Introductory 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 2006

Outline

- Early work
- Notable results
- Canonical coevolutionary algorithms
- Pathologies, Monitoring Progress, Remedies
- What are we really looking for?
- Looking forward

Early Work:

- Samuel 1959, 1967
  - learning checkers through self-play
- Barricelli 1963, Barricelli et al. 1967
  - TacTix (game similar to Nim)
- Axelrod 1987
  - iterated prisoner's dilemma

Some Notable Results

- Several good results have been obtained through the use of self-play in general and coevolution in particular
- Below we sample some of these results
- (Not an exhaustive list!)

Sorting Networks: Hillis 1990

- Learner-teacher paradigm
- Coevolves sorting networks against inputs
- Obtains 61-comparator network (just one more comparator than best known for 16-input problem)

Virtual Creatures: Sims 1994

- Virtual creatures in simulated physics environment
- Pair-wise competitions to gain control over a cube in the middle of the arena
- Coevolution of agent morphology and control
- Variety of interesting body plans and behaviors obtained

Backgammon: Tesauro 1995

- Neural network trained to evaluate board positions; achieves "strong master" level
- Temporal Difference learning used—not coevolution
- But, a compelling demonstration of learning through self-play
- Follow-up work by Pollack & Blair 1998 that uses coevolution

Intertwined Spirals: Juillé & Pollack '96

- Difficult classification problem motivated by study of neural networks
- 194 data points to classify
- Coevolves genetic-program classifiers, where payoff to Player *i* is:
- G(i, j) = #points "covered" by Player *i* that are not covered by Player *j*
- Found modular solutions to problem

CA Rules: Juillé & Pollack 1998

- Density classification task in CA:
  - if #1 > #0 in initial condition, ⇒ all 1
  - otherwise, ⇒ all 0
- No perfect rule exists [Land & Belew 1995]
- Previous best known performance: 82.4% [Andre et al. 1996]
- Discovered rule with performance: 86.3%

HIV Resistance: Rosin et al. '98, '99

- Coevolves (in simulation) HIV-1 virus against anti-virals
- Finds highly resistant forms of HIV-1 protease
- Finds effective protease inhibitors

Checkers: Chellapilla & Fogel '99, '00

- Coevolves weights of neural network used to evaluate game boards
- Combined with four-ply lookahead
- Initial work achieved "Class-A" designation
- Subsequent work produced "Expert" level
- (Just below Master and Grand Master)

Coevolutionary Algorithms

- "Competitive coevolution"
- "Cooperative coevolution" ⇒

  "Compositional coevolution"
- Game theory provides some common ground

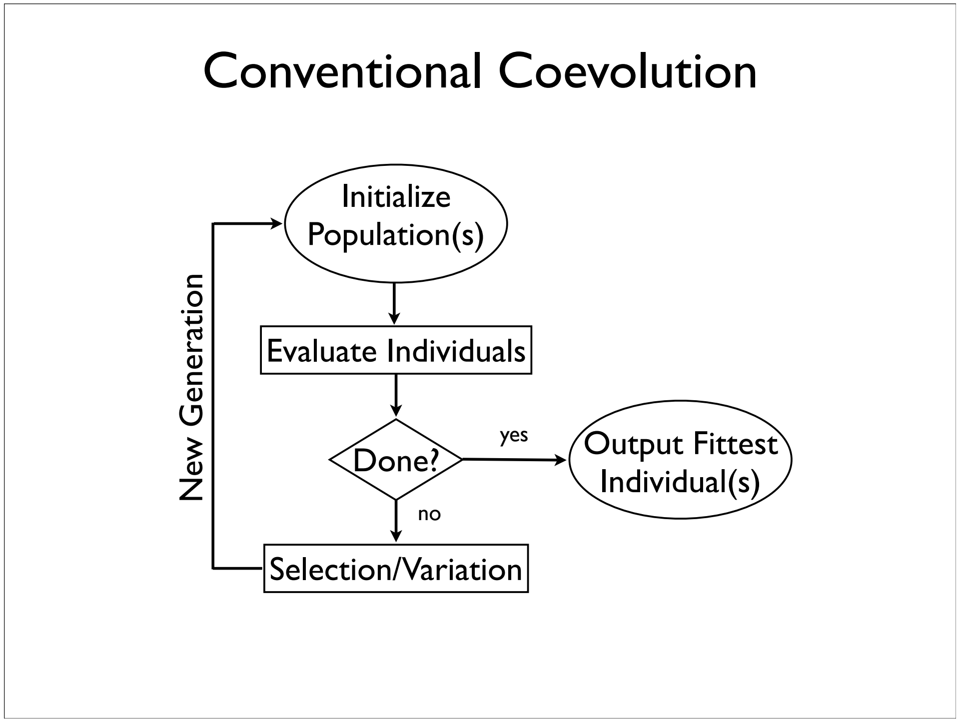
Conventional Coevolution
Initialize Population(s)
Evaluate Individuals
Done?
yes
Output Fittest Individual(s)
no
Selection/Variation
New Generation

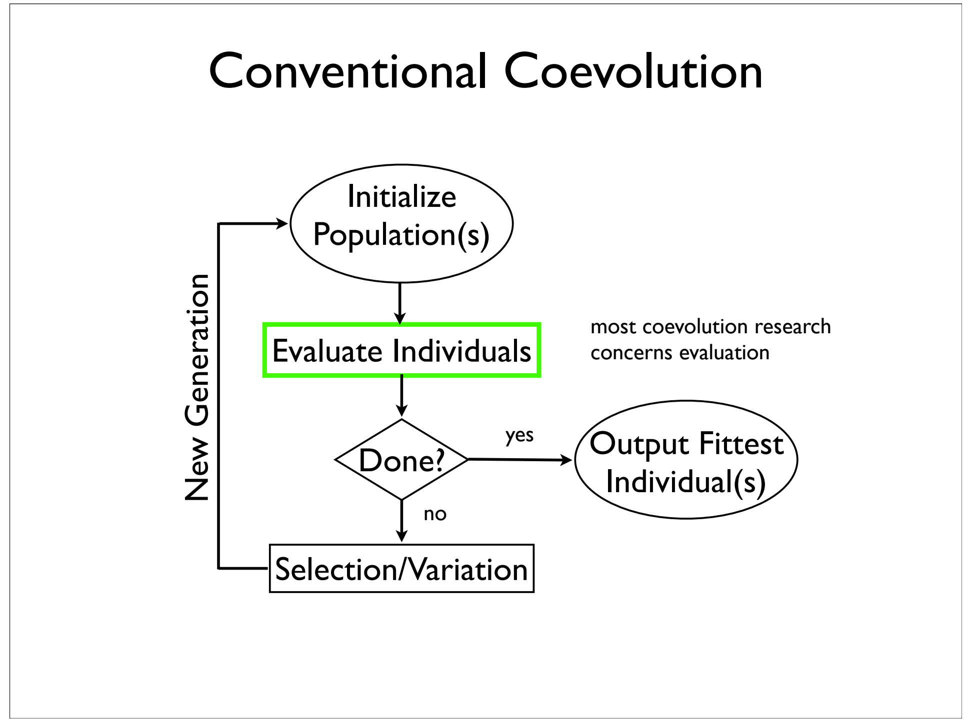
Conventional Coevolution
Initialize Population(s)
Evaluate Individuals
most coevolution research concerns evaluation
Done?
yes
Output Fittest Individual(s)
no
Selection/Variation
New Generation

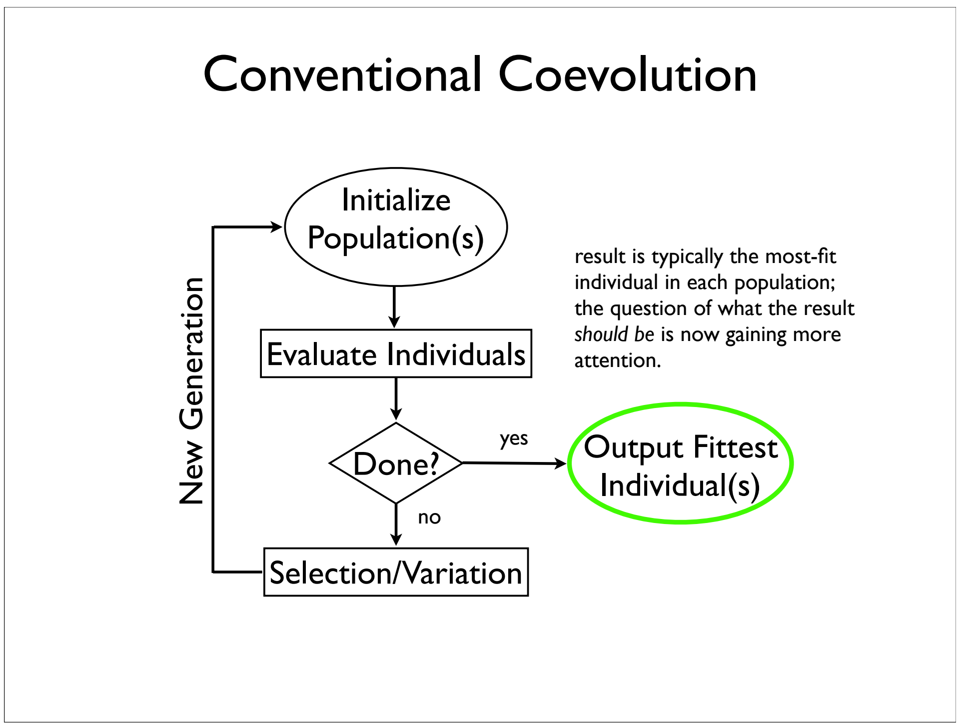
Conventional Coevolution
Initialize Population(s)
Evaluate Individuals
result is typically the most-fit individual in each population; the question of what the result *should be* is now gaining more attention.
Done?
yes
Output Fittest Individual(s)
no
Selection/Variation
New Generation

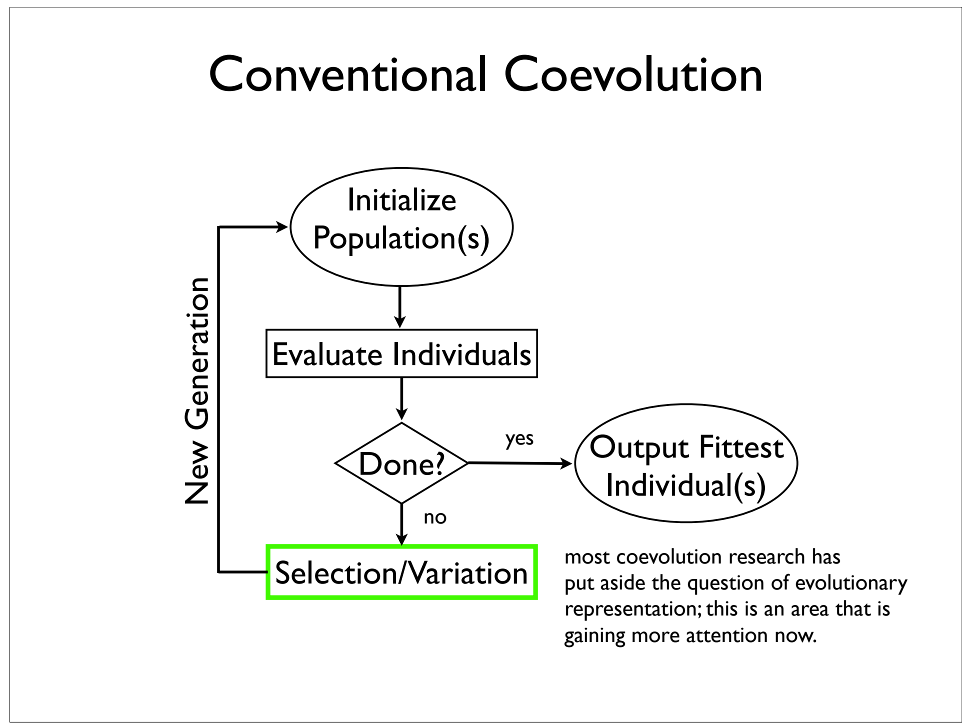
Conventional Coevolution
Initialize Population(s)
Evaluate Individuals
Done?
yes
Output Fittest Individual(s)
no
Selection/Variation
most coevolution research has put aside the question of evolutionary representation; this is an area that is gaining more attention now.
New Generation

Standard Evolutionary Algorithm

- Evaluation uses a fixed objective function

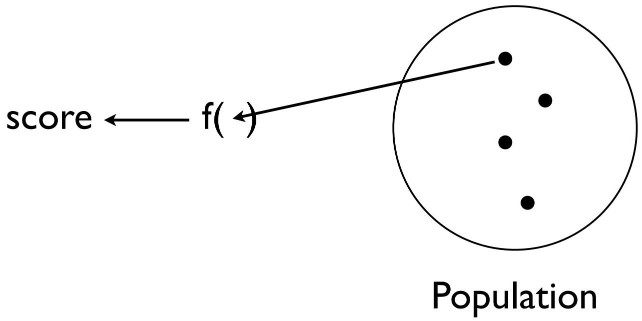

Conventional Coevolution

- Individuals are evaluated by having them interact with each other

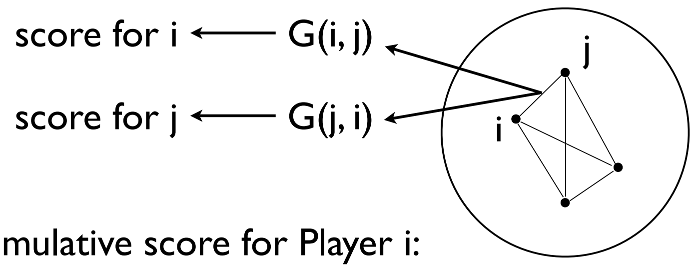

Cumulative score for Player i:
$s_i = \sum_{j=1, N} G(i, j)$

Population

Conventional Coevolution

- In asymmetric games, each member of Pop. 1 interacts with each member of Pop. 2

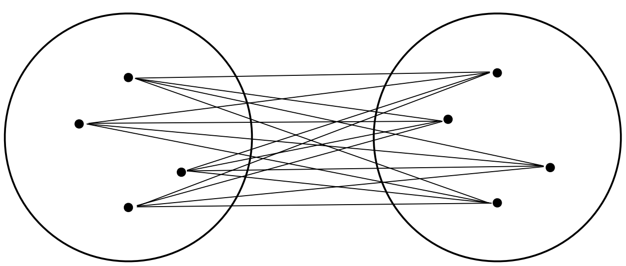

Population 1 Population 2

Interaction Patterns

- All vs. all is "canonical" but expensive
- All vs. previous-best
- Tournament
- See Angeline & Pollack 1993, Sims 1994
- Shared sampling [Rosin & Belew 1997]

All vs. Best

- Individuals interact with "best" individual(s) from previous generation
- Feasible for one or more populations

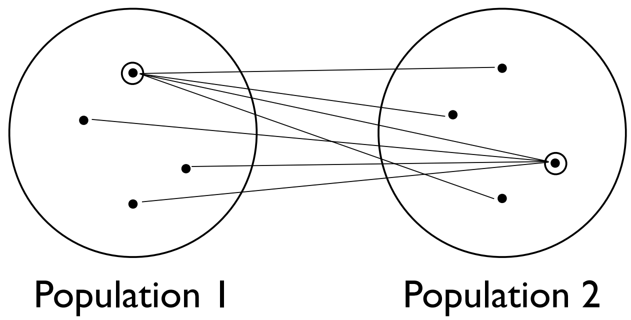

Tournament Evaluation

- Pairwise interactions in single-elimination tournament (single-population)
- Each individual's score determined by how far individual progresses in tournament

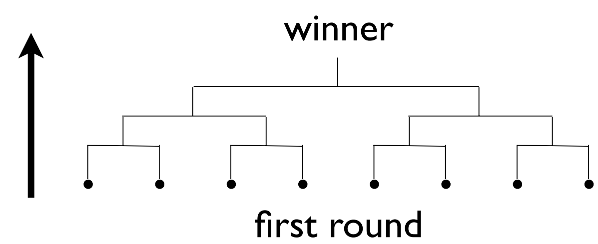

Shared Sampling: Rosin 1997

- Purpose to enhance diversity in evaluation
- Based on their Competitive Fitness Sharing method (discussed below)
- Bias sampling of individuals with whom interaction (during evaluation) takes place
- Sample "redundant" individuals less (relative to uniform); "rare" individuals more
- "Redundant" and "rare" determined by similarity in performance

Spatial Coevolution: Hillis 1990, Pagie & Hogeweg 1997

- Individuals spatially arranged on a lattice
- Individuals interact only with neighbors
- In two-population system, interact with individuals in corresponding neighborhood of other population

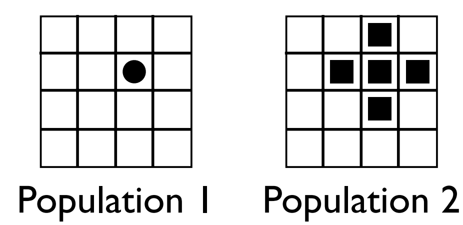

Spatial Coevolution: Hillis 1990, Pagie & Hogeweg 1997

- Spatial structure used to determine who interacts with whom (local neighborhoods)
- Spatial structure also used to control selection; e.g., each individual X is replaced by best individual in the local neighborhood
- Local neighborhood structure helps maintain population diversity; thus can provide help against various pathologies (discussed below)

Cooperative Coevolution

Potter and De Jong 1994

- Decompose problem into simpler sub-tasks that are easier to solve
- Combine sub-solutions to form solution to original problem

$$\text{score for } s_1 = s_2 = \ldots = s_N \longleftarrow G(s_1, s_2, \ldots, s_N)$$

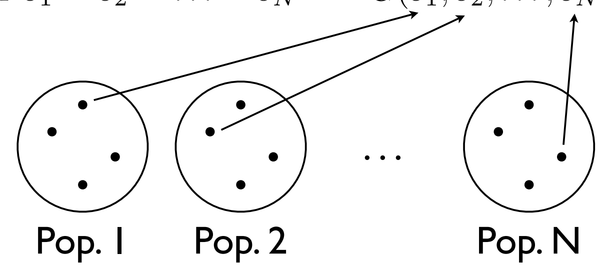

Cooperative Coevolution

Potter and De Jong 1994

- Exhaustive mixing (all vs. all) too expensive
- Sampling is more feasible; cheapest to interact with "best" from each population

$$\text{score for } s_1 \longleftarrow G(s_1, s_2^*, \ldots, s_N^*)$$

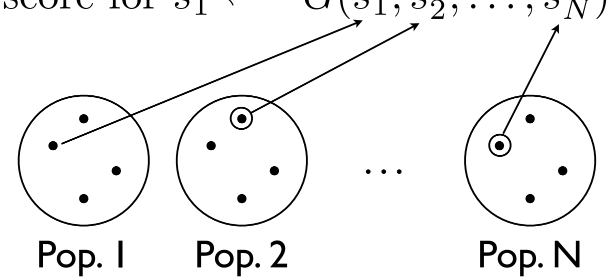

Game Theory and Coevolution

- "Competitive" coevolution epitomized by two-player zero-sum game, e.g., checkers
- "Cooperative" coevolution epitomized by N-player variable-sum coordination game:
  - all players obtain payoff when they play a certain *joint strategy profile*; otherwise they obtain no (or poor) payoff

Pathologies, Monitoring, Remedies

- Early results sparked interest in coevolution, but various pathologies quickly became evident
- Why coevolution fails to produce desired results is often unclear
- We discuss these pathologies and outline several attempts to remedy them

Concept: Gradient

- The evaluation of individuals depends on other, coevolving individuals.
- *Gradient* refers to the evaluational information provided by those coevolving individuals, particularly to the ability to distinguish individuals on the basis of their interactions with coevolving individuals.
- Roughly, gradient allows an algorithm to tell which individuals appear better.

Disengagement

- The event that gradient is lost; i.e., individuals can no longer be distinguished.
- Typically, the algorithm *stalls* or *drifts*, as it can no longer tell which individuals are better.
- Imagine a school with grade inflation. All students receive an A. Then grades can no longer be used to distinguish the better students. The students and the curriculum are disengaged.

Stalling/Drift

- When lack of gradient persists over evolutionary time, *stalling* or *drift* can occur.
- If the algorithm only replaces individuals with strictly better ones (e.g., a hillclimber), it will stall. The population stops changing.
- Otherwise, the algorithm will essentially perform a random walk.

Disengagement, Stalling, and Drift

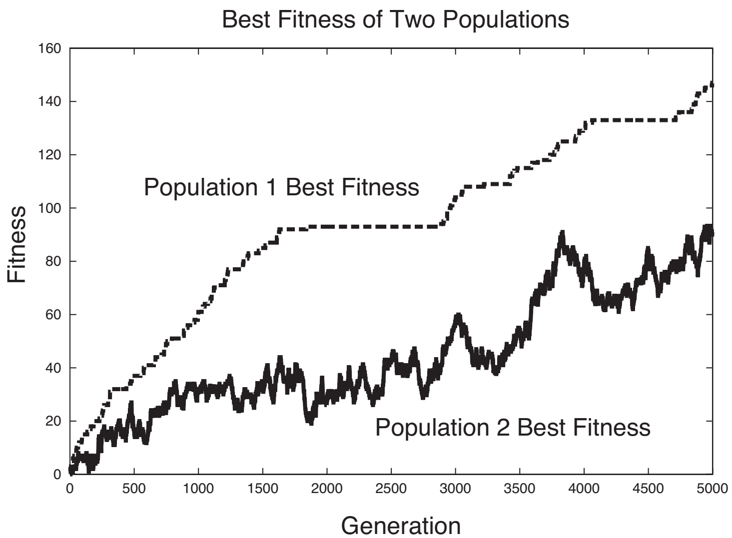

Cycling/Intransitivity

- *Cycling* typically refers to an oscillation in some metric of algorithm performance.
- If there is an offline metric of performance, we may observe the performance of coevolved individuals going up and down through evolutionary time.
- Or, we may observe that present individuals beat some past individuals but lose to others.

Cycling/Intransitivity

- *Intransitivity* is a characteristic of a problem domain.
- Rock-paper-scissors is a canonical example of an intransitive domain.
- Coevolutionary algorithms have been observed to cycle on intransitive domains, but may cycle on any domain.

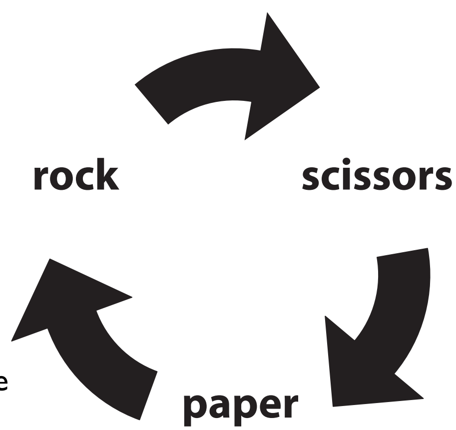

Cycling/Intransitivity

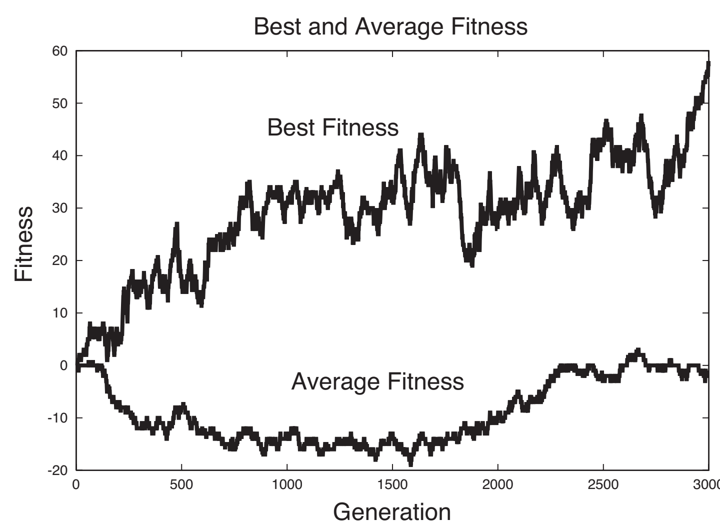

Evolutionary Forgetting

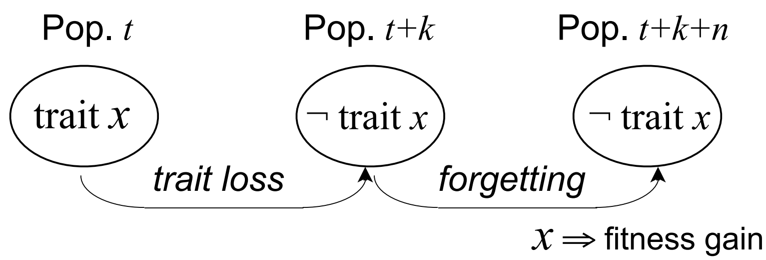

selected against → intransitivity

selection-neutral → drift/focusing

selected for → no elitism

Synopsis

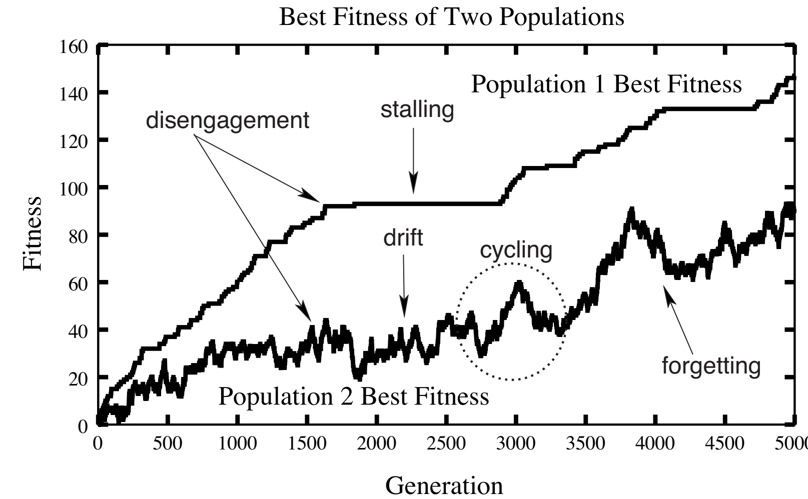

Concept: Underlying Objectives

- Multiobjective algorithms simultaneously optimize several different objective functions.
- Consider "capabilities" as objectives.
- Similarly, coevolutionary domains might have a *set of underlying objectives* that must be optimized to produce good individuals.

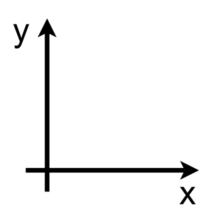

Overspecialization/Focusing

- When individuals improve on some underlying objectives at the expense of others.
- For instance, coevolving game players may *focus* on defeating certain (types of) opponents and not evolve to defeat others.

Overspecialization/Focusing

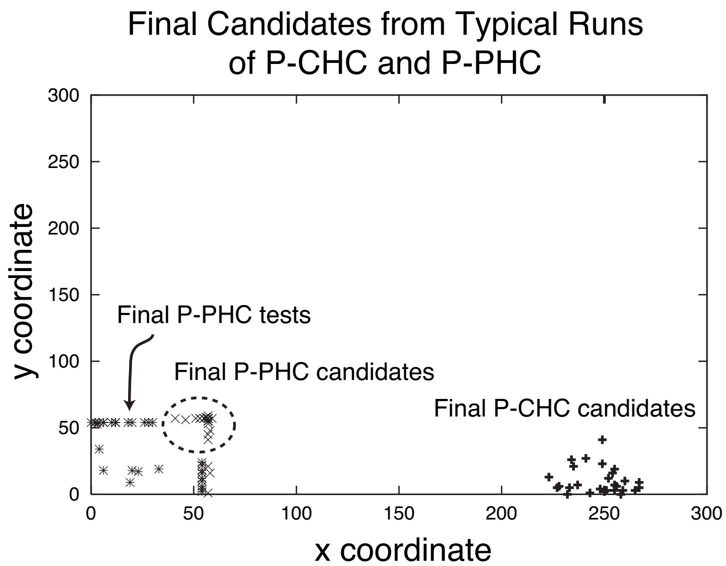

The Red-Queen Effect:

van Valen 1973

- In biological coevolution, the observation that despite constant genetic change, the extinction probability of a species does not change because of changes in the environment.

- In evolutionary computation, the observation that changes which improve the quality of an individual do not increase its selection probability because of changes to other coevolving individuals.

The Red-Queen Effect

- Most troubling is that the Red-Queen Effect prevents us from distinguishing improvement from stall/drift when monitoring an algorithm online. New individuals appear as capable as previous ones relative to the present context.

Relative Overgeneralization

- The phenomenon in Cooperative Coevolution where components that perform well with a large number of other individuals are favored over components that are part of an optimal solution.

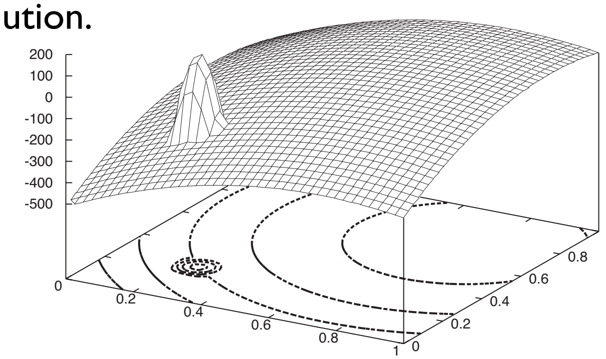

Relationships: Gradient

- *Disengagement* is a loss of gradient.
- *Stalling* or *drifting* can result from a lack of gradient which persists through evolutionary time.
- *Forgetting* or *overspecialization* may result from drift.

Relationships: Underlying Objectives

- *Overspecialization* is focusing on one underlying objective at the expense of others.
- *Cycling* may result from oscillating between two underlying objectives.
- *Relative overgeneralization* has been argued to result from the loss of an underlying objective in Cooperative Coevolution.

About Remedies

- *Forgetting* remedies are typically about distinguishing individuals.
  - If individuals cannot be distinguished, some might be lost to drift and forgetting may occur.
- *Disengagement* remedies have traditionally kept suboptimal individuals in the population.
  - Empirically, greedy algorithms which consolidate around present best tend to disengage.
  - Suboptimal individuals may provide gradient.

Cycling

Ways to address cycling include:

- Fitness sharing
- Memory mechanisms
- Enrich the environment
- Multiple populations

Cycling: Bullock 1995

- "True" coevolution is direct reciprocal evolution between two populations
- "Diffuse" coevolution entails evolutionary change in response to traits in several other populations
- Diffuse coevolution leads to more robust strategies
- Follow-up by Hornby & Mirtich 1999

Cycling: Rosin & Belew 1995

- Zero-sum games (symmetric or asymmetric)
- Competitive fitness sharing
- Score you get against an opponent is divided by sum of all scores obtained against that opponent

Cycling: Rosin & Belew 1995

Standard fitness calculation:

	A	B	C	D	
W	1	1	1	0	$1 + 1 + 1 + 0 = 3$
X	1	1	0	0	$1 + 1 + 0 + 0 = 2$
Y	1	0	0	0	$1 + 0 + 0 + 0 = 1$
Z	0	0	1	1	$0 + 0 + 1 + 1 = 2$

Cycling: Rosin & Belew 1995

Competitive fitness sharing:

	A	B	C	D	
W	1	1	1	0	$1/3 + 1/2 + 1/2 + 0 = 4/3$
X	1	1	0	0	$1/3 + 1/2 + 0 + 0 = 5/6$
Y	1	0	0	0	$1/3 + 0 + 0 + 0 = 1/3$
Z	0	0	1	1	$0 + 0 + 1/2 + 1 = 3/2$
	3	2	2	1	

Cycling: Juillé & Pollack 1996

- Fitness based on unique "covering"
- Individuals in Population 1 interact with opponents in Population 2
- Fitness of an individual determined by comparing performance with other individuals in *same* population
- Points for beating opponents that others do not beat

Cycling: Juillé & Pollack 1996

Standard fitness calculation:

	A	B	C	D
W	1	1	1	0
X	1	1	0	0
Y	1	0	0	0
Z	0	0	1	1

$$1 + 1 + 1 + 0 = 3$$
$$1 + 1 + 0 + 0 = 2$$
$$1 + 0 + 0 + 0 = 1$$
$$0 + 0 + 1 + 1 = 2$$

Cycling: Juillé & Pollack 1996

Covering calculation:

	A	B	C	D
W	1	1	1	0
X	1	1	0	0
Y	1	0	0	0
Z	0	0	1	1

	W	X	Y	Z
W	0	1	2	2
X	0	0	1	2
Y	0	0	0	1
Z	1	2	2	0

Cycling: Juillé & Pollack 1996

Covering calculation:

	W	X	Y	Z
W	0	1	2	2
X	0	0	1	2
Y	0	0	0	1
Z	1	2	2	0

$$0 + 1 + 2 + 2 = 5$$
$$0 + 0 + 1 + 2 = 3$$
$$0 + 0 + 0 + 1 = 1$$
$$1 + 2 + 2 + 0 = 5$$

Cycling: Equilibria & Dynamics

- Rosin & Belew 1997 prove that any fitness equilibrium without fitness sharing is also a fitness equilibrium with fitness sharing (in zero-sum game)
- Juillé & Pollack 1996 show that their "covering" method can lead to stable *polymorphisms*

Cycling: Nolfi & Floreano 1998

- Robotic pursuit and evasion
- Observe cyclic dynamics
- Hypothesize that a more complex environment may dampen cyclic dynamic
- Added obstacles and walls
- Found to provide significant performance boost in some runs
- On average, though, delays onset of cycling

Cycling: Hornby & Mirtich '99

- Virtual pursuit and evasion with simulated physics of wheeled car-like agents
- Round arena with large obstacle in center

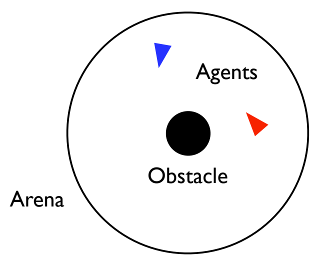

Cycling: Hornby & Mirtich '99

- Use multiple populations for each role of the game (c.f. Bullock 1995)
- Pursuers and evaders obtained under "diffuse" coevolution were more effective than those obtained from "direct" coevolution

Cycling: Hornby & Mirtich '99

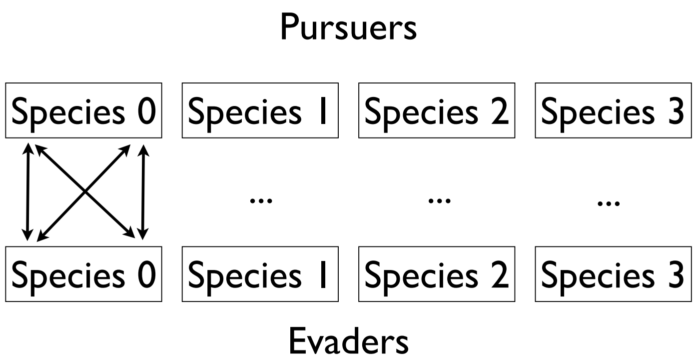

Cycling: Hornby & Mirtich '99

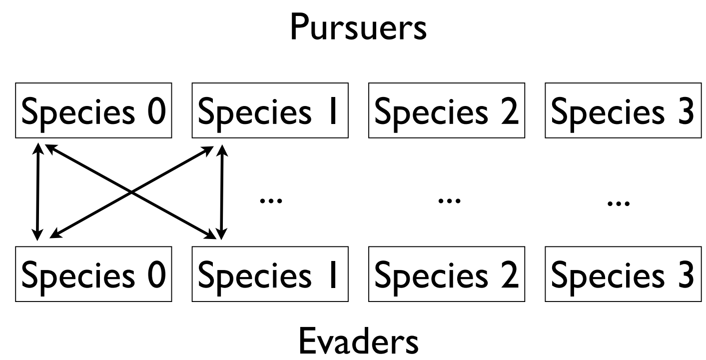

Cycling: Hornby & Mirtich '99

- Runs using direct coevolution exhibit cyclic behavior and disengagement
- Runs using diffuse coevolution stay close to 50% wins for pursuers and evaders

Disengagement: Rosin & Belew '97

- "Phantom parasite" used with competitive fitness sharing to handle disengagement

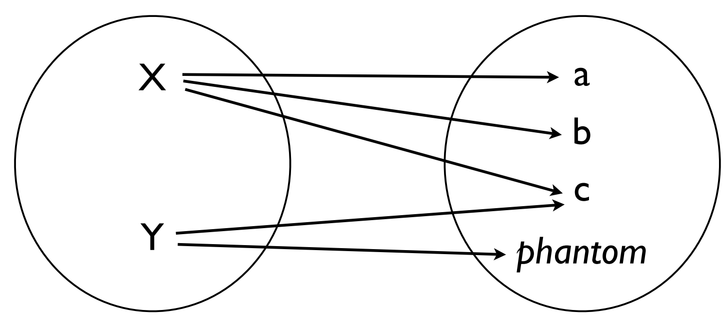

Disengagement: Juillé & Pollack '98

- Density classification task in CA
- Purely competitive evaluation ⇒ cycling
- Competitive fitness sharing ⇒ disengage
- Penalize initial conditions with densities that cause rules to perform near random
- Should be applicable to other domains, e.g., sorting networks

Disengagement: Juillé & Pollack '98

$$f(\mathrm{IC}_j) = \sum_{i=1}^{n_R} W(R_i') \times E(R_i, \rho(\mathrm{IC}_j)) \times \overline{\mathrm{covered}(R_i, \mathrm{IC}_j)})$$

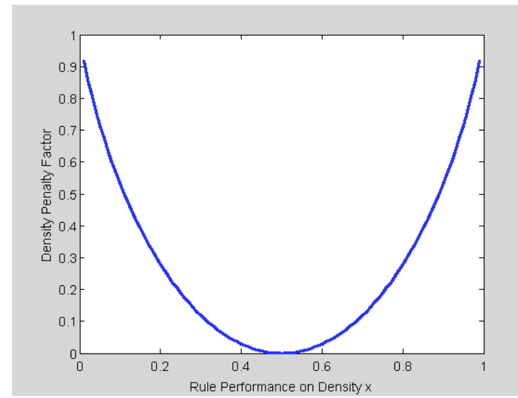

Disengagement: Olsson '98

- Asymmetric zero-sum games
- Evolve only one population, leaving the other population fixed
- Evolve Pop. 1 until individual found that beats all individuals in Pop. 2
- Then evolve Pop. 2 until individual found that beats all individuals in Pop. 1

Disengagement: Paredis '99

- Asymmetric zero-sum games
- Steady-state algorithm
- "X-method" to decide which population gets a new individual

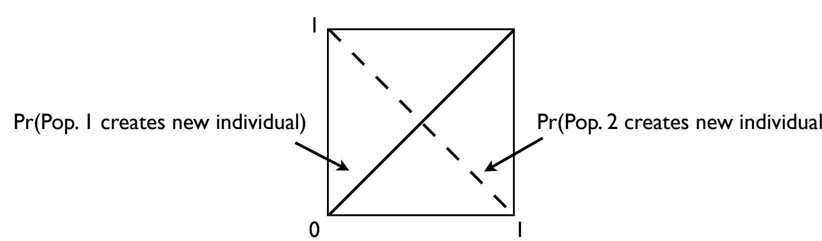

Proportion of Evaluations Won by Pop. 1 Individuals

Disengagement: Cartlidge & Bullock '02

- Moderating "parasite virulence"
- Non-monotonic function of performance

$$f(x,\lambda) = \frac{2x}{\lambda} - \frac{x^2}{\lambda^2}$$

f' indicates peak fitness at λ

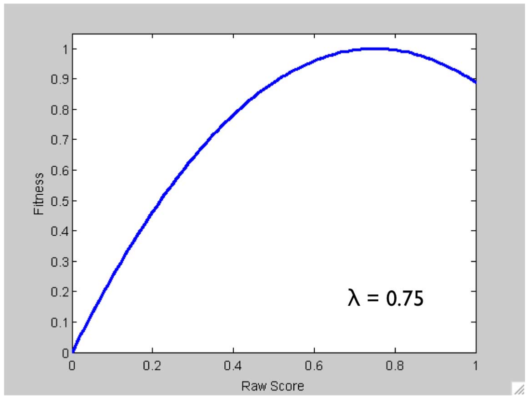

Forgetting: Boyd 1989

- Studies IPD where players can make mistakes
- Tit-For-Tat enters mutual retaliation
- Contrite Tit-For-Tat is resistant to invasion
- All-Cooperate cannot invade via drift
- Noise distinguishes CTFT from All-C

Forgetting: Pollack & Blair '98

- Backgammon naturally resists forgetting
- All aspects of skill are continuously needed
- A simple hill-climber is thus able to achieve fairly impressive performance
- Estimated to achieve skill comparable to TD-Gammon rev. 1992

Forgetting: Memory Mechanisms

- Augment evaluation by interacting with individuals stored in the memory

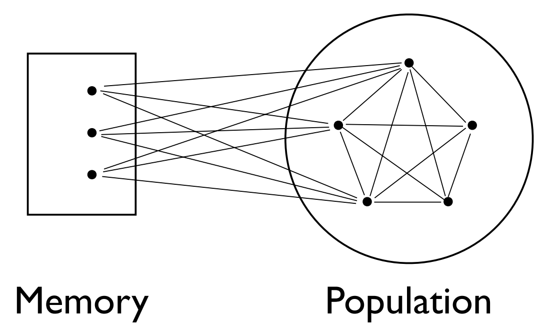

Forgetting: Memory Mechanisms

- Best-of-Generation (BOG) methods
- Most-fit individual from the m most recent generations retained in memory
- Sample n of the m individuals with replacement to augment evaluation of population

Forgetting: Memory Mechanisms

- Sims 1994, Cliff & Miller 1995: m = 1, n = 1
- Potter & De Jong 1994: m = 1, n = 1
- Rosin & Belew 1997: m = ∞, n = 25 & 50
- Nolfi & Floreano 1998a: m = 10, n = 10

Forgetting: Memory Mechanisms

- BOG memory is shown to help
- Broadens selection pressure
- Stabilize algorithm behavior
- Alleviate forgetting

Monitoring: Best Elite Opponent Sims 1994

- Introduces Best Elite Opponent, the model for most subsequent monitoring techniques.
  - individuals of population 1 all compete against most-fit of previous generation from population 2 (best elite).
  - "The most 'interesting' results occurred when the all vs. best competition pattern was used."
- While this is used as a competition pattern, it can be used to monitor progress: track the outcome of each generation's competitions.

Monitoring: *CIAO* Plots
Cliff and Miller 1995

- Introduce *Current Individual vs. Ancestral Opponents* (CIAO) plots.
- Pursuers chase evaders in a 2-D simulated world.
- Two-population coevolutionary algorithm: one of pursuers, the other of evaders.
- "We use the term *fitness ambiguities* to refer to such cases where qualitative trends in time-series of instantaneous fitness measures could feasibly be interpreted as either continuing progress or as a breakdown of the co-evolutionary process."

Monitoring: *CIAO* Plots
Cliff and Miller 1995

- Play current elites against elite opponents from all previous generations.
- Display outcomes in a bitmap image.
- Used as a monitor of progress: if progress has occurred, present elites should be able to defeat elite opponents from previous generations.

Example Pursuer CIAO Plot

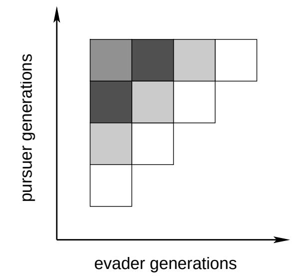

Monitoring: Master Tournament
Floreano and Nolfi 1997

- Also a predator/prey robot experiment.
- Two populations, one of predators, one of prey.
- Introduce *Masters Tournament*: all best predators compete against all best prey.
- Masters Tournament reveals more than CIAO plots:
  - shows at which generation the overall best of each population occurred
  - shows at which generation the most 'interesting' tournaments occur

Monitoring: Hall of Fame
Rosin and Belew 1997

- "To ensure progress, we may want to save individuals for an arbitrarily long time and continue testing against them."
- Introduces *Hall of Fame*
  - stores best of each generation
  - new individuals are tested against a sample of Hall of Fame members.
- While this is used as a memory mechanism, it can also function as a monitor: track performance of new individuals by testing against the members of the hall of fame.

Summary

- Outlined early work and notable results.
- Discussed work on pathological algorithm behavior and proposed remedies
- Raised the question: what do we really want coevolution to do?

What Do We Want Coevolutionary Algorithms To Be Doing?

- **Creating Arms Races** (Ficici & Pollack 1998). "The key to successful coevolutionary learning is a *competitive arms race* between opposed participants."
- **Optimizing Robustness** (Wiegand 2003). "CCEAs…are adaptive optimizers of robustness."

What Do We Want Coevolutionary Algorithms To Be Doing?

- **Complexifying** (Stanley & Miikkulainen 2004). "Complexification encourages continuing innovation by elaborating on existing solutions."
- **Implementing Solution Concepts** (Ficici 2004). "We assert that pathologies in coevolutionary optimization arise when algorithms fail to implement the required (or desired) solution concepts."

Looking Forward

- Solution Concepts
  - Addresses question of what a coevolutionary algorithm should output
- Pareto Coevolution
  - Treats evaluational issues
- Cooperative Coevolution and Robustness
  - Treats composing evolved subparts into wholes
- NEAT and Complexification
  - Treats issues of representation

Looking Forward: Solution Concepts

- Formally specifies which individuals are part of solutions
- Fundamental questions:
  - Are common/intuitive notions of solution reasonable?
  - What solution concepts do we know, and how can we find new ones?
  - Given a solution concept, how do we know if an algorithm actually approximates it?

Looking Forward: Pareto Coevolution

- Focuses on discriminating among and evaluating candidate solutions.
- Fundamental questions:
  - Which individuals are “good,” and why?
  - How do we turn the Pareto Optimal Set into a working solution?
  - How can we deal with the “curse of dimensionality”?
  - Are memory or archive mechanisms necessary?

Looking Forward: Cooperative Coevolution

- Evolving populations of parts which can be assembled into capable wholes.
- Fundamental questions:
  - What makes a good subpart?
  - What makes a good whole?
  - Does CCEA find global optima?
  - Should CCEA be producing “robust” individuals?

Looking Forward: NEAT and Complexification

- Focuses on representing complicated objects in open-ended domains.
- Fundamental questions:
  - Can we remedy pathologies by elaborating on/ complexifying present solutions, versus simply altering them?
  - Can continuous, open-ended progress be achieved?

Introductory Tutorial on Coevolution—References[1]

0.1 Early Work

[Samuel, 1959], [Barricelli, 1963], [Reed et al., 1967], [Samuel, 1967], [Axelrod, 1987]

Notes: The papers by Samuel are available on-line from IBM at `http://www.research.ibm.com/journal/`. The remaining papers are reprinted in [Fogel, 1998].

0.2 Notable Results

[Hillis, 1990], [Angeline and Pollack, 1993], [Sims, 1994], [Tesauro, 1995], [Juillé and Pollack, 1996], [Moriarty and Miikkulainen, 1997], [Eriksson and Olsson, 1997], [Pagie and Hogeweg, 1997], [Rosin and Belew, 1997], [Juillé and Pollack, 1998], [Pollack and Blair, 1998], [Rosin et al., 1998], [Rosin et al., 1999], [Chellapilla and Fogel, 1999], [Chellapilla and Fogel, 2000], [Juillé and Pollack, 2000], [Potter and Jong, 2000], [Pagie and Hogeweg, 2000], [Ficici and Pollack, 2001]

0.3 Pathologies, Monitoring Progress

0.3.1 Cycling, Red-Queen Effect

[Reed et al., 1967], [van Valen, 1973], [Bullock, 1995], [Cliff and Miller, 1995], [Cliff and Miller, 1996], [Floreano and Nolfi, 1997a], [Floreano and Nolfi, 1997b], [Floreano et al., 1998], [Nolfi and Floreano, 1998b], [Nolfi and Floreano, 1998a], [Ficici and Pollack, 1998], [Hornby and Mirtich, 1999], [Watson and Pollack, 2001], [Floreano et al., 2001]

0.3.2 Evolutionary Forgetting

[Boyd, 1989], [Hillis, 1990], [Bullock, 1995], [Floreano and Nolfi, 1997a], [Floreano and Nolfi, 1997b], [Floreano et al., 1998], [Nolfi and Floreano, 1998b], [Nolfi and Floreano, 1998a], [Pollack and Blair, 1998], [Ficici and Pollack, 1998], [Hornby and Mirtich, 1999], [Eriksson and Lindgren, 2001], [Floreano et al., 2001], [Stanley and Miikkulainen, 2002]

0.3.3 Disengagement

[Bullock, 1995], [Rosin and Belew, 1996], [Rosin and Belew, 1997], [Floreano and Nolfi, 1997a], [Floreano and Nolfi, 1997b], [Floreano et al., 1998], [Nolfi and Floreano, 1998b], [Nolfi and Floreano, 1998a], [Ficici and Pollack, 1998], [Olsson, 1998], [Hornby and Mirtich, 1999], [Paredis, 1999], [Paredis, 2000], [Olsson, 2000], [Watson and Pollack, 2001], [Floreano et al., 2001], [Cartlidge and Bullock, 2002]

0.3.4 Overspecialization

[Angeline and Pollack, 1993], [Epstein, 1994], [Pagie and Hogeweg, 1997], [Pagie and Hogeweg, 2000], [Darwen and Yao, 1996], [Watson and Pollack, 2001]

0.3.5 Monitoring Progress

[Reynolds, 1994], [Cliff and Miller, 1995], [Cliff and Miller, 1996], [Ficici and Pollack, 1998], [Floreano and Nolfi, 1997a], [Floreano and Nolfi, 1997b], [Floreano et al., 1998], [Nolfi and Floreano, 1998b], [Nolfi and Floreano, 1998a], [Floreano et al., 2001], [Stanley and Miikkulainen, 2002]

0.4 Methodologies

0.4.1 Competitive Coevolution

[Angeline and Pollack, 1993], [Sims, 1994], [Rosin and Belew, 1995], [Rosin and Belew, 1996], [Rosin and Belew, 1997]

0.4.2 Cooperative Coevolution

[Potter and Jong, 1994], [Moriarty and Miikkulainen, 1997], [Potter and Jong, 2000], [Bull, 2001]

0.4.3 Pareto Coevolution

[Watson and Pollack, 2000], [Ficici and Pollack, 2000], [Ficici and Pollack, 2001], [Noble and Watson, 2001]

0.5 ALife Simulation

[Ray, 1994], [Hraber et al., 1997], [Lenski et al., 1999], [Taylor, 2001]

0.6 No Free Lunch

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

Bibliography

[Adami et al., 1998] Adami, C. et al., editors (1998). *Artificial Life VI (1998)*. MIT Press.

[Alliot et al., 1995] Alliot, J.-M., Lutton, E., Ronald, E., Schoenauer, M., and D., S., editors (1995). *Artificial Evolution (AE 95)*. Springer-Verlag.

[Angeline et al., 1999] Angeline, P., Michalewicz, Z., Schoenauer, M., Yao, X., and Zalzala, A., editors (1999). *Proceedings of the 1999 Congress on Evolutionary Computation (CEC)*. IEEE Press.

[Angeline and Pollack, 1993] Angeline, P. J. and Pollack, J. B. (1993). Competitive environments evolve better solutions for complex tasks. In [Forrest, 1993], pages 264–270.

[Axelrod, 1987] Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In Davis, L., editor, *Genetic Algorithms and Simulated Annealing*, pages 32–41. Morgan Kaufmann.

[Banzhaf et al., 1999] Banzhaf, W., Daida, J. M., Eiben, Á. E., Garzon, M. H., Honavar, V., Jakiela, M. J., and Smith, R. E., editors (1999). *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)*. Morgan Kaufmann.

[Barricelli, 1963] Barricelli, N. A. (1963). Numerical testing of evolution theories Part II: Preliminary tests of performance, symbiogenesis and terrestrial life. *Acta Biotheoretica*, 16(3–4):99–126.

[Belew and Booker, 1991] Belew, R. and Booker, L., editors (1991). *Proceedings of the Fourth International Conference on Genetic Algorithms*. Morgan Kaufmann.

[Blum and Kearns, 1996] Blum, A. and Kearns, M., editors (1996). *Proceedings of the Ninth Annual ACM Conference on Computational Learning Theory (COLT 1996)*. ACM Press.

[Boyd, 1989] Boyd, R. (1989). Mistakes allow evolutionary stability in the repeated prisoner's dilemma game. *Journal of Theoretical Biology*, 136:47–56.

[Brooks and Maes, 1994] Brooks, R. A. and Maes, P., editors (1994). *Artificial Life IV*. MIT Press.

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

[Bullock, 1995] Bullock, S. (1995). Co-evolutionary design: Implications for evolutionary robotics. Technical Report CSRP 384, School of Cognitive and Computing Sciences, University of Sussex. Presented as a poster at the Third European Conference on Artificial Life (ECAL 1995).

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

[Cartlidge and Bullock, 2002] Cartlidge, J. and Bullock, S. (2002). Learning lessons from the common cold: How reducing parasite virulence improves coevolutionary optimization. In [Fogel et al., 2002], pages 1420–1425.

[Chellapilla and Fogel, 1999] Chellapilla, K. and Fogel, D. (1999). Evolving neural networks to play checkers without expert knowledge. *IEEE Transactions on Neural Networks*, 10(6):1382–1391.

[Chellapilla and Fogel, 2000] Chellapilla, K. and Fogel, D. B. (2000). Anaconda defeats Hoyle 6-0: A case study competing an evolved checkers program against commercially available software. In [Zalzala et al., 2000], pages 857–863.

[Cliff et al., 1994] Cliff, D., Husbands, P., Meyer, J.-A., and Wilson, S. W., editors (1994). *From Animals to Animats III*. MIT Press.

[Cliff and Miller, 1995] Cliff, D. and Miller, G. F. (1995). Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In [Morán et al., 1995], pages 200–218.

[Cliff and Miller, 1996] Cliff, D. and Miller, G. F. (1996). Co-evolution of pursuit and evasion II: Simulation methods and results. In [Maes et al., 1996], pages 506–515.

[Cohen and Hirsh, 1994] Cohen, W. W. and Hirsh, H., editors (1994). *Machine Learning: Proceedings of the Eleventh International Conference*. Morgan Kaufmann.

[Darwen and Yao, 1996] Darwen, P. and Yao, X. (1996). Automatic modularization by speciation. In [Fukuda et al., 1996], pages 88–93.

[De Jong, 1992] De Jong, K. A. (1992). Are genetic algorithms function optimizers? In Männer, R. and Manderick, B., editors, *Parallel Problem Solving from Nature II*, pages 3–13. Elsevier.

[De Jong, 1993] De Jong, K. A. (1993). Genetic algorithms are NOT function optimizers. In Whitely, L. D., editor, *Foundations of Genetic Algorithms 2 (FOGA)*, pages 5–18. Morgan Kaufmann.

[De Jong et al., 2003] De Jong, K. A., Poli, R., and Rowe, J. E., editors (2003). *Proceedings of the Foundations of Genetic Algorithms 2003 Workshop (FOGA 7)*. Morgan Kaufmann Publishers.

[Epstein, 1994] Epstein, S. L. (1994). Toward an ideal trainer. *Machine Learning*, 15(3):251–277.

[Eriksson and Lindgren, 2001] Eriksson, A. and Lindgren, K. (2001). Evolution of strategies in repeated stochastic games with full information of the payoff matrix. In [Spector et al., 2001], pages 853–859.

[Eriksson and Olsson, 1997] Eriksson, R. and Olsson, B. (1997). Cooperative coevolution in inventory control optimisation. In Smith, G., Steele, N., and Albrecht, R., editors, *Proceedings of the Third International Conference on Artificial Neural Networks and Genetic Algorithms*. Springer.

[Eshelman, 1995] Eshelman, L. J., editor (1995). *Proceedings of the Sixth International Conference on Genetic Algorithms*. Morgan Kaufmann.

[Fawcett and Mishra, 2003] Fawcett, T. and Mishra, N., editors (2003). *Machine Learning: Proceedings of the Twentieth International Conference (ICML 2003)*. AAAI Press.

[Ficici and Pollack, 1998] Ficici, S. G. and Pollack, J. B. (1998). Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In Adami, C. et al., editors, *Proceedings of the Sixth Conference on Artificial Life*, pages 238–247. MIT Press.

[Ficici and Pollack, 2000] Ficici, S. G. and Pollack, J. B. (2000). A game-theoretic approach to the simple coevolutionary algorithm. In [Schoenauer et al., 2000], pages 467–476.

[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.

[Floreano and Nolfi, 1997a] Floreano, D. and Nolfi, S. (1997a). Adaptive behavior in competing co-evolving species. In [Husbands and Harvey, 1997], pages 378–387.

[Floreano and Nolfi, 1997b] Floreano, D. and Nolfi, S. (1997b). God save the Red Queen! Competition in co-evolutionary robotics. In [Koza et al., 1997], pages 398–406.

[Floreano et al., 1998] Floreano, D., Nolfi, S., and Mondada, F. (1998). Competitive co-evolutionary robotics: From theory to practice. In [Pfeifer et al., 1998], pages 515–524.

[Floreano et al., 2001] Floreano, D., Nolfi, S., and Mondada, F. (2001). Co-evolution and ontogenetic change in competing robots. In Patel, M., Honavar, V., and Balakrishnan, K., editors, *Advances in the Evolutionary Synthesis of Intelligent Agents*, chapter 10. MIT Press.

[Fogarty, 1994] Fogarty, T. C., editor (1994). *Evolutionary Computing (AISB Workshop)*. Springer.

[Fogel, 1998] Fogel, D. B. (1998). *Evolutionary Computation: The Fossil Record*. IEEE Press.

[Fogel et al., 2002] Fogel, D. B., El-Sharkawi, M. A., and Yao, X., editors (2002). *Proceedings of the 2002 Congress on Evolutionary Computation*. IEEE Press.

[Fogel et al., 1996] Fogel, L. J., Angeline, P. J., and Bäck, T., editors (1996). *Proceedings of the Fifth Annual Conference on Evolutionary Programming*. MIT Press.

[Forrest, 1993] Forrest, S., editor (1993). *Proceedings of the Fifth International Conference on Genetic Algorithms*. Morgan Kaufmann.

[Freund and Wolter, 1991] Freund, H. and Wolter, R. (1991). Evolution of bit strings: Some preliminary results. *Complex Systems*, 5:279–298.

[Fukuda et al., 1996] Fukuda, T., Furuhashi, T., Bäck, T., Kitano, H., and Michalewicz, Z., editors (1996). *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation*. IEEE Press.

[Grefenstette, 1987] Grefenstette, J. J., editor (1987). *Proceedings of the Second International Conference on Genetic Algorithms*. Lawrence Earlbaum Associates.

[Grefenstette, 1988] Grefenstette, J. J., editor (1988). *Proceedings of the First International Conference on Genetic Algorithms (1985 Carnegie-Mellon University)*. Lawrence Earlbaum Associates.

[Harvey, 1997] Harvey, I. (1997). Cognition is not computation: Evolution is not optimisation. In Gerstner, W., Germond, A., Hasler, M., and Nicoud, J.-D., editors, *Proceedings of the 7th International Conference on Artificial Neural Networks (ICANN97)*, pages 685–690. Springer-Verlag.

[Hillis, 1990] Hillis, D. (1990). Co-evolving parasites improves simulated evolution as an optimization procedure. *Physica D*, 42:228–234. Reprinted in [Hillis, 1992].

[Hillis, 1992] Hillis, D. (1992). Co-evolving parasites improves simulated evolution as an optimization procedure. In [Langton et al., 1992], pages 313–324.

[Hornby and Mirtich, 1999] Hornby, G. and Mirtich, B. (1999). Diffuse versus true coevolution in a physics-based world. In Banzhaf, Daida, Eiben, Garzon, Honavar, Jakiela, and Smith, editors, *Proceedings of 1999 Genetic and Evolutionary Computation Conference*, pages 1305–1312. Morgan Kaufmann.

[Hraber et al., 1997] Hraber, P. T., Jones, T., and Forrest, S. (1997). The ecology of echo. *Artificial Life*, 3(3):165–190.

[Husbands and Harvey, 1997] Husbands, P. and Harvey, I., editors (1997). *Proceedings of the Fourth European Conference on Artificial Life*. MIT Press.

[Juillé and Pollack, 1996] Juillé, H. and Pollack, J. (1996). Dynamics of co-evolutionary 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.

[Koza et al., 1997] Koza, J. R., Deb, K., Dorigo, M., Fogel, D. B., Garzon, M., Iba, H., and Riolo, R. L., editors (1997). *Proceedings of the Second Conference on Genetic Programming*. Morgan Kaufmann.

[Langdon et al., 2002] Langdon, W. B. et al., editors (2002). *Proceedings of the 2002 Genetic and Evolutionary Computation Conference (GECCO)*. Morgan Kaufmann.

[Langton et al., 1992] Langton, C., Taylor, C., Farmer, J., and Rasmussen, S., editors (1992). *Artificial Life II: Proceedings of the 1990 Workshop on Artificial Life*, SFI Studies in the Sciences of Complexity, vol. X. Addison-Wesley.

[Langton, 1989] Langton, C. G., editor (1989). *Artificial Life I (1987)*. Addison-Wesley.

[Langton, 1994] Langton, C. G., editor (1994). *Artificial Life III (1992)*. Addison-Wesley.

[Langton and Shimohara, 1997] Langton, C. G. and Shimohara, K., editors (1997). *Artificial Life V (1996)*. MIT Press.

[Lenski et al., 1999] Lenski, R. E., Ofria, C., Collier, T. C., and Adami, C. (1999). Genome complexity, robustness and genetic interactions in digital organisms. *Nature*, 400(6745):661–664.

[Lindgren, 1992] Lindgren, K. (1992). Evolutionary phenomena in simple dynamics. In [Langton et al., 1992], pages 295–312.

[Maes et al., 1996] Maes, P. et al., editors (1996). *From Animals to Animats IV*. MIT Press.

[Meyer et al., 1993] Meyer, J.-A., Roitblat, H. L., and Wilson, S. W., editors (1993). *From Animals to Animats II (1992)*. MIT Press.

[Meyer and Wilson, 1991] Meyer, J.-A. and Wilson, S. W., editors (1991). *From Animals to Animats (1990)*. MIT Press.

[Morán et al., 1995] Morán, F. et al., editors (1995). *Proceedings of the Third European Conference on Artificial Life*. Springer.

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

[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.

[Nolfi and Floreano, 1998a] Nolfi, S. and Floreano, D. (1998a). Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution? *Artificial Life*, 4(4):311–335.

[Nolfi and Floreano, 1998b] Nolfi, S. and Floreano, D. (1998b). How coevolution can enhance the adaptation power of artificial evolution: Implications for evolutionary robotics. In Husbands, P. and Meyer, J.-A., editors, *Proceedings of the First European Workshop on Evolutionary Robotics (EvoRobot98)*, pages 22–38. Springer.

[Olsson, 1998] Olsson, B. (1998). A host-parasite genetic algorithm for asymmetric tasks. In Nédellec, C. and Rouveirol, C., editors, *Proceedings of the Ninth European Conference on Machine Learning*, pages 346–351. Springer.

[Olsson, 2000] Olsson, B. (2000). *NK*-landscapes as test functions for evaluation of host-parasite algorithms. In [Schoenauer et al., 2000], pages 487–496.

[Pagie and Hogeweg, 1997] Pagie, L. and Hogeweg, P. (1997). Evolutionary consequences of coevolving targets. *Evolutionary Computation*, 5(4):401–418.

[Pagie and Hogeweg, 2000] Pagie, L. and Hogeweg, P. (2000). Information integration and red queen dynamics in coevolutionary optimization. In [Zalzala et al., 2000], pages 1260–1267.

[Paredis, 1999] Paredis, J. (1999). Coevolution, memory, and balance. In Dean, T., editor, *International Joint Conference on Artificial Intelligence (IJCAI 99)*. Morgan Kaufmann.

[Paredis, 2000] Paredis, J. (2000). Towards balanced coevolution. In [Schoenauer et al., 2000], pages 497–506.

[Pfeifer et al., 1998] Pfeifer, R. et al., editors (1998). *From Animals to Animats V*. MIT Press.

[Pollack and Blair, 1998] Pollack, J. B. and Blair, A. D. (1998). Co-evolution in the successful learning of backgammon strategy. *Machine Learning*, 32(3):225–240.

[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.

[Ray, 1994] Ray, T. S. (1994). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life*, 1(1/2):195–226.

[Reed et al., 1967] Reed, J., Toombs, R., and Barricelli, N. A. (1967). Simulation of biological evolution and machine learning: I. selection of self-reproducing numeric patterns by data processing machines, effects of hereditary control, mutation type and crossing. *Journal of Theoretical Biology*, 17:319–342.

[Reynolds, 1994] Reynolds, C. W. (1994). Competition, coevolution and the game of tag. In [Brooks and Maes, 1994], pages 59–69.

[Rosin and Belew, 1995] Rosin, C. D. and Belew, R. (1995). Methods for competitive co-evolution: Finding opponents worth beating. In [Eshelman, 1995], pages 373–381.

[Rosin and Belew, 1996] Rosin, C. D. and Belew, R. (1996). A competitive approach to game learning. In [Blum and Kearns, 1996], pages 292–302.

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

[Rosin et al., 1998] Rosin, C. D., Belew, R. K., Morris, Garrett M.and Olson, A. J., and Goodsell, D. S. (1998). Computational coevolution of antiviral drug resistance. *Artificial Life*, 4:41–59.

[Rosin et al., 1999] Rosin, C. D., Belew, R. K., Morris, G. M., Olson, A. J., and Goodsell, D. S. (1999). Coevolutionary analysis of resistance-evading peptidomimetic inhibitors of hiv-1 protease. *Proc. Natl. Acad. Sci. USA*, 96:1369–1374.

[Samuel, 1959] Samuel, A. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3):211–229.

[Samuel, 1967] Samuel, A. (1967). Some studies in machine learning using the game of checkers. II—Recent progress. *IBM Journal of Research and Development*, 11(6):601–617.

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

[Sims, 1994] Sims, K. (1994). Evolving 3d morphology and behavior by competition. In [Brooks and Maes, 1994], pages 28–39.

[Spector et al., 2001] Spector, L. et al., editors (2001). *Proceedings of the 2001 Genetic and Evolutionary Computation Conference (GECCO 2001)*.

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

[Taylor, 2001] Taylor, T. (2001). Creativity in evolution: Individuals, interactions and environments. In Bentley, P. J. and Corne, D. W., editors, *Creative Evolutionary Systems*, pages 79–108. Morgan Kaufmann.

[Tesauro, 1995] Tesauro, G. (1995). Temporal difference learning and TD-gammon. *Communications of the ACM*, 38(3):58–68.

[van Valen, 1973] van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1:1–30.

[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.

[Watson and Pollack, 2001] Watson, R. A. and Pollack, J. B. (2001). Coevolutionary dynamics in a minimal substrate. In [Spector et al., 2001], pages 702–709.

[Whitley et al., 2000] Whitley, L. D. et al., editors (2000). *Proceedings of the 2000 Genetic and Evolutionary Computation Conference*. 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.

[Zalzala et al., 2000] Zalzala, A. et al., editors (2000). *Proceedings of the 2000 Congress on Evolutionary Computation*. IEEE Press.