Co-evolutionary Design II

Jason Noble

jasonn@comp.leeds.ac.uk

Biosystems group, School of Computing

Last time

- Co-evolution and “arms races” in nature
- A co-evolutionary GA
- Co-evolved sorting networks
- Further applications in biology and engineering

This time

- Things that can go wrong in a co-evolutionary GA:
  1. disengagement
  2. over-specialization
  3. relativism
- Parasite virulence
- Diffuse and true co-evolution
- Co-evolution and multi-objective optimization

Advantages of co-evolution

Compared to a static fitness function, co-evolution supposedly offers:

- A reachable optimization target
- A relevant target
- A moving target

But. . .

- Co-evolutionary algorithms don’t always work
- Sometimes the hoped-for arms races don’t seem to take off
- Why?

A model system for co-evolution

- Watson and Pollack (2001) looked at "co-evolution in a minimal substrate."
- Trivial task: selection pressure for high integers (0–100).
- Distinction between subjective and objective fitness.

When everything works

- Two populations of 25 individuals.
- Fitness assessed by competing against 15 individuals from the other population.
- Co-evolution successfully pushes both populations to the maximum objective fitness.

Normal coevolution

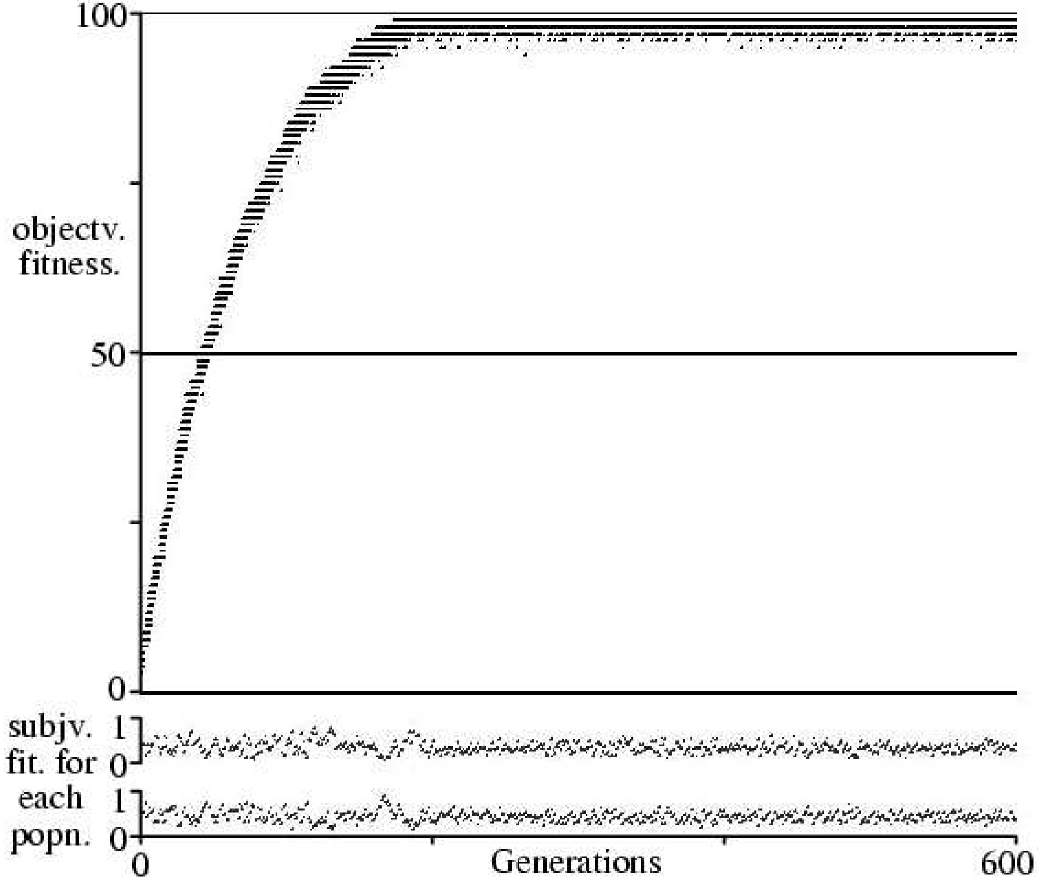

Disengagement

- What if fitness was assessed by competing against just one individual from the other population?
- The co-evolving populations periodically *disengage*, and fail to provide a fitness gradient for their opponents.

Disengagement

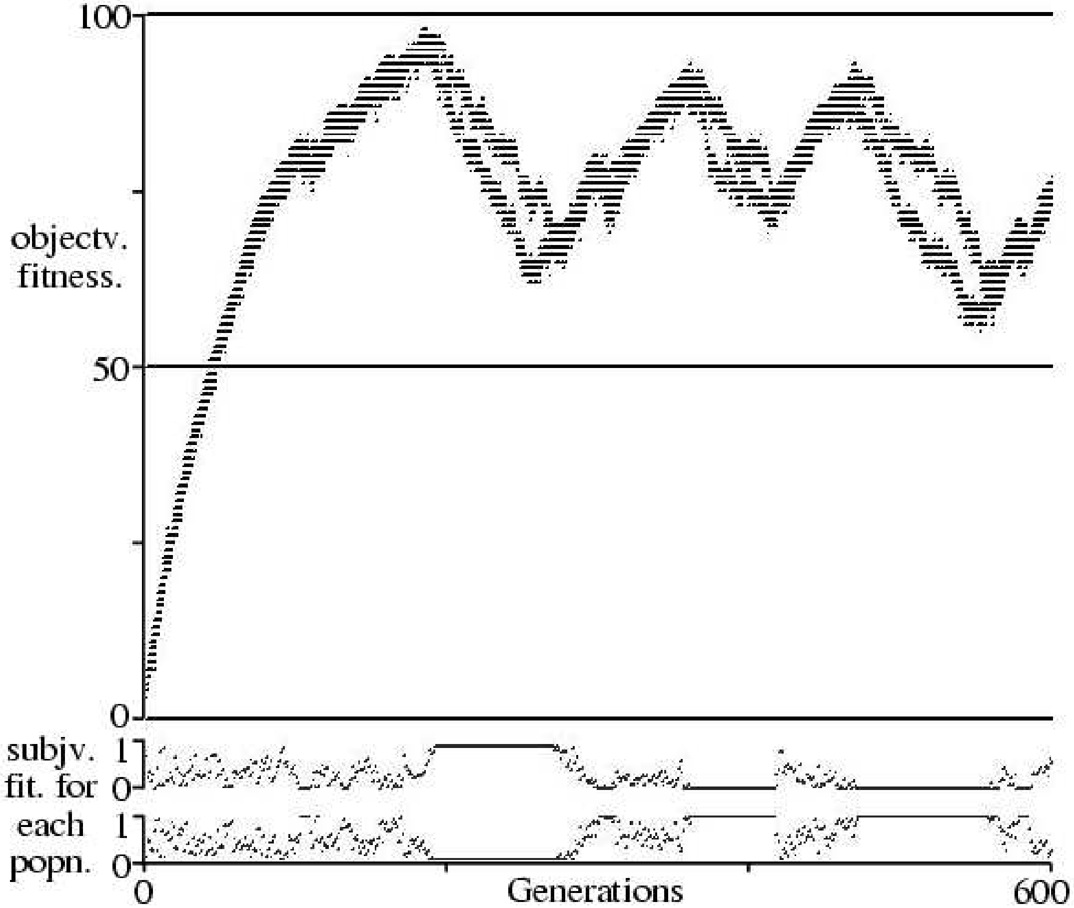

Over-specialization

- Change the game: now the objective fitness function is to be high on ten integer dimensions.
- The outcome of a contest is decided by comparing two individuals on the dimension on which they are most different.
- This leads to specialization on a small number of dimensions at a time, and no true “generalists” evolve.

Over-specialization

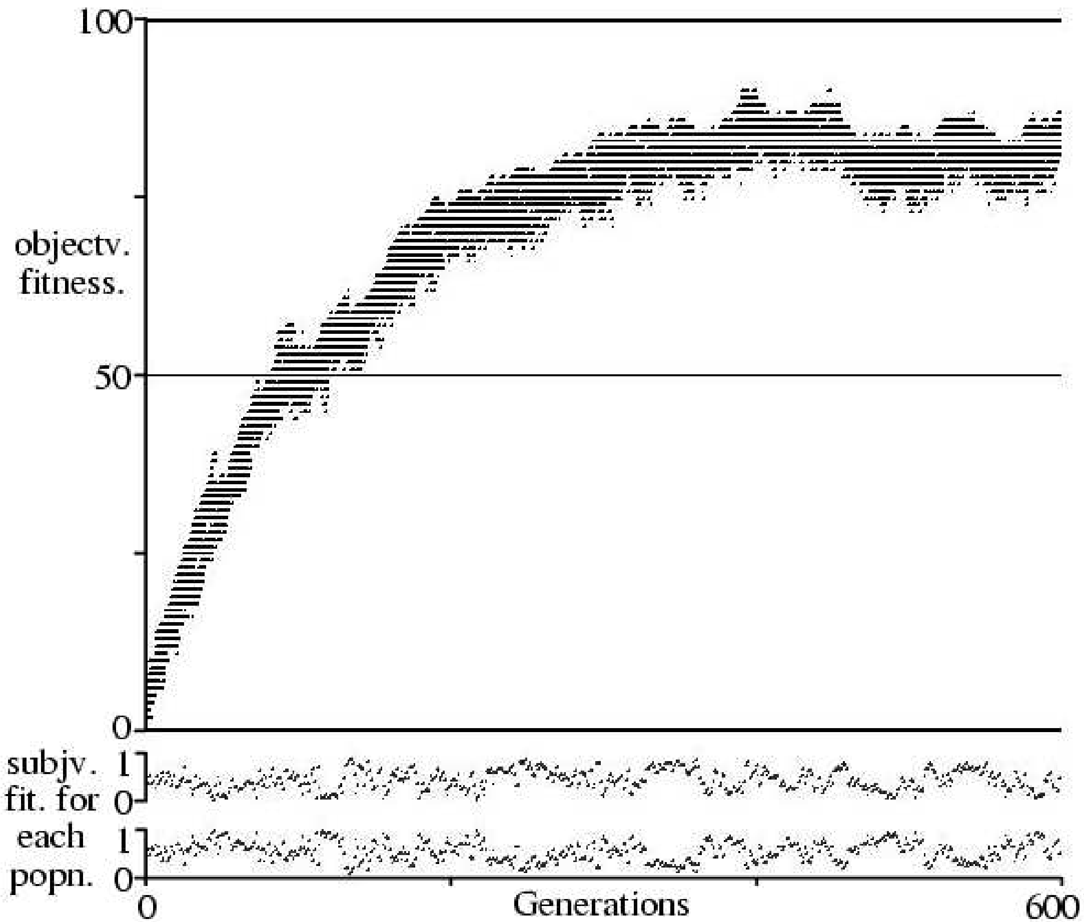

Relativism

- If we use multiple dimensions but decide contests based on the dimension on which two individuals are most *similar*, we can have intransitive dominance relationships.
- Success is now relative to your opponent, even though the objectively best genotype is to score 10 on all 10 dimensions.
- This setup is not good for co-evolutionary progress on objective fitness.

Intransitive dominance?

- The game of paper-rock-scissors provides a good example.

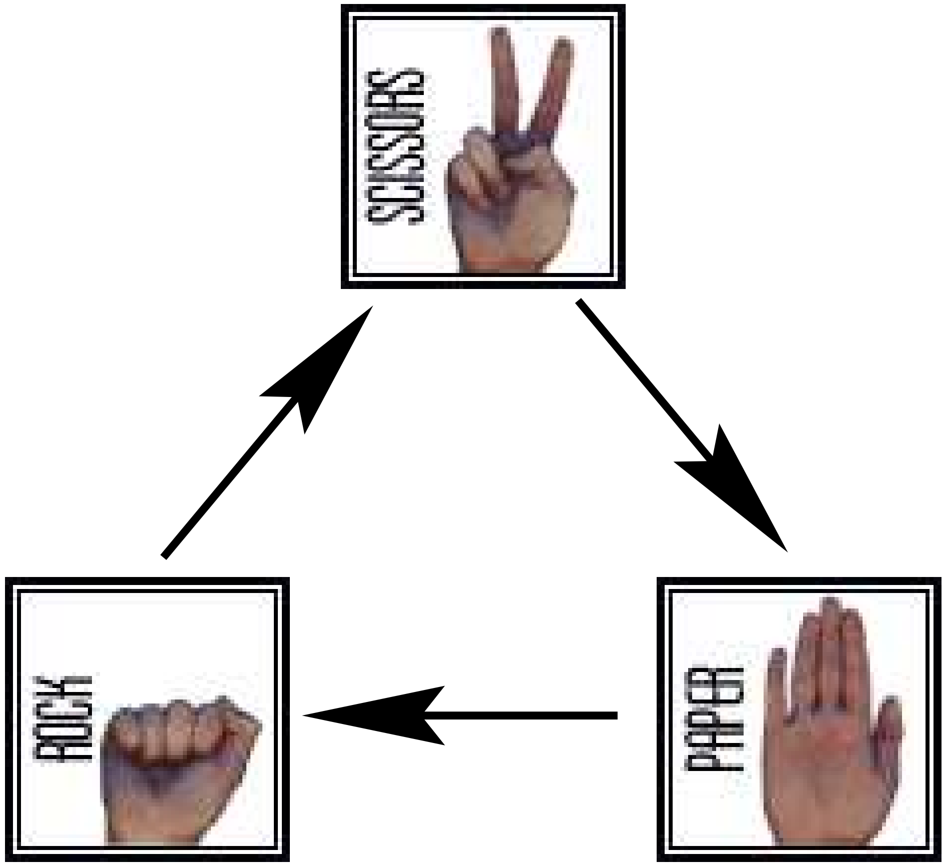

Intransitive dominance?

- Intransitive dominance refers to the fact that P beats R, and R beats S, but S beats P.
- If co-evolving populations played paper-rock-scissors against each other, the frequency of the three strategies would cycle.

Relativism

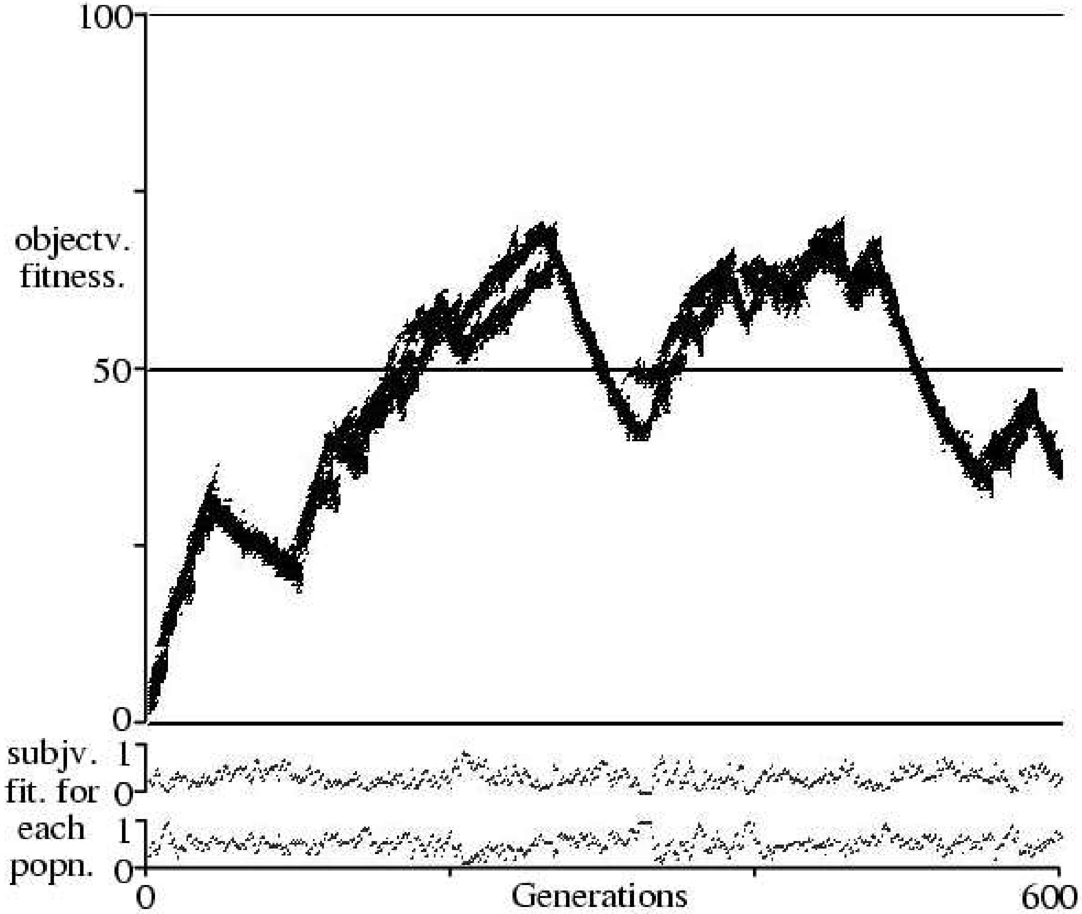
100
50
0
objectv.
fitness.
subjv. 1
fit. for 0
each 1
popn. 0
0
Generations
600

Relativism

- From the graph we can see that relativism can even push objective fitness down — this is not genetic drift.
- Leads to the perverse situation where subjective and objective fitness are opposed.

Parasite virulence

- Modifying parasite virulence is a way of addressing disengagement.
- In nature, parasites often experience selection pressure against being too hard on their hosts: if you kill the host, there's nothing left to parasitize.
- What if parasite fitness was based not on being as difficult as possible for hosts, but on being as discriminating as possible?
- Note parallel to teacher-student relationship.

Virulence and sorting networks

- Cartlidge and Bullock (2002) has re-addressed Hillis's sorting network problem.
- Implemented a non-linear transformation of fitness for parasites:

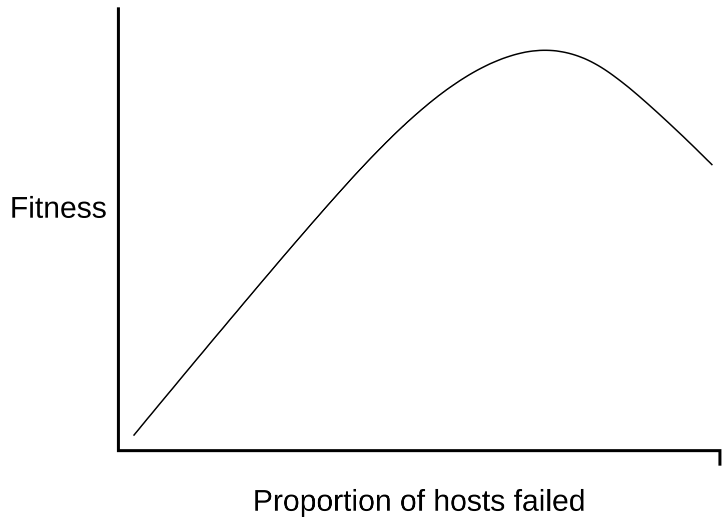

- Simulations show that reduced virulence helps with disengagement.

True and diffuse co-evolution

- Diffuse co-evolution is a way of addressing over-specialization.
- Bullock (1995) draws attention to the distinction made by Janzen (1980) between *true* and *diffuse* co-evolution, both within the context of asymmetric, inter-specific competition.
- True co-evolution is typical of parasites and hosts, where the mutual adaptation is on a single trait or a complementary pair of traits (e.g., egg mimicry and egg discrimination systems in brood parasitism).

True and diffuse co-evolution

- Diffuse co-evolution is more typical of predators, where co-evolution has affected a whole group of traits (e.g., the hard shells of crustaceans as a response to a variety of different predators with distinct shell-breaching tactics).
- Bullock suggests that the second variety of co-evolution is a more appropriate model for evolutionary engineers. True co-evolution leads to fragile or brittle solutions whereas diffuse co-evolutionary solutions are more robust.

Multi-objective optimization

- Taking multiple fitness objectives seriously may be a way to address the relativism problem.
- Is it practical to treat performance against different opponents as dimensions of success, and then apply Pareto optimization?

Pareto optimality

- Pareto optimal solutions are those for which no dimension of success can be improved without reducing performance on one of the other dimensions.
- Consider a proposed car design that needs to be both fast and economical . . .

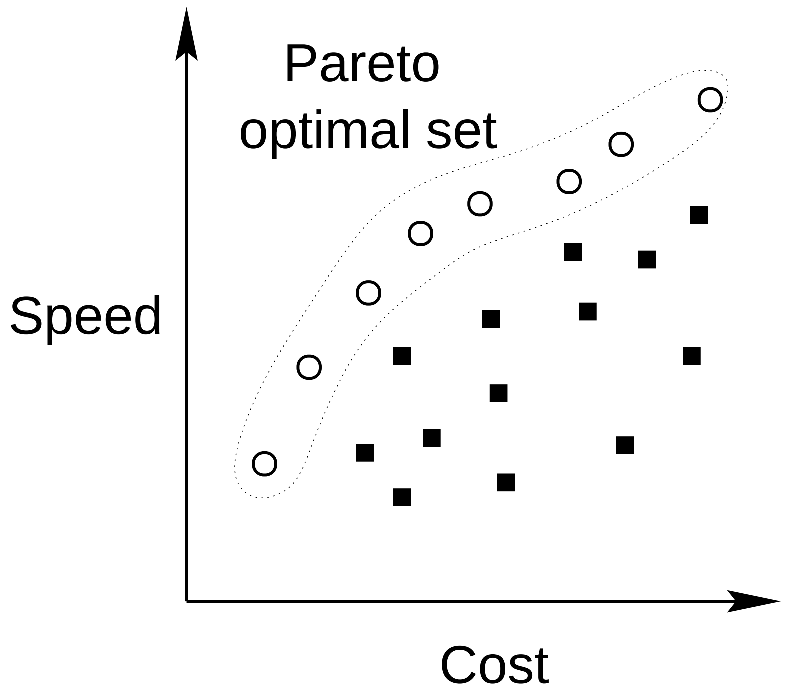

Pareto selection

- In theory, if we knew the performance of each strategy against every other, we could compute the pareto optimal set (POS).

	A	B	C	D
A	0	2	3	-1
B	-2	0	2	-1
C	-3	-2	0	0
D	1	1	0	0

Pareto selection

- In practice, this is computationally infeasible. But in a GA, we can find the POS for each generation, and use membership as the selection criterion.
- This gives us a noisy, partial window onto the true payoff matrix.

Coevolving poker strategies

- In evolving Texas Hold'em poker strategies, Noble and Watson (2001) have shown that a Pareto GA outperforms a standard co-evolutionary GA.
- Poker interesting because it's clear that the success of a strategy depends heavily on the opponents present, e.g., extreme bluffing. A lot of room for relativism.
- Pareto selection in co-evolutionary games seems to be a promising idea, and can be seen as a way of using the information from each game more effectively.

References

Bullock, S. (1995). Co-evolutionary design: Implications for evolutionary robotics. Cognitive science research paper 384, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.

Cartlidge, J., & Bullock, S. (2002). Learning lessons from the common cold: How reducing parasite virulence improves coevolutionary optimization. In Fogel, D. (Ed.), *Proceedings of the Congress on Evolutionary Computation*, pp. 1420–1425. IEEE Press.

Janzen, D. H. (1980). When is it co-evolution?. *Evolution*, *34*(3), 611–612.

Noble, J., & Watson, R. A. (2001). Pareto coevolution: Using performance against coevolved opponents in a game as dimensions for Pareto selection. In Spector, L., Goodman, E., Wu, A., Langdon, W. B., Voigt, H.-M., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M., & Burke, E. (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*, pp. 493–500. Morgan Kauffman, San Francisco.

Watson, R. A., & Pollack, J. B. (2001). Coevolutionary dynamics in a minimal substrate. In Spector, L., Goodman, E., Wu, A., Langdon, W. B., Voigt, H.-M., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M., & Burke, E. (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*. Morgan Kauffman, San Francisco.