

Nest-Building Using Place Cells for Spatial Navigation in an Artificial Neural Network

Thomas E. Portegys

Abstract—An animal behavior problem is presented in the form of a nest-building task that involves two cooperating virtual birds, a male and female. The female builds a nest into which she lays an egg. The male's job is to forage in a forest for food for both himself and the female. In addition, the male must fetch stones from a nearby desert for the female to use as nesting material. The task is completed when the nest is built, and an egg is laid in it. A goal-seeking neural network and a recurrent neural network were trained and tested with little success. The goal-seeking network was then enhanced with “place cells”, allowing the birds to spatially navigate the world, building the nest while keeping themselves fed. Place cells are neurons in the hippocampus that map space.

Keywords—Artificial animal intelligence, artificial life, goal-seeking neural network, nest-building, place cells, spatial navigation.

I. INTRODUCTION

A task is presented to two virtual birds, a male and a female, that cooperate in navigation, foraging, communication, and nest-building activities. These activities are commonly performed in many animal species to ensure survival and reproduction. The task is a division of labor for the birds: the female builds a nest into which she lays an egg, completing the task; the male forages in a forest for food for both himself and the female, and obtains stones from a desert for the female to use as nesting material.

The nest-building task was recently proposed as a game providing an artificial animal intelligence challenge to evaluate machine learning systems [1]. While the task itself is independent of how it is tackled, here artificial neural networks (ANNs) are chosen, as ANNs are capable of generalized learning, and are intended to perform functions of the biological neural networks of animals.

The task was originally introduced in 2001 [2], a solution for which was obtained using Mona, a goal-seeking ANN. Mona is designed to control autonomous artificial organisms by learning cause-and-effect relationships in the environment along with responses that enact these causations to navigate toward goals. However, the solution at that time employed domain-specific macro-responses, such as “Go to mate”.

In planning to attack the problem again, this time with spatial enhancements that would obviate the need for domain-specific responses, place cells seemed to be a good choice. Place cells are neurons located within the hippocampus which are triggered in response to an animal entering specific places in its local environment, aiding spatial navigation [3]-[5]. In

this project an abstraction of place cells is incorporated into Mona to allow it to navigate a grid world to accomplish foraging and nest-building. It is the author's belief that training an ANN to build a nest is a novel contribution to machine learning.

There is a significant body of work on using place cell inspired neural functionality in ANNs, much of it involved with robotic navigation [6], [7]. These systems are aimed at solving specific tasks with models that mimic biological place cells. They are not intended to be general-purpose ANNs, such as Mona, which are designed to learn arbitrary domain-independent tasks. General-purpose ANNs borrow functionality from brains, such as neural connection updating, but are not intended to be models of biological brains.

Historically, AI has mostly focused on human-like intelligence, for which there are now numerous success stories: games, self-driving cars, stock market forecasting, medical diagnostics, language translation, image recognition, etc. The impact of ChatGPT [8] as a generative language model is a recent example. Yet the elusive goal of artificial general intelligence (AGI) seems as far off as ever, likely because these success stories lack the “general” property of AGI, operating as they do within narrow, albeit deep, domains. A language translation application, for example, does just that and nothing else.

Zador [9] expresses this succinctly: “We cannot build a machine capable of building a nest, or stalking prey, or loading a dishwasher. In many ways, AI is far from achieving the intelligence of a dog or a mouse, or even of a spider, and it does not appear that merely scaling up current approaches will achieve these goals.” It is evident that nature has more to teach us toward the development of AGI. In general, AIs are currently designed to predict the future, whereas animals have evolved to cause the future in the pursuit of goals.

We can go back to Braitenberg's vehicles [10] as a possible starting point for the study of how simple networks are capable of controlling goal-directed behavior in automata. Dyer [11] delineated the contrasting aims of ANNs for animal behavior vs. traditional artificial intelligence (AI) by framing the former as having biological goals: “Survivability is the overriding task”. As an example of a simulated animal brain, Coleman et al. [12] developed an ANN-based cognitive-emotional forager that performed well in a task that required not only foraging, but also the avoidance of predators.

The foraging problem is a common task for swarm robotics, since a swarm of cooperating agents can cover an area efficiently. Ericksen et al. [13] developed a successful swarm seed forager using NEAT (Neuroevolution of Augmented

Thomas E. Portegys is with Dialectek, DeKalb, Illinois USA (e-mail: portegys@dialectek.com).

Topologies) created by Stanley and Miikkulainen [14]. NEAT evolves a neural network topology that controls behavior which maximizes a fitness score, in their case defined by one point for picking up a seed and two points for returning it to the nest. In contrast, the Mona ANN in this project is trained from experience with the environment rather than through evolution. Mona additionally employs needs and goals as a motivation mechanism.

Another example of artificial creatures controlled by ANNs is Yaeger's Polyworld artificial life system [15]. Polyworld is an environment in which a population of agents search for food, mate, have offspring, and prey on each other in a two-dimensional world. An individual makes decisions based on its neural network which is derived from its genome, which in turn is subject to evolution.

In addition to simulating life forms, ANNs have been used as modeling and analysis tools for animal behavior [16], [17].

As a comparison, an LSTM (Long short-term memory) recurrent neural network (RNN) [18] was also trained on the task, without spatial enhancement.

II. DESCRIPTION

The code with instructions is available at [19].

World

The world is a 21x21 two-dimensional grid of cells. Each cell has a *locale*, and an *object* attribute. *Locale* describes the type of terrain: *plain*, *forest*, and *desert*. An *object* is an item to be found in the world: *mouse* (food), and *stone* (nest-building material). A forest exists in the upper left of the world, populated by mice, which randomly move about, providing an element of surprise for a foraging bird. A desert is found in the lower right of the world, where stones are to be found at various locations. The birds are initially located on the plain in the center of the world.

Birds

There is a male and a female bird. The female builds the nest and the male forages for mice and stones. The nest is a stone ring in the center of the world in which the female lays her egg. The birds have four components: senses, internal state, needs, and responses. These are sex-specific to suit the different roles of the birds.

Male

Senses: locale, mouse-proximity, stone-proximity, mate-proximity, goal, has-object, flying, female-needs-mouse, female-needs-stone.

Locale pertains to the current location of the male and has a value of *plain*, *forest*, or *desert*.

The *proximity* sensor values are *present*, *left*, *right*, *forward*, or *unknown*. The *mouse-proximity* sensor senses a mouse when in the forest, the *stone-proximity* sensor senses a stone when in the desert, and the *female-proximity* sensor senses the female within the bounds of the nest.

The *goal* sensor values are *eat-mouse*, *mouse-for-female*, *stone-for-female*, and *attend-female*.

The *has-object* sensor indicates an object carried by the bird and can be *mouse*, *stone*, or *no-object*.

The *flying sensor* is true when the male is in flight; otherwise false.

Female-needs-mouse is sensed when the female expresses a corresponding response of *want-mouse* in the presence of the male. This is the only time this sense is active; when not in the presence of the female it is in the off state. A similar process exists for the *female-needs-stone* sense and *want-stone* female response. Only one of the *female-needs/wants* is sensed/expressed at a time.

Internal state: *food*.

Food is initialized to a parameterized value. When food reaches zero, the need for a mouse is surfaced as a goal. Upon eating a mouse, food is increased by a random value.

Needs: mouse-need, female-mouse-need, female-stone-need, attend-female-need.

Needs correspond to goals: *mouse-need* to the goal *eat-mouse*, *female-mouse-need* to *mouse-for-female*, *female-stone-need* to *stone-for-female*, and *attend-female-need* to *attend-female*. Upon completion of an errand to satisfy a goal, signified by returning to the female, the next goal is determined by current needs. As discussed, *mouse-need* is raised when food reaches 0. The *female-mouse-need* and *female-stone-need* are signaled to the male by the female when she desires a mouse to eat or a stone to place in the nest, respectively. If none of the above are raised, the *attend-female-need* is raised, causing the male to move to the female's location.

Responses:

- *do-nothing*: a no-op response.
- *move-forward*: move forward in the orientation direction. Movement off the grid is a no-op.
- *turn-right/left*: change orientation by 90 degrees.
- *eat-mouse*: eat mouse if *has-object* is a mouse. If no mouse, this is a no-op.
- *get-object*: if *has-object* is empty and an object in current cell, pick up the object and set it to *has-object*.
- *put-object*: if *has-object* not empty and no object at current cell, put object in cell and clear *has-object*.
- *give-mouse*: if *has-object* is mouse, and female present with empty *has-object*, transfer mouse to female.
- *give-stone*: if *has-object* is stone, female present with empty *has-object*, transfer stone to female.
- *fly*: take flight. This activates a place motor neuron which will move to a specific location in the world. The specific motor is determined by the current mediator neuron context (see Artificial neural networks section).
- *alight*: terminate flight after arriving at a location in the world determined by the active place motor neuron.

Female

Senses: object-sensors, orientation, goal, has-object.

Object-sensors: the female senses the object values in its Moore (3x3) neighborhood.

Orientation: north, south, east and west.

Goal: lay-egg, brood-egg, eat-mouse.

Has-object: identical to male.

Internal state: food.

Food is initialized to a parameterized value. When food reaches zero, the need for a mouse is surfaced as a goal. Upon eating a mouse, food is increased by a random value.

Needs: lay-egg-need, brood-egg-need, mouse-need.

The *mouse-need* need is raised by food level and sets the *eat-mouse* goal. It will cause the female to express the *want-mouse* response. While building the nest and not hungry, the *lay-egg-need* is kept raised with the associated *lay-egg* goal. The female asserts the *want-stone* response when she is located in a cell where the nest requires a stone to be placed. After a stone is placed, the female proceeds to the next location in the nest and repeats the process, continuously motivated by the *lay-egg-need* and *lay-egg* goal. When the nest is built, the female lays her egg in the center of the nest. After that *brood-egg-need* is kept raised and the *brood-egg* goal keeps the female brooding on the egg.

Responses:

- *do-nothing*: a no-op response.
- *move-forward*: move forward in the orientation direction. Movement off the grid is a no-op.
- *turn-right/left*: change orientation by 90 degrees.
- *eat-mouse*: eat mouse if has-object is a mouse. If no mouse, this is a no-op.
- *get-object*: if has-object is empty and object in current cell, pick up the object and set it to has-object.
- *put-object*: if has-object and no object in current cell, put object in cell and clear has-object.
- *want-mouse*: when the eat-mouse goal is sensed, the want-mouse response signals the male to retrieve a mouse from the forest for the female to eat.
- *want-stone*: when the lay-egg goal is sensed, and the female is ready to place a stone in the nest, the want-stone response signals the male to retrieve a stone from the desert for the female to place in the nest.
- *lay-egg*: when the female has completed the nest and has moved to its center, she lays the egg with this response.

Artificial Neural Networks

Mona

Mona learns cause-and-effect chains and hierarchies of neurons that represent corresponding relationships in the environment. A detailed description of the architecture can be found in [2]. An overview is provided here.

Three types of neurons are defined, as shown in Fig. 1:

- Receptor neuron: represents a sensory event.
- Motor neuron: expresses response to the environment.
- Mediator neuron: represents a predictive relationship between neurons. The firing of its cause neuron *enables* its motor neuron to fire, producing a response. If the network fires the motor neuron, the effect neuron will probabilistically fire. A mediator that mediates lower level mediators serves as a context that recursively affects the probability of causation in its components.




Fig. 1 A simple Mona network

Mona is a goal-seeking network, falling into the category of a model-based reinforcement learning system [20]. Needs arising from internal and external sources are satisfied by the firing of associated goal neurons. For example, a need for water is satisfied by firing neurons involved with drinking water. Need drives backward through the network from a goal neuron as *motive*, following enabled pathways to fire a motor response that will navigate toward the goal.

Multiple needs can vie for control. For example, food and thirst might both act on the network simultaneously. The winner will be a function of the strength of the need and the enablement of the network to achieve the goal.

In the nest-building task, the male learns a network consisting of chains of mediators that drives it to the forest when it needs a mouse to eat, then returns it to the female for further goal setting. It also learns mediators that orchestrate fetching mice and stones for the female. A powerful feature of Mona is that mediators are modular, that is, they can take part in multiple goal activities. For example, the mediators that fly the male to the forest and hunt for a mouse are used when both the male and female require a mouse.

The female's need to lay an egg, in conjunction with the goal of sensing the egg in the nest, drives through a chain of mediators that orchestrate a series of nest-building actions, each of which involves expressing a want for a stone that signals the male to fetch a stone, placing the stone, and moving to the next location until the nest is entirely constructed. The female then moves to the center of the nest and lays her egg, satisfying the egg laying need. Then the need to brood the egg is raised, which keeps the female sitting on the egg.

Bird responses are trained by overriding incorrect responses with correct ones. The correct responses are incorporated into the network. During testing, responses are generated by the network. Mona is also capable of being trained in an animal-like manner through operant conditioning [21], a learning process where behaviors are modified through the association of stimuli with reinforcement. However, as animal trainers will attest, this process is slow and laborious.

Place Motor Neurons

To enhance Mona with place cell functionality, a place motor neuron was implemented. A biological place neuron fires when a specific location in the world is reached. Mona's place neurons fire similarly. However, they also navigate to

specific places in the world. A related capability seems to exist in the form of route encoding in rats [22].

In Mona, a place neuron is implemented as a motor neuron that records a specific location in the world and will navigate to that location when motivated to respond. It fires when the location is reached. For example, when the male's goal is to fetch a mouse for the female, a mediator with a cause of sensing the female's want of a mouse, and an effect of sensing the proximity of a mouse in the forest will fire its place motor to enact a series of primitive movement responses that navigate to a prescribed location in the forest.

Place motors can be learned while exploring the world by marking significant locations. A special response initiates the commencement of a sequence of movements that terminate at some location, at which time another special response marks the location, creating a new place motor neuron. In the nesting birds, these two special responses are mapped to the male *fly* and *alight* responses, respectively.

LSTM

The LSTM, introduced in 1997 [18], is a recurrent neural network which has established itself as a workhorse for sequential pattern recognition. LSTMs address a problem with other recurrent neural networks in which corrective information vanishes as the time lag between the output and the relevant input increases, leading to the inability to train long-term state information.

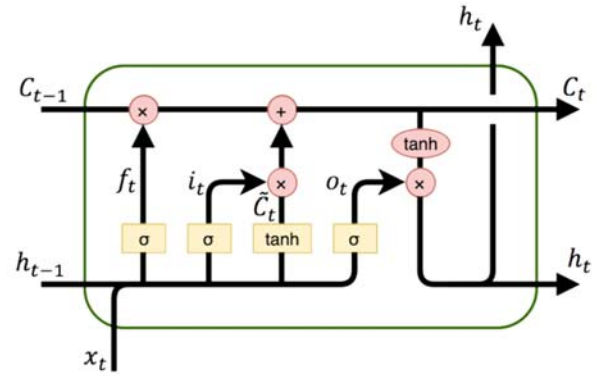


Fig. 2 LSTM memory block

In the LSTM network, the hidden units of a neural network are replaced by memory blocks, each of which contains one or more memory cells. A memory block is shown in Fig. 2. The block can latch and store state information indefinitely, allowing long-term temporal computation. What information to store, and when to output that information are part of the training process.

Scenario

The scenario is taken from the game proposal [1], which illustrates the task from a general viewpoint. It shows intermittent states of the world, from initial state to egg-laying in the completed nest. A video is available [23].

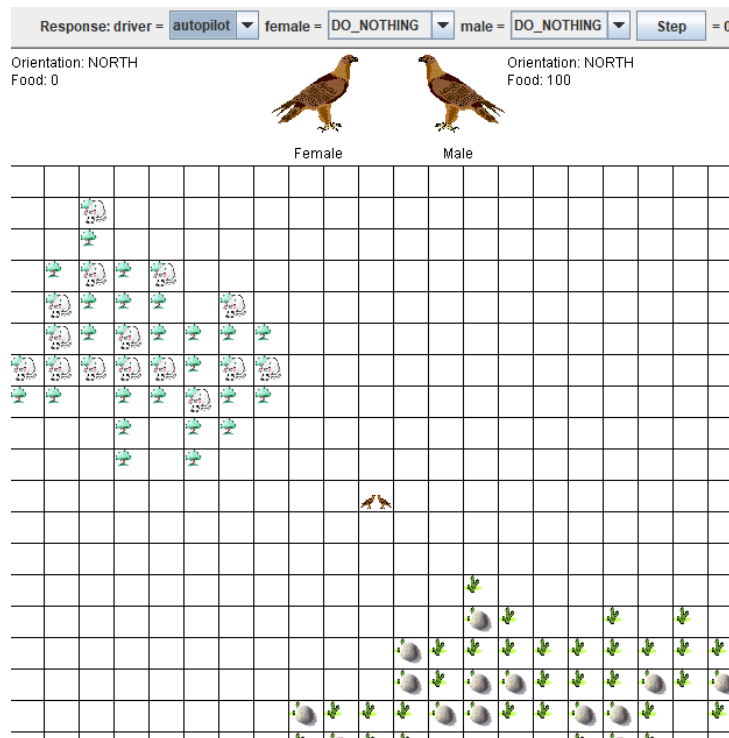


Fig. 3 Beginning state: female is hungry (0 food), male has maximum food; initial response for both is “do nothing”; both are located at center of world; upper left is forest with mice (food); lower right is desert with stones for nest-building

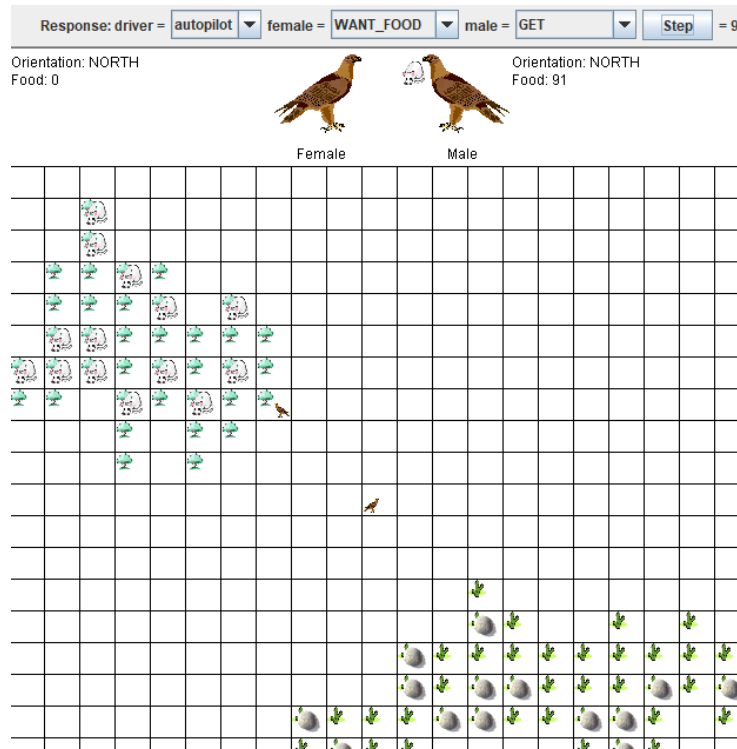


Fig. 4 While co-located, female signals to male with “want food” response; male flies to forest and picks up a mouse to feed to her

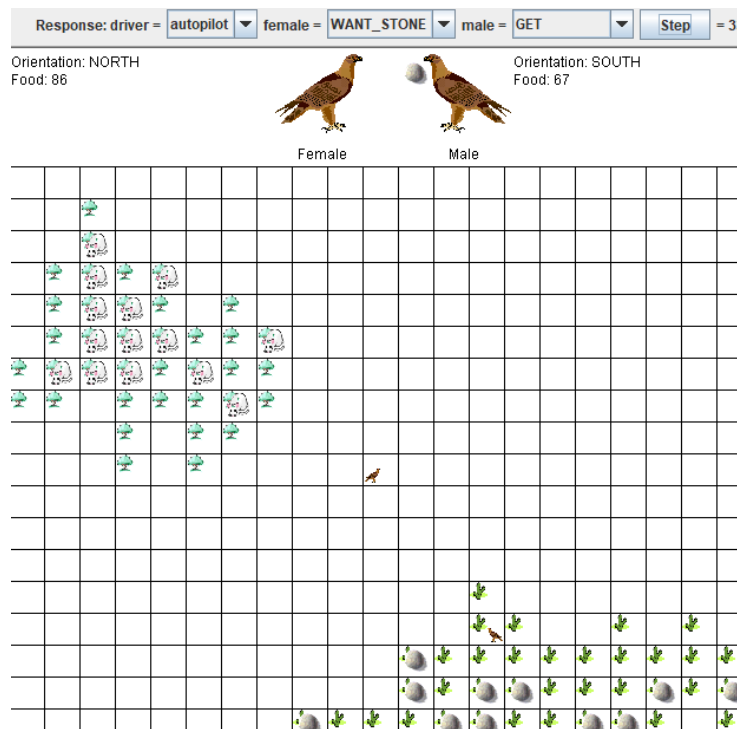


Fig. 5 Female moves to location of first nesting stone; male follows her; she signals to male that she wants a stone; male flies to desert and picks up a stone

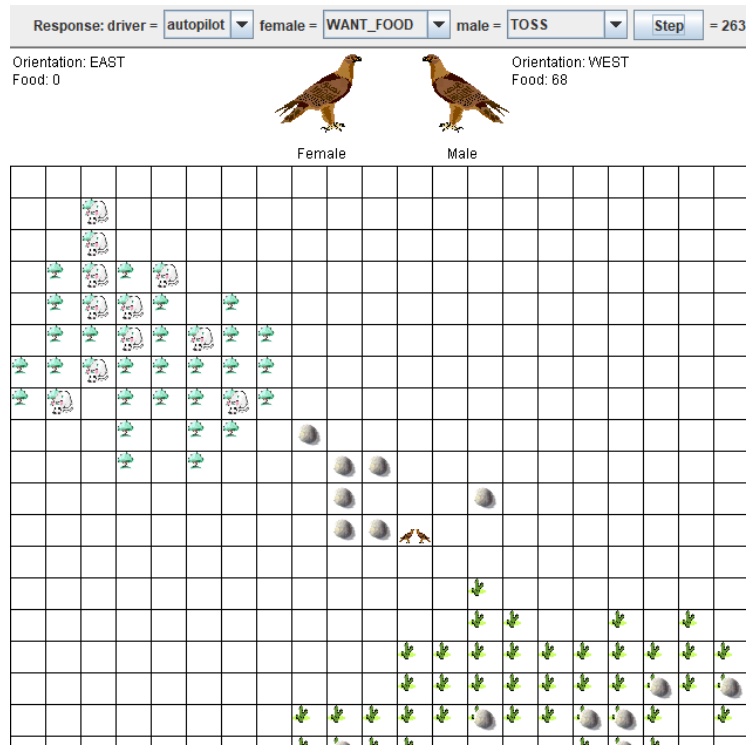


Fig. 6 Male returns to female with stone; discovers she is hungry; he flies to forest for mouse for her

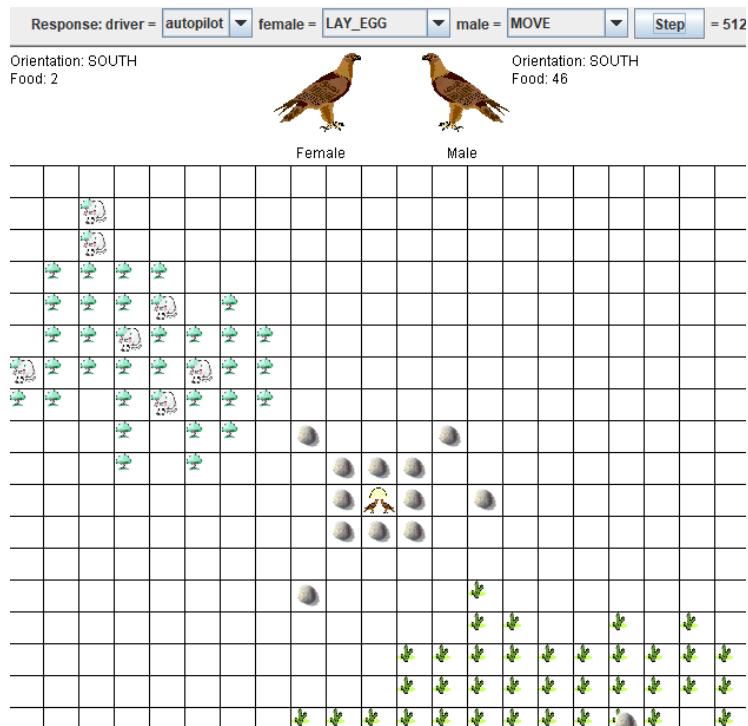


Fig. 7 Nest completed; egg laid

III. RESULTS

Two ANNs were trained and tested on the task under varying conditions:

1. Mona version 6.1. There are a maximum of 500 mediator neurons. There is a maximum mediator level of 0,

meaning mediators mediated receptor and motor neurons exclusively; higher level mediators that mediate lower level mediators were not needed.

2. An LSTM recurrent neural network using the Keras 2.6.0 python package. There are 128 neurons in a hidden layer, with a mean squared error loss function. Input and output

were one-hot encoded. Training was conducted with 500 epochs.

Number of Training Datasets

Performance was measured based on the number of training datasets, as shown in Figs. 8 and 9. A training dataset is generated by running the nesting birds with optimal bird responses until the nest is completed while keeping the birds

fed.

Each run is seeded with a different random number that controls the movements of the mice in the forest and how long eating a mouse serves to satisfy a bird's hunger. The testing accuracy was calculated as the percent of correct responses out of 1000 steps. Measurements were based on the mean values of 20 trials. Very little variance in values was observed.

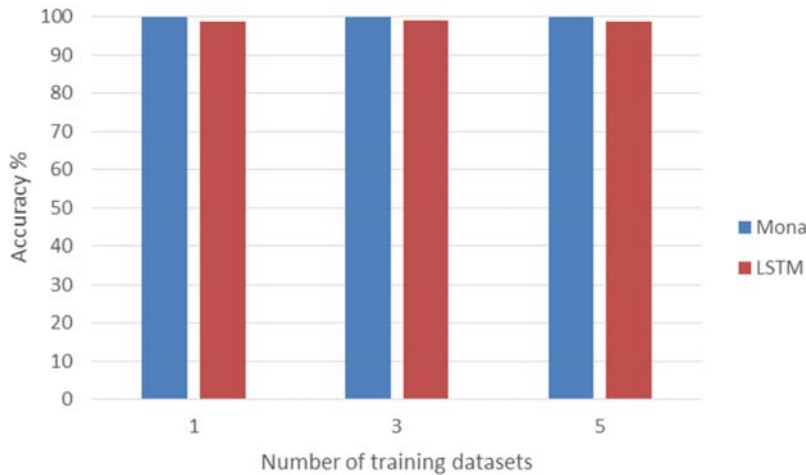


Fig. 8 Female test performance with number of training datasets




Fig. 9 Male test performance with number of training datasets

Both networks show excellent performance with only a single training dataset.

Dynamic Testing

Mona performance is measured as it interacts with the world. That is, responses output to the world cause changes in the world that are subsequently input to the bird. For example, if the bird turns to the left, the world reacts by altering the sensory state of the bird accordingly. With place motor neurons, Mona solves the task every time, but without place motors the male bird, which must perform complex navigation to fetch mice and stones, repeatedly becomes lost, causing the entire task to fail.

When the same “dynamic” interface is applied to the RNN network the male bird repeatedly becomes lost while fetching mice or stones. This means that even the few errors that the male makes are crucial, preventing successful completion. If the male cannot return with a mouse for the female, for example, the female cannot proceed with nest building.

RNN Epoch Testing

Mona trains in a single epoch, a skill frequently seen in human learning [24]. The RNN training is significantly affected by the number of epochs of training, especially for the male, as shown in Fig. 10.

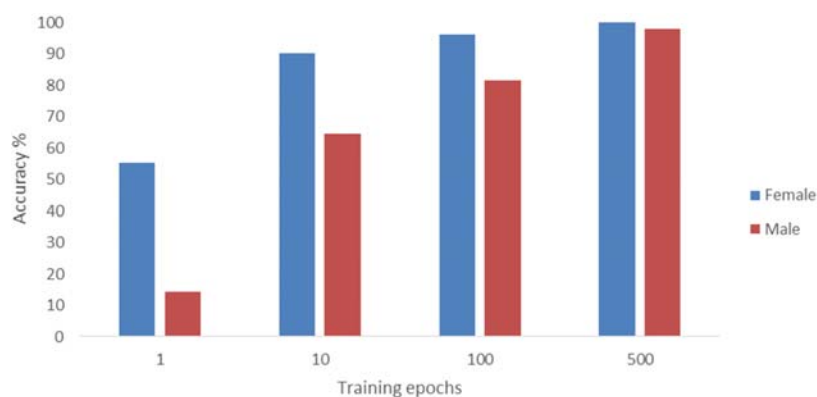


Fig. 10 Training epochs testing

IV. CONCLUSION

Enhanced with place motor neurons, Mona is capable of solving the nest-building task every time. The RNN performs well with its typical usage, which is to predict upcoming responses, but fails as a navigator for the male bird, causing nest-building to be unsuccessful. How place neuron functionality might be incorporated into an RNN is an interesting topic.

Place motor neurons and goal-seeking causation learning are a powerful combination of capabilities for the nest-building task which demands both spatial and sequential learning. This animal learning task exposes shortcomings in a deep learning ANN that researchers interested in AGI should be aware of. We recommend that further research is conducted to (1) further simulate animal behaviors, and (2) adopt mechanisms from neurobiology, such as place cells, that allow machines to acquire animal-like capabilities. These we think are essential to the achievement of AGI.

CONFLICT OF INTEREST

There are no conflicts of interest concerning this work.

REFERENCES

- Portegys, T.E. (2022). I want to play a game. https://www.researchgate.net/publication/364089546_I_WANT_TO_PLAY_A_GAME
- Portegys, T.E. (2001). Goal-Seeking Behavior in a Connectionist Model. *Artificial Intelligence Review* 16, 225–253 (2001). <https://doi.org/10.1023/A:1011970925799>
- Moser M.B., Rowland D.C., Moser E.I. Place cells, grid cells, and memory. (2015). *Cold Spring Harbor Perspectives in Biology*. Feb 2;7(2):a021808. doi: 10.1101/cshperspect.a021808. PMID: 25646382; PMCID: PMC4315928.
- Robinson, N. T. M., Descamps, L. A. L., Russell, L. E., Buchholz, M. O., Bicknell, B. A., Antonov, G. K., Lau, J. Y. N., Nutbrown, R., Schmidt-Hieber, C., Häusser, M. (2020). Targeted activation of hippocampal place cells drives memory-guided spatial behavior. *Cell*, 183, pp. 1586-1599.
- Xu, H., Baracska, P., O'Neill, J., and Csicsvari, J. (2019). Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze. *Neuron* 101, 119–132.
- Milford, M. and Wyeth, G., (2010). Persistent navigation and mapping using a biologically inspired SLAM system, *Int. J. Robot. Res.* 29:1131–1153.
- Zhou, X., Weber, C. and Wermter, S., (2017). Robot localization and orientation detection based on place cells, *Proc. ICANN 2017*, Springer pp. 137–145.
- OpenAI. (2023). *ChatGPT* (Mar 14 version) Large language model. <https://chat.openai.com/chat>
- Zador, A. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications* volume 10, Article number: 3770. <https://www.nature.com/articles/s41467-019-11786-6>
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press. "*Vehicles - the MIT Press*". Archived from the original on 2010-01-29. Retrieved 2012-06-18.
- Dyer, M.G. (1993). Toward Synthesizing Artificial Neural Networks that Exhibit Cooperative Intelligent Behavior: Some Open Issues in Artificial Life. *Artificial Life*, vol. 1, no. 1_2, pp. 111-134, Oct. 1993, doi: 10.1162/artl.1993.1.1_2.111.
- Coleman S.L., Brown V.R., Levine D.S., Mellgren R.L. (2005). A neural network model of foraging decisions made under predation risk. *Cogn Affect Behav Neurosci.* 2005 Dec;5(4):434-51. doi: 10.3758/cabn.5.4.434. PMID: 16541813.
- Erickson, J., Moses, M. and Forrest, S. (2017). Automatically evolving a general controller for robot swarms. *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, Honolulu, HI, USA, 2017, pp. 1-8, doi: 10.1109/SSCI.2017.8285399
- Stanley, K. O. and R. Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, vol. 10, no. 2, pp. 99–127.
- Lizier, J.T., Piraveenan, M., Pradhana, D., Prokopenko, M., Yaeger, L.S. *Functional and Structural Topologies in Evolved Neural Networks. ECAL 2009.*
- Enquist, M. and Ghirlanda, S. (2006). *Neural Networks and Animal Behavior*. Volume 33 in the series *Monographs in Behavior and Ecology* Published by Princeton University Press. <https://doi.org/10.1515/9781400850785>
- Wijeyakulasuriya, D.A., Eisenhauer, E.W., Shaby, B.A., Hanks E.M. (2020). Machine learning for modeling animal movement. *PLoS ONE* 15(7): e0235750. <https://doi.org/10.1371/journal.pone.0235750>
- Hochreiter, S., Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
- Code with instructions: <https://github.com/morphognosis/NestingBirds>
- Moerland, T.M., Broekens, J., Plaat, A., Jonker, C.M. (2023). *Model-based Reinforcement Learning: A Survey*. Foundations and Trends in Machine Learning Series. Now Publishers. <https://books.google.com/books?id=FimgzweEACAAJ>
- Jenkins, H. M. (1979). *Animal Learning & Behavior Theory*. Ch. 5 in Hearst, E. *The First Century of Experimental Psychology* Hillsdale N. J., Earlbaum.
- Grieves R. M., Wood E. R., Dudchenko P. A. (2016). Place cells on a maze encode routes rather than destinations. *Elife*. Jun 10;5:e15986. doi: 10.7554/eLife.15986.
- Scenario video: <https://youtu.be/d13hxhltsGg>
- Lee, S.W., O'Doherty, J.P., Shimojo, S., (2015). Neural computations mediating one-shot learning in the human brain. *PLoS biology* 13.