Learning to Select Mates in Evolving Non-playable Characters

Dylan R. Ashley*, Valliappa Chockalingam*, Braedy Kuzma*, and Vadim Bulitko

2019-08-23

University of Alberta

Background

What is PCG?

Procedural Content Generation tries to procedurally generate assets for video games. With the increasing complexity of video games this is becoming increasingly important and valuable.

https://kotaku.com/a-look-at-how-no-mans-skys-procedural-generation-works-1787928446

Figure 1: No Man's Sky made heavy use of PCG.

What is A-life?

Artificial life studies **lifelike systems** in computer simulations.

Frequently A-life systems are multi-agent systems where the survival of individual agents is dependent on some fitness function implicitly defined by the environment.

https://cottagelife.com/outdoors/brains-of-mice-and-bats-could-help-researchers-understand-human-interactions/

Figure 2: How likely are these mice to be eaten by an owl tonight?

How do PCG and A-life relate?

Many video games try to emulate **lifelike systems**. One common example would be populations of NPCs interacting with one another. We can imagine generating these populations procedurally.

Figure 3: Vilmonic is a game at the intersection of A-life game and PCG.

What is mate selection?

Many A-life simulations involve evolving populations of agents. When evolving agents in artificial life, some strategy has to be employed to decide how agents will mate to produce the next generation.

https://imgur.com/QCvP6xt

Figure 4: What strategy do these sheep use to pick their mates?

Why is mate selection important?

We often seek to have populations of fit agents but if a fit agent selects a bad mate then the offspring may end up being feeble.

However, it is often unclear what makes two agents compatible.

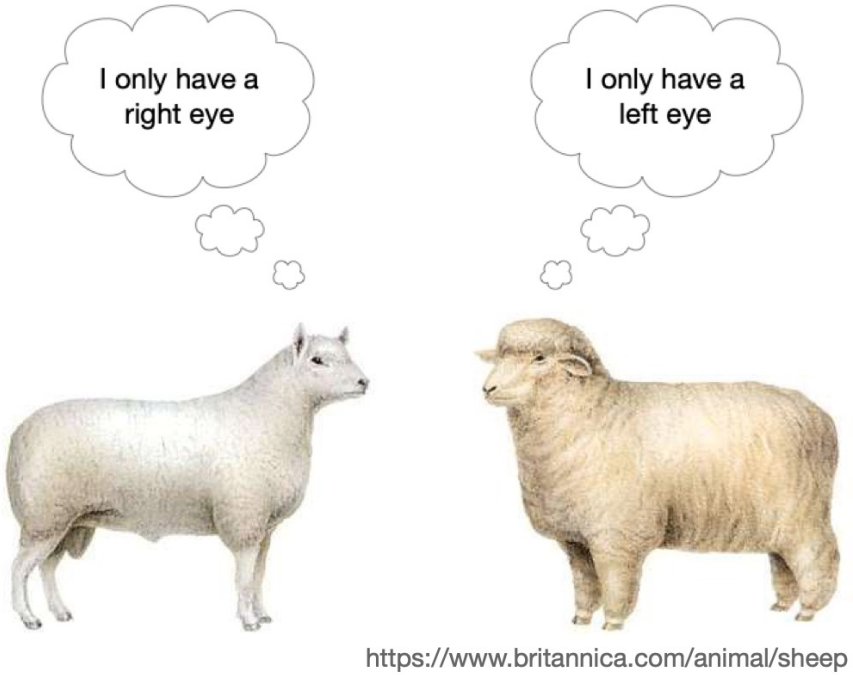

Figure 5: Will the offspring have zero, one, or two eyes?

Proposed Approach

How do we tackle the mate selection problem?

We investigate how evolving mate-selection strategies can increase the efficiency of evolutionary search in an A-life setting.

https://blogs.plos.org/blog/2016/10/04/darwins-demons-better-video-games-through-natural-selection/

Figure 6: In Darwin's Demons, the efficiency of the evolutionary search was paramount.

Methodology

What is the domain we use to analyze this?

We use a wolf-sheep predation model. Grass grows and is eaten by sheep; sheep are eaten by wolves, starve to death, or die from natural causes:

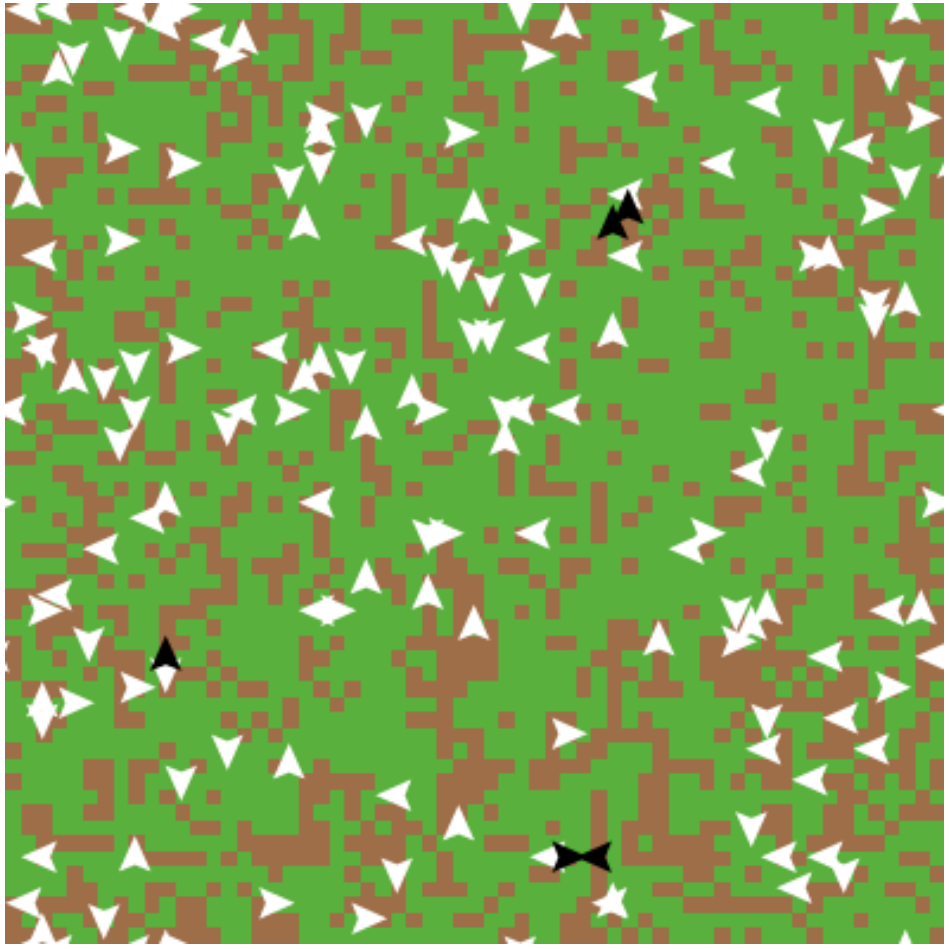

What are the agents of interest here?

We focus on the sheep. We treat them as agents in an *Evolutionary Reinforcement Learning* setting and use the expected lifetime of a population of sheep as a proxy for gauging the efficiency of the evolutionary search.

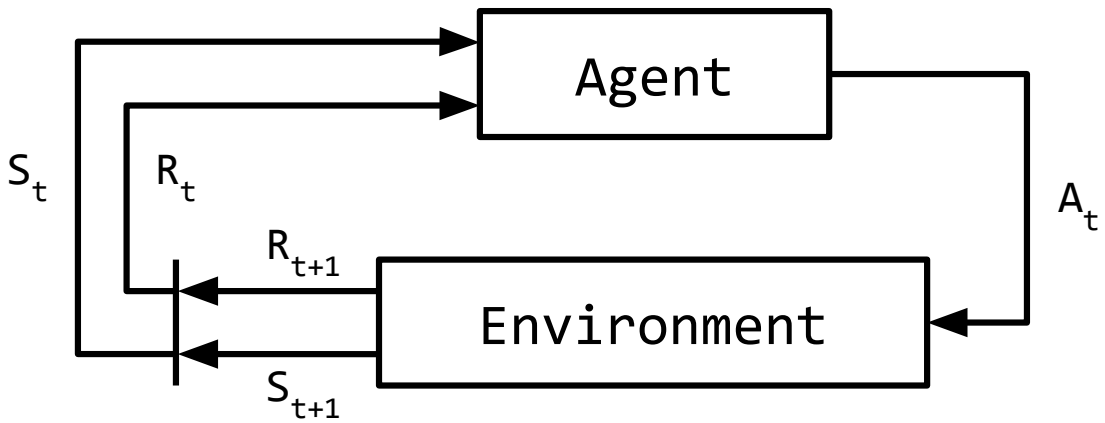

Figure 7: Evolutionary Reinforcement Learning is Reinforcement Learning with an evolved reward function.

How do we construct the state for the sheep?

Because the world is small and in the spirit of A-life, we provide the sheep with the following state features:

- how much energy it has,
- the distance and heading to the nearest patch of grass,
- the distance and heading to the nearest wolf,
- the distance and heading to the nearest sheep,
- the maximum of the last known age of any of its children, and
- their preference for the nearest sheep.

What actions do the sheep have?

At each step a sheep can either

- move forwards,
- turn left or right,
- eat the patch of grass it's standing on, or
- mate with the closest sheep.

To decide what to do, each sheep has

- a learned and evolved action network,
- an evolved evaluation network,
- an evolved preference network.

How is a sheep's preference for another sheep calculated?

The preference of one sheep for another is a linear combination with evolved coefficients which takes the genome of the potential mate as input:

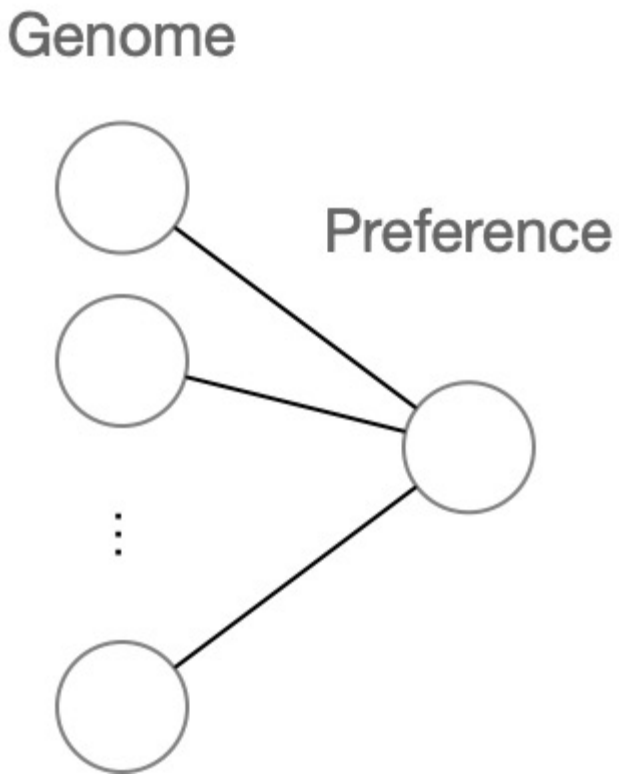

Do we really care only about the other agent's genome?

As an agent selecting a mate you may care about how similar the potential mate is to you. So instead of just feeding the preference network the genome of the potential mate we also consider feeding it

- the elementwise absolute difference of the genomes of the selecting agent and the potential mate,
- the elementwise squared difference of the genomes of the selecting agent and the potential mate, or
- the euclidean distance between the genomes of the selecting agent and the potential mate.

Results

Does evolving a preference network help?

When using a preference network we see a significant improvement in the population survival time compared to random mating:

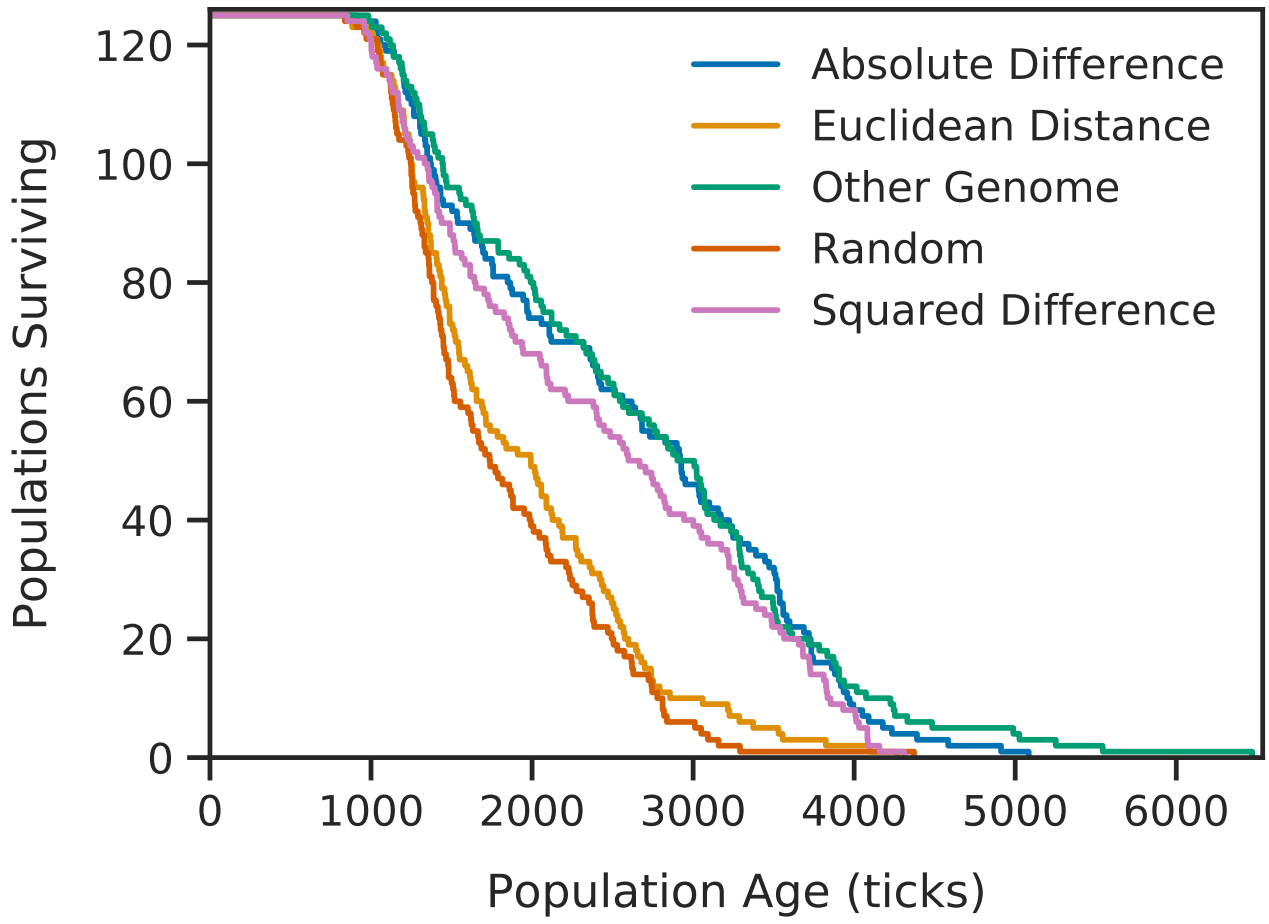

So what does that all mean?

We note that

- absolute difference, squared difference, and just using the other agent's genome did the best;
- just using the other agent's genome did as well as absolute difference or squared difference; and
- euclidean distance did no better than random.

We can see this more clearly when we look at the distribution:

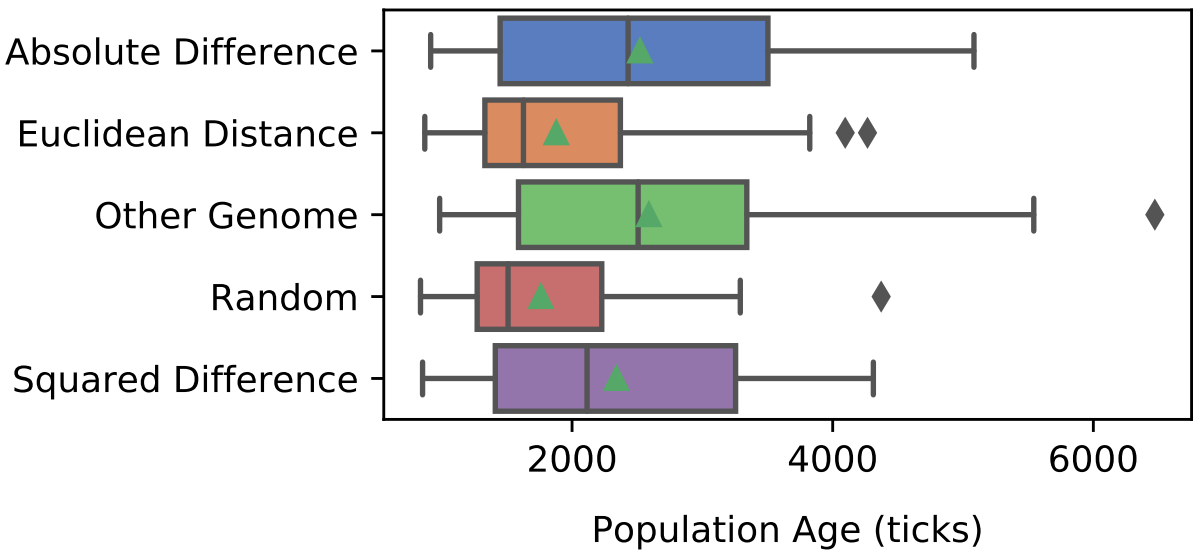

But what are they actually learning?

When looking at some of the trends in the populations we see they tend to evolve to favour survival traits:

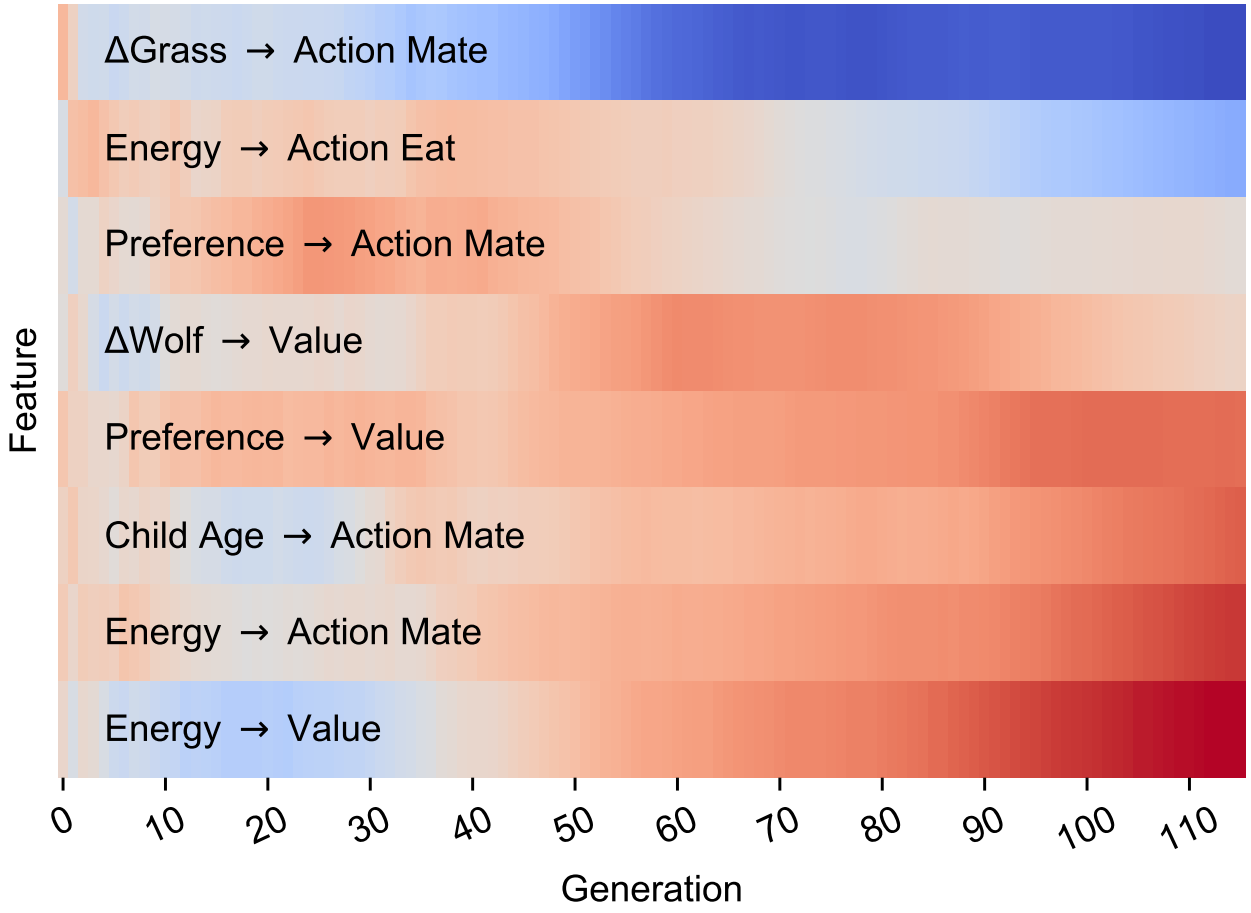

Conclusion

So what are the major takeaways?

In this work we conclude that

1. evolving mate-selection strategies can increase the expected survival time of a population of agents;
2. when evolving mate-selection strategies, agents tend to learn to favour mates with survival traits; and
3. in the context of evolving mate-selection strategies, how agents compare their genomes to other agents' genomes matters.

What's next?

In an artificial life situation it makes sense that agents should both control **how they are presented** to others and **how they interpret the presentation of others**.

This could be seen as a form of emergent communication.

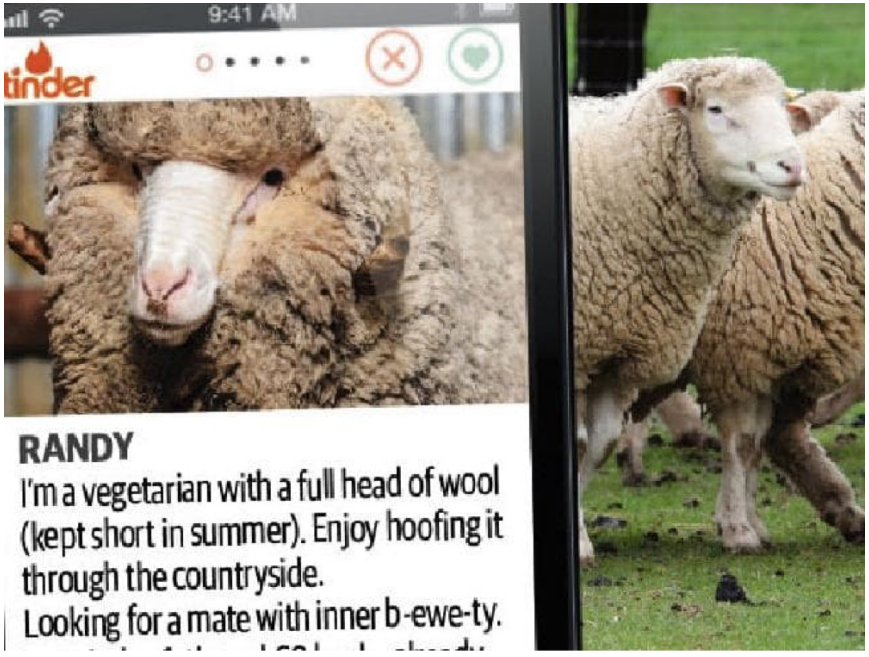

Figure 8: Would you date this sheep?

What else is next?

We would also like to

- see how including information about how potential mates calculate preferences when calculating preferences affects mate selection, and
- implement these techniques in a video game with the player's actions informing the evolution.

Questions?