collaboration methods in cooperative coevolution. In *Parallel Problem Solving from Nature, PPSN-VII*, pages 257–268.

[Wiegand et al., 2001] Wiegand, R. P., Liles, W., and De Jong, K. (2001). An empirical analysis of collaboration methods in cooperative coevolutionary algorithms. In *Proceedings from the Genetic and Evolutionary Computation Conference*, pages 1235–1242.

[Wiegand et al., 2002b] Wiegand, R. P., Liles, W., and De Jong, K. (2002b). Analyzing cooperative coevolution with evolutionary game theory. In *Proceedings of the 2002 Congress on Evolutionary Computation CEC2002*, pages 1600–1605. IEEE Press.

[Wiegand et al., 2003] Wiegand, R. P., Liles, W. C., and De Jong, K. (2003). Modeling variation in cooperative coevolution using evolutionary game theory. In *Foundations of Genetic Algorithms 7*, pages 203–220. Morgan Kaufmann.

[Wolpert and Macready, 1997] Wolpert, D. and Macready, W. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82.

[Wolpert and Macready, 2005] Wolpert, D. H. and Macready, W. G. (2005). Coevolutionary free lunches. *IEEE Transactions on Evolutionary Computation*, 9(6):721–735.