

# Modeling Prehistoric Economic Behavior With AI

Claes Strannegård<sup>1,2</sup>

<sup>1</sup>Dept. of Applied Information Technology, University of Gothenburg, Sweden

<sup>2</sup>Dept. of Molecular Medicine and Surgery, Karolinska Institutet, Sweden

claes.strannegard@gu.se

## Introduction

Agent-based models (ABMs) are becoming increasingly popular in a wide range of scientific fields (Macal and North, 2010). In economics, they have been used to represent human behavior (Hamill and Gilbert, 2015), e.g., in finance (Axtell and Farmer, 2025), game theory (Koster et al., 2025), and human-ecosystem models (An, 2012). Moreover, in archaeology, they have been applied to modeling prehistoric human populations in North America (Janssen, 2010) and Crete (Chliaoutakis and Chalkiadakis, 2016). Likewise, in ecology, they are used for modeling individual animals living in ecosystems and interacting with other organisms and the physical environment (DeAngelis and Mooij, 2005), thus shedding light on population dynamics (Geary et al., 2020) and the ecological consequences of economic activities (Strannegård et al., 2024).

Hand-coded models of behavior are widely used in psychology, economics, ecology, human-ecosystem models, and artificial life, not least in the context of ABMs (Macal and North, 2010). However, there are limitations associated with hand-coding behavioral models, particularly when the goal is to model animals or humans. In many cases, we simply do not know enough about animals or humans to be able to hand-code their behavior (Tohyama et al., 2025). In other cases, we may know in principle but struggle to translate our knowledge into code (DeAngelis and Diaz, 2019).

While hand-coded models remain common, machine learning, particularly in the form of deep reinforcement learning (DRL) (Barto, 2021), has proven to be superior in many domains. For example, DRL models excel at Atari games (Mnih et al., 2015), Go (Silver et al., 2017), and Minecraft (Hafner et al., 2025). Other applications of DRL include algorithmic trading (Sun et al., 2023) and modeling animal behavior in a predator-prey context (Sunehag et al., 2019; Yamada et al., 2020).

According to modern biology, *Homo sapiens* is one animal species among others (Darwin, 2008). From that perspective, it is natural to model prehistoric humans as animals living in ecosystems, while taking into account their ability to use tools and cooperate. To address the challenges

of hand-coding and take advantage of the power of DRL, our earlier work introduced an agent-based framework for ecosystem modeling (Strannegård et al., 2025). The contribution of this work is an agent-based framework for ecosystem modeling that leverages DRL and incorporates human agents capable of tool use and trade, enabling the study of emergent economic behavior in prehistoric environments.

## Framework for ecosystem modeling

Our framework for ecosystem modeling relies on four key ideas. First, geographical areas are represented by grids of cells, where each cell contains explicitly modeled populations of organisms of different species and has physical properties such as temperature, altitude, and land-cover class. Second, the behavior of each animal species is controlled by a neural network (policy network), which determines what each individual of that species will do at each time step, based on its observations of the cells surrounding it and its own internal variables (Fig. 1). Third, DRL is used




Figure 1: Policy network of a sheep agent. The agent observes its own internal variables, e.g., energy, protein, and hydration (left). It also observes the properties of the cells in the  $3 \times 3$  cell neighborhood surrounding it, informing it about the presence of grass, water, obstacles, other sheep, and wolves (middle). Based on those observations, the agent decides which action to take (right).

for constructing the policy networks, with a reward function that encourages survival. To survive, the agents must in turn learn to find food, avoid predators, conserve energy, and navigate efficiently through the terrain. Finally, the agents are trained across a range of different ecosystem models to

enable them to survive in diverse environments.

### Adding human agents

To add human agents to the above framework, we need to specify: (1) their internal variables, e.g., energy, protein, and water; (2) the external objects that they can observe, e.g., all objects in the  $3 \times 3$  cell neighborhood surrounding them, plus aggregated smells in the same cells; (3) what they can ingest, e.g., fruits, nuts, and water; (4) how they can move, e.g., North, East, South, West; (5) how moving and ingesting affect their internal variables; (6) how they reproduce; and (7) which events will cause their death. Thus, we can model




Figure 2: View of an ecosystem model with grass (green), water (blue), sheep (white), wolves (red), and humans (yellow). The amounts of fruit and nuts in each cell are not shown in this view.

“naked” humans interacting with each other and with their environment (Fig. 2).

### Adding production and trade

Next, we will extend the human models by adding extra actions for production and trade. First, we introduce *resources* of three kinds: *Foods*, *Commodities*, and *Tools*. For the sake of concreteness, we give numerous examples of resources that could be used in the extended models:

**Foods:** Apples, Nuts, Berries, Mushrooms, Roots, Eggs, Cooked potato, Cooked fish, Cooked meat, and Bread

**Commodities:** Grain, Meal, Raw fish, Raw meat, Water, Firewood, Sand, Stone, Iron, Gold, and Diamond

**Tools:** Spear for hunting, Hook for fishing, Net for fishing, Boat for fishing, Axe for chopping firewood, and Fence for breeding animals or growing crops.

Each human agent has a set of possessions called *assets*, consisting of certain amounts of resources. For human agents, the *observation space* may be defined similarly to the one in Fig. 1, with the addition that they can observe their own assets. Their *action space* may contain the following actions:

- $Move(D)$ , where  $D$  is North, East, South, or West
- $Consume(F)$ , where  $F$  is a food item
- $Produce_i(R)$ , where  $R$  is a resource that is being produced and  $i$  is an index representing a method
- $Trade$

The action  $Consume(F)$  means removing a fixed quantity of food  $F$  from the agent’s assets and updating its internal variables according to the nutritional value of  $F$ .

$Produce_i(R)$  might represent actions such as hunting, fishing, or crafting an axe. These actions might add resources to the agent’s assets and some actions require assets such as a spear, a net, or a piece of iron. More advanced production actions may require collaboration with several agents performing the same action.

$Trade$  is the action of exchanging some of the agent’s own assets with assets owned by another agent. For each pair of resources  $R$  and  $R'$  and each human agent  $A$ , the conditions under which  $A$  will trade its own asset of type  $R$  for  $R'$  are specified by two additional outputs of its neural network:  $Quantity(R, R')$  and  $Price(R, R')$ . Here  $Quantity(R, R')$  represents the amount of its own assets of  $R$  that  $A$  is willing to trade for  $R'$ , while  $Price(R, R')$  represents the minimum amount of  $R'$  that  $A$  wants to receive for each unit of  $R$ . Now, suppose two human agents  $A$  and  $B$  (i) perform the action  $Trade$  simultaneously, (ii) have matching trading conditions, and (iii) are in the same cell. Then, one matching trade is randomly selected and executed between  $A$  and  $B$ .

The agents of the extended model can now be trained with reinforcement learning across a range of environments of varying difficulty, using a reward signal that encourages survival. In this way, behavioral models of all agents, including humans, can be constructed in a uniform manner.

### Emergent economic behavior

With these models, one can run simulations to study a variety of emergent economic behaviors. For example, one may show that survival in certain environments is impossible without fishing, hunting, saving, or trading. Our next step is to implement the framework described and conduct initial simulations.

### Acknowledgements

This research was supported by the Sten A Olsson Foundation for Research and Culture and Stiftelsen Sävstaholm. ChatGPT was used for language improvement.

## References

- An, L. (2012). Modeling human decisions in coupled human and natural systems: Review of agent-based models. *Ecological Modelling*, 229:25–36.
- Axtell, R. L. and Farmer, J. D. (2025). Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature*, 63(1):197–287.
- Barto, A. G. (2021). Reinforcement learning: An introduction. *SIAM Rev*, 6(2):423.
- Chliaoutakis, A. and Chalkiadakis, G. (2016). Agent-based modeling of ancient societies and their organization structure. *Autonomous agents and multi-agent systems*, 30(6):1072–1116.
- Darwin, C. (2008). *The descent of man, and selection in relation to sex*. Princeton University Press.
- DeAngelis, D. L. and Diaz, S. G. (2019). Decision-making in agent-based modeling: A current review and future prospectus. *Frontiers in Ecology and Evolution*, 6:237.
- DeAngelis, D. L. and Mooij, W. M. (2005). Individual-based modeling of ecological and evolutionary processes. *Annual Review of Ecology, Evolution, and Systematics*, 36:147–168.
- Geary, W. L., Bode, M., Doherty, T. S., Fulton, E. A., Nimmo, D. G., Tulloch, A. I., Tulloch, V. J., and Ritchie, E. G. (2020). A guide to ecosystem models and their environmental applications. *Nature Ecology & Evolution*, 4(11):1459–1471.
- Hafner, D., Pasukonis, J., Ba, J., and Lillicrap, T. (2025). Mastering diverse control tasks through world models. *Nature*, pages 1–7.
- Hamill, L. and Gilbert, N. (2015). *Agent-based modelling in economics*. John Wiley & Sons.
- Janssen, M. A. (2010). Population aggregation in ancient arid environments. *Ecology and Society*, 15(2).
- Koster, R., Píslar, M., Tacchetti, A., Balaguer, J., Liu, L., Elie, R., Hauser, O. P., Tuyls, K., Botvinick, M., and Summerfield, C. (2025). Deep reinforcement learning can promote sustainable human behaviour in a common-pool resource problem. *Nature Communications*, 16(1):2824.
- Macal, C. M. and North, M. J. (2010). Tutorial on agent-based modelling and simulation. *Journal of Simulation*, 4(3):151–162.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., and Hassabis, D. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359.
- Strannegård, C., Engsner, N., Lindgren, R., Olsson, S., and Endler, J. (2024). AI tool for exploring how economic activities impact local ecosystems. In Arai, K., editor, *Intelligent Systems and Applications*, pages 690–709, Cham. Springer Nature Switzerland.
- Strannegård, C., Palak, M., Engsner, N., Stocco, A., Antonelli, A., and Silvestro, D. (2025). Predicting ecosystem resilience using multi-agent reinforcement learning. *bioRxiv*.
- Sun, S., Wang, R., and An, B. (2023). Reinforcement learning for quantitative trading. *ACM Transactions on Intelligent Systems and Technology*, 14(3):1–29.
- Sunehag, P., Lever, G., Liu, S., Merel, J., Heess, N., Leibo, J. Z., Hughes, E., Eccles, T., and Graepel, T. (2019). Reinforcement learning agents acquire flocking and symbiotic behaviour in simulated ecosystems. In *Artificial life conference proceedings*, pages 103–110. MIT Press.
- Tohyama, T., Nagashima, U., Higashino, I., Arima, F., Hiyoshi, K., Nagase, M., Yada, Y., Honda, N., and Watabe, A. M. (2025). Aversive experiences induce valence plasticity of instructive signals to change future learning rules in mice. *Communications Biology*, 8:1002.
- Yamada, J., Shawe-Taylor, J., and Fountas, Z. (2020). Evolution of a complex predator-prey ecosystem on large-scale multi-agent deep reinforcement learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.