Co-evolutionary Design I

Jason Noble

jasonn@comp.leeds.ac.uk

Biosystems group, School of Computing

Last time

- Search spaces, fitness landscapes, and other metaphors…
- What are the preconditions for a GA to work well?
- Evolutionary robotics: using GAs to design robot brains and bodies

This time

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

Co-evolution in nature

- Evolutionary arms races (Dawkins & Krebs, 1979).

- Successive adaptations in one group put pressure on another group to catch up.
- Each group may become bigger, faster, more lethal, more intelligent, etc.

Varieties of arms races

- Some arms races are symmetric:
  - For example, the arms race for huge body size and tusk weaponry in male elephant seals.
  - Participants all “striving” for the same goal.

Varieties of arms races

- Other arms races are asymmetric:
  - For example, the co-evolution of speed in the cheetah and evasive agility in the gazelle.
  - Each side becomes more accomplished at one of two complementary behaviours.
  - Parasite-host relationships also fit this category.

Varieties of arms races

Arms races can be inter- or *intra*-specific:

- Inter-specific arms races are usually asymmetric, because symmetric competition between species will usually dissolve into divergent evolution: it will pay one or the other species to shift their ecological niche a little.
- Intra-specific arms races can be either symmetric, as in male-male competition, or asymmetric, as in parent-offspring conflict (e.g., in birds: chicks evolve strident begging strategies, parents respond with more conservative feeding strategies, etc.).

Other forms of co-evolution

- Competitive arms races are not the only form that co-evolution can take.
- The co-evolution of cooperation is also possible through mechanisms such as:
  - mutualism, in which the cooperative strategy is in the interest of both organisms (e.g., symbiotic relationships).
  - kin selection, whereby enhancing the fitness of close relatives is a way of maximizing one's own *inclusive fitness* (see Hamilton, 1964).
  - reciprocal altruism, in which organisms take it in turns to be nice to each other.

Artificial co-evolution?

As usual, there are two reasons to be interested:

- On the engineering side, there is a perception that co-evolution might make evolution go faster, and/or produce more robust solutions.
- On the biological modelling side, incorporating co-evolution will produce richer and more plausible models of animals in an ecological context.

Co-evolution and landscapes

- The fitness of a genotype will vary over time according to the genotypes present in the co-evolving population.
- The fitness landscape is no longer static but changes over time.
- For example, a highly fit cheetah of six million years ago might not be able to catch any modern gazelles.

Building a co-evolutionary GA

- Two GAs running in tandem.
- Fitness of an individual in population A is determined by how well it does against one or more members of population B, and vice versa.
- In other respects the GAs are standard.

Co-evolved sorting networks

- Hillis (1990) used co-evolution as an engineering tool to develop sorting networks.

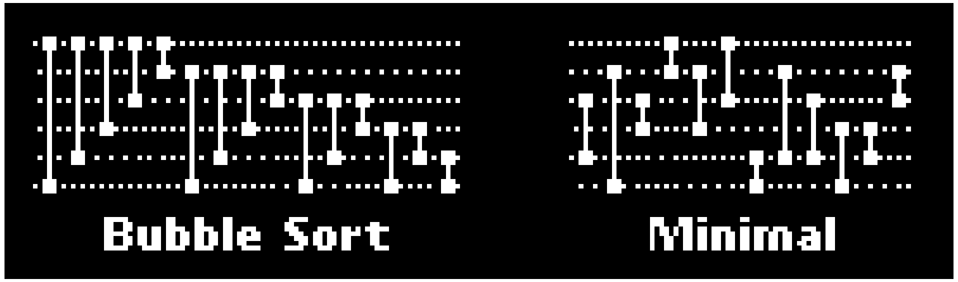

- Sorting networks describe algorithms for sorting a fixed-length list of numbers.
- Much effort has gone into designing sorting networks that are not only correct but *minimal* (15 vs. 12 comparisons, above).

Starting with a normal GA

- Hillis originally used a standard GA to evolve minimal correct networks for sorting lists of 16 numbers.
- The fitness function was the percentage of cases sorted correctly, given a series of random lists to sort. There was therefore only an implicit pressure towards fewer comparisons.
- The smallest network previously discovered (by hand) used 60 comparisons. The best network discovered by the GA used 65 comparisons.

Moving to co-evolution

- Hillis tried co-evolution, hoping to start an arms race between the sorting networks and their “parasite” opponents.
- A parasite was implemented as a collection of 16 lists of numbers to be sorted.
- The fitness function for the parasite population was the average percentage of cases that the networks failed to sort.

Moving to co-evolution

- Co-evolution thus pushed the parasites towards hard-to-sort lists, and pushed the networks towards robust sorting algorithms.
- The smallest correct sorting network to evolve used only 61 comparisons—just short of equalling the best human-designed network.

Co-evolved pursuit and evasion

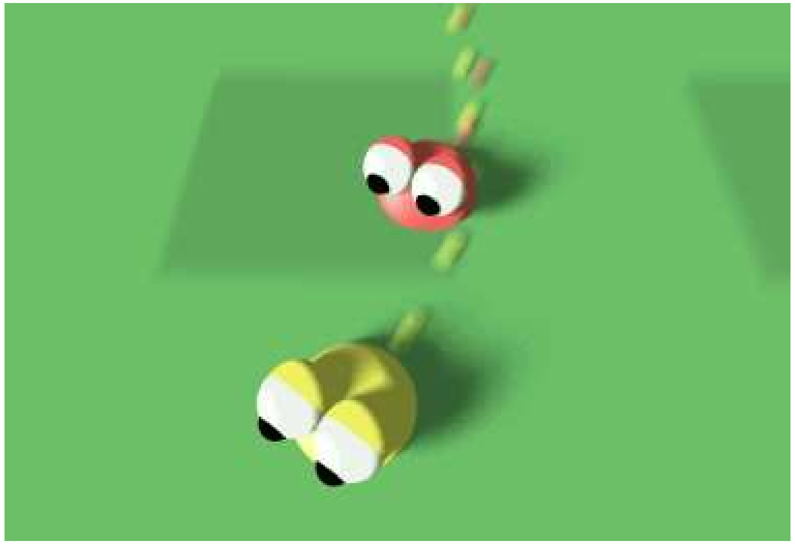

- Miller and Cliff (1994) modelled pursuit and evasion behaviour.

- An inherently co-evolutionary scenario.
- A similar system is modelled as a BEAST demo.

Co-evolved pursuit and evasion

- Miller and Cliff argued that an evolutionary simulation model was necessary because *game-theoretic* models of pursuit and evasion were either over-simplified or mathematically intractable.
- They set up a detailed physical simulation in which one robot would chase another across an infinite plane—both robots had equal top speeds and maneuverability. Friction, turning radius, and energy consumption were all included in the model.

Co-evolved pursuit and evasion

- The simulated robots were controlled using continuous time recurrent neural networks, encoded using a fractal-tree model of neuronal growth (previous lecture).
- The fitness function for the evading robot was proportional to the average distance it maintained between itself and the pursuer. The pursuer's fitness was the inverse of this.
- Plausible-looking pursuit and evasion behaviour certainly evolved. Evaders dodged and weaved, while pursuers followed the evaders and sometimes appeared to anticipate their next move.

Measuring progress

- Cliff and Miller (1995) talk about “tracking the Red Queen”, referring to the Red Queen hypothesis of van Valen (1973).

- Species in co-evolved ecological contexts have to keep adapting just to maintain the status quo.
- In a standard GA, we can assess progress by plotting fitness over time. This is not possible in a co-evolutionary GA.

Comparing new and old

- Cliff and Miller found that if modern pursuers were assessed against ancient evaders, and vice versa, progress was evident.

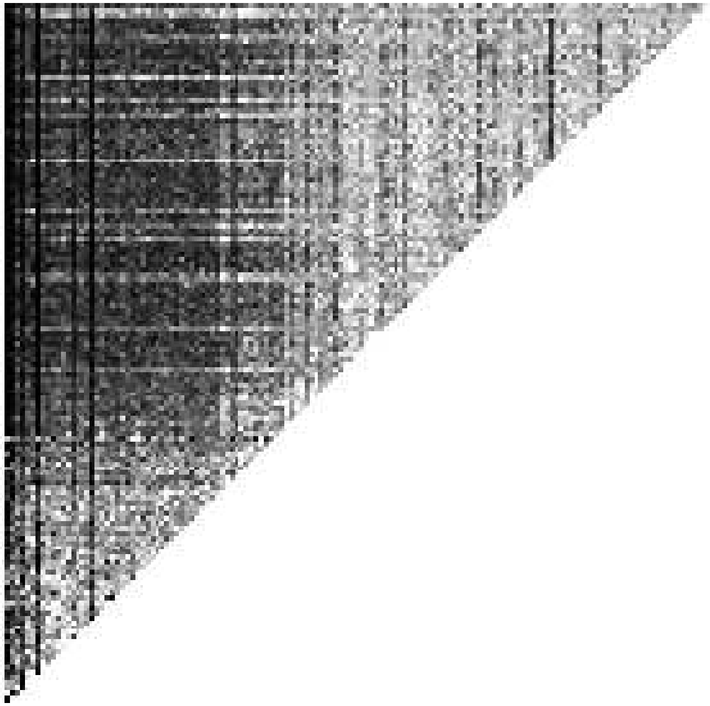

- It is not trivial to interpret this information, however.

Two populations needed?

- The important thing about co-evolution is the fact that fitness varies depending on which other genotypes are around (cf. game theory, Maynard Smith, 1982).
- This means that single-population co-evolution is not a contradiction.
- Doesn’t that make everything co-evolution? No: climbing a fixed landscape is still just evolution.

Next 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

References

Cliff, D., & Miller, G. F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Morán, F., Moreno, A., Merelo, J. J., & Chacón, P. (Eds.), *Advances in Artificial Life: Third European Conference on Artificial Life (ECAL '95)*, Vol. 929 of *Lecture Notes in Artificial Intelligence*, pp. 200–218. Springer, Berlin.

Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London: Biological Sciences*, *205*, 489–511.

Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, *7*, 1–52.

Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, *42*, 228–234.

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

Miller, G. F., & Cliff, D. (1994). Protean behaviour in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pp. 411–420. MIT Press / Bradford Books, Cambridge, MA.

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