Artificial Intelligence Accidentally Learned Ecology through Video Games

Lou Barbe, Cendrine Mony, Benjamin W Abbott

Artificial intelligence accidentally learned ecology through video games

Authors: Lou **Barbe**[1], Cendrine **Mony**[1], Benjamin W. **Abbott**[2]

Affiliations:

[1]ECOBIO, OSUR, CNRS, Université de Rennes 1, 35000 Rennes, France

[2]Brigham Young University, Department of Plant and Wildlife Sciences, Provo, USA

Corresponding author: Lou Barbe (+33 6 67 44 06 56)

Emails:

lou.barbe@univ-rennes1.fr

cendrine.mony@univ-rennes1.fr

benabbott@byu.edu

Type of article: Forum

Abstract word count: 50

Main text word count: 1319

Number of figures and boxes: 2

Number of references: 12

Abstract

An advanced artificial intelligence system defeated the best human players in StarCraft II, a popular real-time strategy game. In a virtual ecosystem, players compete for habitats and resources, unintentionally reproducing many ecological phenomena. We propose to repurpose this A.I. to test ecological hypotheses that have been intractable using traditional approaches.

Video games: a new playground for science and artificial intelligence

Transforming data into predictive understanding is one of the fundamental challenges facing ecology. Artificial intelligence (A.I.) and machine learning techniques have been widely used to interpret complex data streams, but their application in ecology remains very limited compared to other research disciplines and industries (e.g. economics, genetics, and engineering [1]). One domain where artificial intelligence plays a dominant role is the video game industry, which has provided strong support for the development of a new generation of A.I. systems, using artificial neural networks and cutting-edge techniques of machine learning. Here, we report on a recent breakthrough, concerning a popular real-time strategy game (StarCraft II, [2]). In this game, which unintentionally reproduces central ecological dynamics, an advanced A.I. called AlphaStar trained itself to play the game using multi-agent reinforcement learning. Using unexpected tactics, it beat soundly some of the world's best human players in matchups in late 2019. We propose to retool this system: feeding AlphaStar with real socioecological scenarios, it could address persistent and novel issues in the real world, including predicting emergent properties and non-linearities in succession, identifying risk and factors for ecological collapse, and understanding ecological strategies under changing environments. This would build on work by programmers and researchers who are already collaborating to create realistic or "serious" games [3] that advance research via gamification [4,5] and facilitate education [6].

A complex and unpredictable virtual environment reproducing ecological dynamics

StarCraft II is one of the most popular video game franchises and largest games in e-sports. Two or more players confront each other on a game map with the goal of eliminating the opponent, similar to the traditional board games of Chess and Go [7,8]. However, StarCraft II integrates several dimensions that simulate real-world complexity much more closely than traditional games. Real-time matches occur in a high-resolution, heterogeneous landscape with an uneven distribution of resources and terrain, leading to around 10^{26} possible actions every step of the game [2]. Within this virtual ecosystem, players collect resources to create up to ~20 different game pieces, called units. In turn, these units collect resources, provide defense, or attack competitors. Additionally, there is an artificial fog (called the "fog of war"), that constrains the visible portion of the landscape to the area immediately around the player’s units. Like in reality, StarCraft II is a world where individuals struggle for survival with incomplete information about resource location and threats of competition or predation (see Kriegspiel [9] for previous applications of A.I. in imperfect-information games).

In addition to the generally realistic characteristics inherent to most real-time strategy games, StarCraft II unintentionally integrates a large range of ecological dynamics and mechanisms. First of all, players must choose one of the three "races" (Terran, Protoss, and Zerg), which have different life strategies that align strongly with the three main ecological plant strategies laid out by Grime: competitors, stress tolerators, and ruderals (Figure 1). Each player continuously makes trade-offs between colonization of new habitats versus rapid conflict with the opponent, depending on the stoichiometry of available resources and level of external threat. Early in the game, when resource acquisition is very limited and few possible units are available, players often produce cheap and numerous units (*r* strategy). As the match progresses, players generally produce more expensive but powerful units (K strategy) that have specialized and complementary roles, increasing the biodiversity in the virtual ecosystem. Furthermore, interactions among players result in shifts towards units with functionally different traits, while the degree of performance and specialization of the traits themselves can be enhanced, intensifying the arms race throughout the course of the game. Overall, the growth is highly modular and plastic, and many choices are viable only when combined with other decisions and considerations (strong phenotypic integration). Growth and cognition operate within a network of units, similar to natural organisms like clonal plants or superorganisms (social animals, e.g. ants or termites). Last, each player has a maximum number of units simultaneously allowed on the map, enforcing an artificial carrying capacity and associated density dependent strategies.

The unexpected ecological mastery of AlphaStar

On October 30th, 2019, Vinyals *et al.* [2] revealed that their A.I. system, named AlphaStar, promptly defeated many professional players and outperformed around 99.8% of officially ranked human players. Before the confrontations, AlphaStar agents (simulated players) had been trained with data from millions of StarCraft II matches played against each other — the equivalent of 200 years of continuous gaming per agent. During this spin-up process, the algorithm selected and crossed the most successful agents, exerting selective pressure on the variation generated by the machine learning. In their isolation from human influence, many agents developed very aggressive strategies, entirely free of human gaming culture such as tacit or unconscious alliances between players to ensure long-term participation. For example, in confrontations in which both players have the same race, agents used apparently counterproductive strategies, such as over-saturating their resource collection capacity or having a homogeneous army composition. This unconventional approach was so effective that these games almost always finished in a few minutes (compared to tens of minutes for a game among humans). The overall domination of AlphaStar over its human competitors in these situations demonstrated that these counterintuitive behaviors better balanced the ecological trade-offs necessary in a context of strong intraspecific competition (*i.e.* symmetric competition). In this configuration, better performance in resource collection and a strategy of quantity (competitive hierarchy) not quality (limiting similarity) were the most likely to allow domination.

Using AlphaStar to address ecological and evolutionary problems

Given the amount of time and resources invested in its development, and despite some limitations of the game (e.g. very little resource recycling, no biotic pressure other than that of the competitor), AlphaStar is arguably more sophisticated, flexible and robust than any existing ecological model. AlphaStar could be readily repurposed to address fundamental questions in ecology and evolution. Preliminarily, we could study the current games of AlphaStar as a competition model between two individuals (Box 1). We could then make agents directly learn from maps corresponding to realistic configurations of biodiversity and resource distribution. The frequency and the trajectory among games of agents' various ecological strategies could reveal their evolutionary trajectory, which would allow assessment of deterministic and contingent development pathways in different starting conditions. We could also manipulate environmental conditions during games and observe the response of agents with given ecological strategies (*i.e.* ecological response), for instance in terms of trait syndromes or chances of success.

Some could understandably ask, what could we learn from this exercise that is not simply an artefact of the game parameters? First of all, we think that AlphaStar should be modified by limiting the number of the actions it can make per minute (*i.e.* "micromanaging"), to allow testing substantive, ecological processes rather than just differences in processing power. Then, from a plant perspective, we could study the effects of ecological strategies and their response to changing environments, including predicting nonlinearities in succession, niche shifts and trade-off modifications caused by direct disturbance, resource change or biological invasions. From an animal perspective, we could study the role of personality in ecological functioning [10], the ecology of fear (*i.e.* the impact of the *risk* of competition [11]) and the evolution of ecological culture ([12], e.g. emergence and drift of ecological approaches). From a human perspective, AlphaStar could shed light on cultural relations: if many humans from different cultures play against AlphaStar, it would generate a rich dataset relevant to how individual, cultural, and artificial personality can impede or solve complex and crucial ecological problems such as ecological collapse, environmental change, and tolerance of immediate and eventual risk. The data, code, and architecture of AlphaStar's neural network and StarCraft II itself are freely available. Relationships among gamers, ecologists, and programmers would allow co-construction and testing of many more hypotheses than we list above, possibly generating new insights into many stubborn challenges of ecology and evolutionary biology.

Acknowledgments

We thank Blizzard for creating StarCraft and we thank the DeepMind's team for its work and achievements on A.I. training in StarCraft II. We thank Simon Chollet, Sylvain Glémin, Jason Hills, Vincent Jung, Mathias Lherbier, Loïs Morel, Julia Niemeyer, Eric Petit, Pedro Poli, Andreas Prinzing, Léa Uroy, Phillipe Vandenkoornhuyse and Nathan Vannier, for fruitful early discussion and suggestions on the manuscript.

Figures

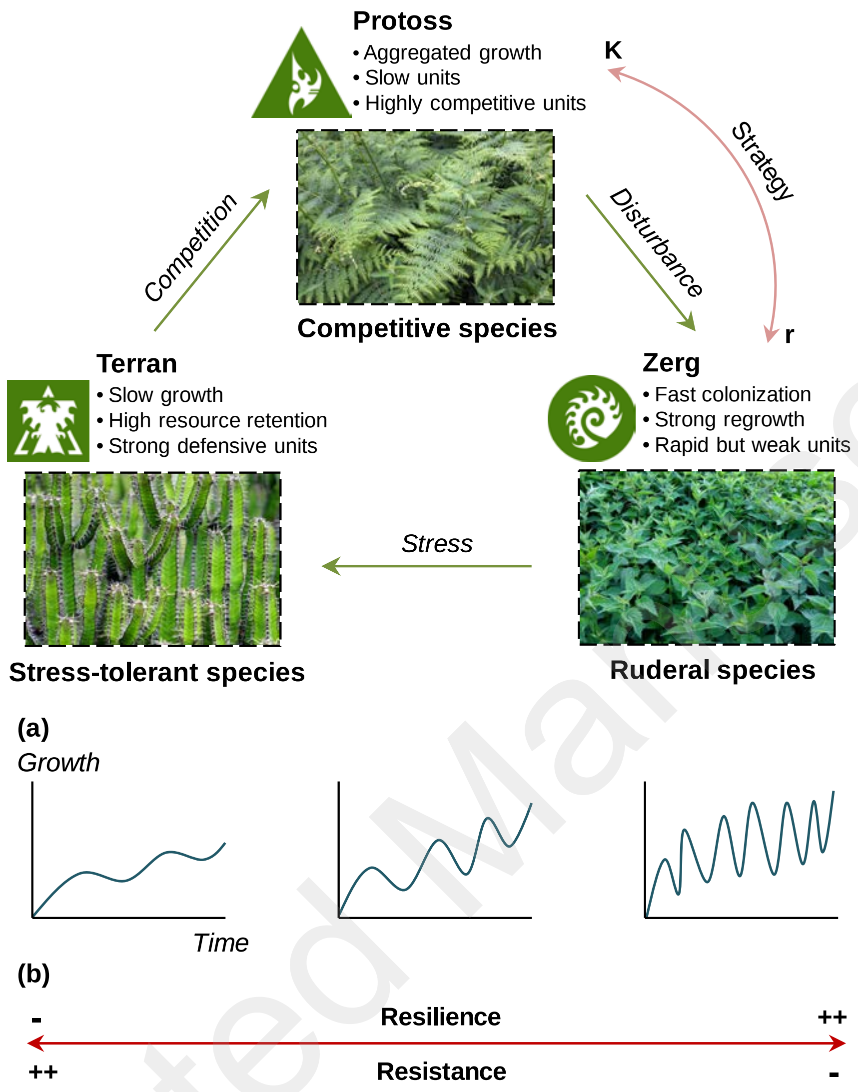

Figure 1. Ecology of the three races of StarCraft II through the prism of Grime's CSR triangle. At the top, a competitive species (e.g. bracken, ferns) represents Protoss, which produce slow and expensive units that need to be locally aggregated but are very competitive. At the left, a stress-tolerant species (e.g. cactus) reflects the strategy of Terrans, which can resist intense stresses by collecting and conserving resources more efficiently under opponent pressure, and use several, strong defensive units and strategies. At the right, a ruderal species (common nettle, *Urtica dioica*) represents the strategy of Zergs, which colonize space rapidly, regrow rapidly and have low cost, aggressive units with generally short lifespans. (a) These ecological strategies mean that each race expresses distinctive growth curve. (b) Each race also has emergent properties such as more or less resilience and resistance.

Box 1. The basical ecological model in StarCraft II

The classic confrontation type in StarCraft II, as played by AlphaStar, consists of two players confronting each other in a particular environment with limited resources − see the "mini-map" (Figure I). This represents an ecological competition between two individuals, which can be intra or interspecific, depending on whether the players choose the same race. The map shows the entire environment, but the players' vision is limited to the lighter, circular zones around their respective units and buildings. Resources are the light blue shapes, and the dark blue and red shapes are the buildings and units of both players (Protoss in blue, Terran in red). The progress and outcome of competition can be monitored through this mini-map, which shows colonization and exploitation of new resource patches, environment exploration, and often, ecological collapse. This map could be used to monitor more realistic ecological models. For instance, several players could compete against each other in a larger and entirely unpredictable environment, which would enable the study of ecological processes at the scale of populations and communities. We could also set up scenarios in which we artificially modify environmental conditions during the game, and then evaluate the consequences on ecosystem functioning and biodiversity dynamics. Note that quantification of several ecological processes like trait changes or trade-off modification are easily feasible with the game parameters.

Figure I. Mini-map of a standard game of StarCraft II.

References

1. Olden, J.D. *et al.* (2008) Machine learning methods without tears: a primer for ecologists. *Q. Rev. Biol.* **83**, 171-193.

2. Vinyals, O. *et al.* (2019) Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575**, 1-5.

3. Chang, A.Y. (2013) Playing nature: The virtual ecology of game environments. Doctoral dissertation, UC Berkeley.

4. Sørensen, J.J.W. *et al.* (2016) Exploring the quantum speed limit with computer games. *Nature* **532**, 210-213.

5. Fenwick, T. (2017) Real exoplanets, virtual universe. *New Sci.* **3122**, 12.

6. Connolly, T.M. *et al.* (2012) A systematic literature review of empirical evidence on computer games and serious games. *Comput. Educ.* **59**, 661-686.

7. Silver, D. *et al.* (2016) Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484-489.

8. Silver, D. *et al.* (2018) A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* **362**, 1140-1144.

9. Ciancarini, P. and Favini, G.P. (2010) Monte Carlo tree search in Kriegspiel. *Artif. Intell.* **174**, 670-684.

10. Start, D. and Gilbert, B. (2017) Predator personality structures prey communities and trophic cascades. *Ecol. Lett.* **20**, 366-374.

11. Clinchy, M. *et al.* (2013) Predator-induced stress and the ecology of fear. *Funct. Ecol.* **27**, 56-65.

12. Mace, R. and Holden, C.J. (2005) A phylogenetic approach to cultural evolution. *Trends Ecol. Evol.* **20**, 116-121.