Learning to Select Mates in Artificial Life

Dylan R. Ashley
University of Alberta
Edmonton, Alberta, Canada
dashley@ualberta.ca

Valliappa Chockalingam
University of Alberta
Edmonton, Alberta, Canada
valliapp@ualberta.ca

Braedy Kuzma
University of Alberta
Edmonton, Alberta, Canada
braedy@ualberta.ca

Vadim Bulitko
University of Alberta
Edmonton, Alberta, Canada
bulitko@ualberta.ca

ABSTRACT

Artificial life (A-life) simulations present a natural way to study interesting phenomena emerging in a population of evolving agents. In this paper, we investigate whether allowing A-life agents to select mates can extend the lifetime of a population. In our approach, each agent evaluates potential mates via a preference function. The role of this function is to map information about an agent and its candidate mate to a scalar preference for deciding whether or not to form an offspring. We encode the parameters of the preference function genetically within each agent, thus allowing such preferences to be agent-specific as well as evolving over time. We evaluate this approach in a simple predator-prey A-life environment and demonstrate that the ability to evolve a per-agent mate-selection preference function indeed significantly increases the extinction time of the population. Additionally, an inspection of the evolved preference function parameters shows that agents evolve to favor mates who have survival traits.

CCS CONCEPTS

• **Computing methodologies** → **Multi-agent reinforcement learning**; **Artificial life**; **Genetic algorithms**; *Multi-agent systems*; *Sequential decision making*;

KEYWORDS

artificial life, mate selection, reinforcement learning, genetic algorithms, multi-agent systems

ACM Reference Format:
Dylan R. Ashley, Valliappa Chockalingam, Braedy Kuzma, and Vadim Bulitko. 2019. Learning to Select Mates in Artificial Life. In *Genetic and Evolutionary Computation Conference Companion (GECCO '19 Companion), July 13–17, 2019, Prague, Czech Republic.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3319619.3322060

1 INTRODUCTION

Our world is becoming saturated with smart, interconnected, Artificial Intelligence (AI)-driven devices. With the emergence of vast networks of such agents comes the potential of interesting emergent behavior among them. This potential raises important questions. What will this behavior look like? Will only simple behavior develop or could we see something complex emerging? How will this impact our society? Artificial life (A-life) simulations are one way to study such phenomena. A-life allows the study of both single- and multi-agent learning at multiple timescales: within an agent's lifetime (e.g., via reinforcement learning) as well as across generations of agents via genetic search [1]. The potential of emergent collective behavior is a powerful provision of A-life since even simple agents acting according to simple rules can produce a population that exhibits complex behavior (e.g., Conway's Game of Life [2] and Wolfram's cellular automaton rule 110 [8]). Another advantage of A-life simulations is that they remove the need for hand-designed fitness function which has applications in procedural content generation for video games [9].

While A-life eliminates the basic genetic search's need for an explicitly designed fitness function (in A-life the fitness function is implicitly induced by the "physics" of the environment), it still requires a mechanism for mate selection for sexual agent reproduction. Mate selection is important for ensuring the diversity of the gene pool as well as protecting innovation as it matures. For instance, the well known NEAT algorithm uses speciation as a part of its mating strategy [5]. One can port speciation and other hand-designed mate-selection strategies to A-life. In the spirit of evolutionary computation, we propose a different approach: we genetically encode a preference function in each agent thus allowing evolution to find good mate-selection strategies. We implement and evaluate our approach in an A-life setting for the reasons listed earlier. We note that a similar approach can be applied in a synchronous evolution scenario with discrete generations such as a scenario where NEAT would be applied.

2 RELATED WORK

Mate selection has been previously studied before in the context of evolutionary methods and A-life. Huang empirically evaluated several methods of performing mate selection in conventional genetic algorithms [4]. However, none of the evaluated methods involves learning to perform mate selection. Guntly and Tauritz proposed a

---

GECCO '19 Companion, July 13–17, 2019, Prague, Czech Republic

ACM ISBN 978-1-4503-6748-6/19/07…$15.00
https://doi.org/10.1145/3319619.3322060

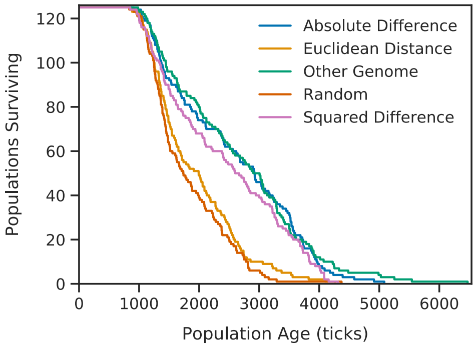

Figure 1: Populations survival curves.

way of directly learning how to perform mate selection [3]. However, their method requires agents to be able to determine the fitness of the offspring of a pairing directly. In Werner and Todd's work on artificial birds, male agents communicate with females through songs and females choose mates based on the songs they prefer [6]. However, there are no additional survival pressures on the agents.

3 PROPOSED APPROACH

We propose extending the Evolutionary Reinforcement Learning paradigm of Ackley and Littman [1] to provide agents with, in addition to an action and an evaluation network, a preference network that learns a linear mapping between some combination of the agents genomes ($\mathbf{G_a} \in \mathbb{R}^n$) with a candidate mate's genomes ($\mathbf{G_c} \in \mathbb{R}^n$) to a scalar value. The weights of this network can be learned evolutionarily. We use four different combinations of genomes as input to the preference network:

Other genome	$\mathbf{G_c} \in \mathbb{R}^n$
Absolute difference	$\lvert\mathbf{G_a} - \mathbf{G_c}\rvert \in \mathbb{R}^n$
Squared difference	$(\mathbf{G_a} - \mathbf{G_c})^2 \in \mathbb{R}^n$
Euclidean distance	$\lVert\mathbf{G_a} - \mathbf{G_c}\rVert \in \mathbb{R}$

4 RESULTS

Figures 1 and Table 1 show the results of an experiment on a wolf sheep predation model implemented in NetLogo [7] in which the four different combination functions, as well as the random mating baseline, were each given 125 unique populations and were allowed to proceed until extinction. Here there is a significant difference between the average survival time of each of absolute difference, other genome, and squared difference and each of Euclidean distance and random ($p < 0.01$ using two-tailed t-test with Bonferroni corrections). The lack of a meaningful gap between the average survival time of each of absolute difference, other genome, and squared difference suggests that some form of co-evolution may be occurring between the whole genome and the portion of the genome that deals with the preference network weights. More investigation is required to determine if this is, in fact, the case.

One reason why Euclidean distance performed so poorly is that the preference network under this way of combining genomes is restricted to learning a single weight, so the preference of an agent for another agent is likely to be dominated by irrelevant features of that agent's genome.

Table 1: Population survival times.

Preference function	Mean population survival time
Other genome	2589.5 ± 100.0
Absolute difference	2521.0 ± 94.3
Squared difference	2340.4 ± 90.0
Euclidean distance	1878.9 ± 64.5
Random	1761.0 ± 57.2

Noting that the other-genome combination function performed well, we conducted a post hoc analysis of what traits agents in the 30 longest living populations evolved to seek or avoid in potential mates. We noted that within these populations sheep evolved to seek other sheep that had survival traits. For example, they evolved to seek mates who assigned a higher value to states where they had more energy. They also evolved to avoid mates who preferred to mate when distant to grass.

5 CONCLUSIONS

Evolutionary methods are commonly used in A-life. However, the performance of these methods can depend on the mate-selection strategy employed by agents. In this work we allowed agents to evolve a mapping from a combination of their genomes and a potential mate's genomes to a scalar preference value that they could then use to decide whether or not to select that mate. We showed that this leads to a significantly higher population extinction time when compared to randomly selecting mates. We additionally noted that agents evolved to prefer mates with survival traits.

We believe that altogether this work represents a principled step towards avoiding needing to use hand-coded mating strategies in A-life domains.

Future work will allow agents to select how they present themselves to others; they will be able to design their own profiles rather than other agents being able to observe their genome directly. These profiles could be co-evolved alongside the preference networks.

ACKNOWLEDGMENTS

The authors would like to acknowledge that this work was made possible through resources provided by Calcul Québec and Compute Canada.

REFERENCES

[1] David Ackley and Michael Littman. 1991. Interactions between learning and evolution. *Artificial life II* 10 (1991), 487–509.
[2] Martin Gardner. 1970. Mathematical games: The fantastic combinations of John Conway's new solitaire game "life". *Scientific American* 223, 4 (1970), 120–123.
[3] Lisa M. Guntly and Daniel R. Tauritz. 2011. Learning individual mating preferences. In *Proceedings of the Genetic and Evolutionary Computation Conference.* ACM, 1069–1076.
[4] Chien-Feng Huang. 2002. *A study of mate selection in genetic algorithms.* Ph.D. Dissertation. University of Michigan.
[5] Kenneth O. Stanley and Risto Miikkulainen. 2002. Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation* 10, 2 (2002), 99–127.
[6] Gregory M. Werner and Peter M. Todd. 1997. Too many love songs: Sexual selection and the evolution of communication. In *Proceedings of the European Conference on Artificial Life.* MIT Press/Bradford Books, 434–443.
[7] Uri Wilensky. 1999. NetLogo. (1999). http://ccl.northwestern.edu/netlogo/
[8] Stephen Wolfram. 2002. *A new kind of science.* Wolfram media.
[9] Georgios N. Yannakakis and Julian Togelius. 2018. *Artificial Intelligence and Games.* Springer.