bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

MORPHOGNOSTIC HONEY BEES COMMUNICATING NECTAR LOCATION THROUGH DANCE MOVEMENTS

Thomas E. Portegys
portegys@gmail.com, Dialectek, DeKalb, IL, USA

ABSTRACT

Honey bees are social insects that forage for flower nectar cooperatively. When an individual forager discovers a flower patch rich in nectar, it returns to the hive and performs a "dance" in the vicinity of other bees that consists of movements communicating the direction and distance to the nectar source. The bees that receive this information then fly to the location of the nectar to retrieve it, thus cooperatively exploiting the environment. This project simulates this behavior in a cellular automaton using the Morphognosis model. The model features hierarchical spatial and temporal contexts that output motor responses from sensory inputs. Given a set of bee foraging and dancing exemplars, and exposing only the external input-output of these behaviors to the Morphognosis learning algorithm, a hive of artificial bees can be generated that forage as their biological counterparts do.

Keywords: Honey bee foraging dance, artificial animal intelligence, Morphognosis, cellular automaton, artificial neural network.

INTRODUCTION

Honey bees, *apis mellifera*, are fascinating social insects. They are also smart, even able to count and add (Fox, 2019). However, it is their ability to communicate symbolically in the form of a "dance" indicating the direction and distance to a nectar source that is truly astonishing (Chittka and Wilson, 2018; Hooper, 2017; Nosowitz, 2016; Schürch et al., 2019; Von Frisch, 1967), especially considering that the use of symbols is rare even in more neurologically complex animals. The dance, done by a bee in the presence of other bees in the hive after discovering nectar at a locale outside the hive, recruits bees to forage at the indicated location, thus acquiring more nectar than solitary foraging would otherwise.

This paper describes artificial honey bees that gather nectar and perform the foraging dance. It employs a general machine learning system, *Morphognosis*, which acquires behaviors by example and enables an artificial organism to express those behaviors. It will be shown that simulating nectar foraging is a difficult task for unaugmented machine learning methods, but with the support of the spatial and temporal contextual information that Morphognosis provides, it can be accomplished.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

As a disclaimer, it should be noted that this project is not intended to offer new or additional findings about honey bees. This is an artificial intelligence project, not a biology study. Its goal is to simulate a biologically inspired behavior in a dynamic environment. It is only functional, "black box" behaviors that are simulated; modeling internal mechanisms is not the aim here.

Honey bees have been the focus and inspiration for a number of simulation initiatives:

- Detailed colony behavior (Betti et al., 2017).
- Swarming and group behavior algorithms (Karaboga and Akay, 2009).
- Flight neural network (Cope et al., 2013).
- Visual system neural network (Roper et al., 2017).
- Odor learning circuits (MaBouDi et al., 2017).
- Spiking neural network that reacts to nectar (Fernando and Kumarasinghe, 2015)

The colony simulation allows a user to observe how bees are affected by various environmental conditions, such as weather. Algorithms for a number of group behaviors, optimal foraging strategies among them, are cited in the Karaboga and Akay paper. The other projects simulate bee-specific neural mechanisms. For example, the odor learning project found that simulated honey bees lacking mushroom bodies, the insect equivalent of the cerebral cortex, may still be able to learn odors. The spiking neural network measures how an abstracted model of a bee's nervous system reacts to nectar-related stimuli.

In contrast, the contribution of this project is to simulate honey bee behavior with a general purpose connectionist model that learns from external observations and which is applicable to arbitrary behavioral simulation tasks, not just the honey bee foraging task.

A number of years ago I explained to a coworker how my dissertation program (Portegys, 1986), a model of instrumental/operant conditioning, could learn various tasks through reinforcement. He then asked me how "smart" it was. I put him off, not having a ready answer. He persisted. So I blurted out that it was as smart as a cockroach (which it is not). To which he replied, "Don't we have enough *real* cockroaches?" Fast forward to this project. Don't we have enough *real* honey bees? (Although, come to think of it, maybe we don't (Oldroyd, 2007)!)

The point of this story is that the question of why anyone should work on artificial animal intelligence is, at least on the surface, a reasonable one, given our species unique intellectual accomplishments. Thus, historically, AI has mostly focused on human-like intelligence, for which there are now innumerable success stories: games, self-driving cars, stock market forecasting, medical diagnostics, language translation, image recognition, etc. Yet the elusive goal of artificial general intelligence (AGI) seems as far off as ever. This is because these success stories lack the

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

"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.

Anthony Zador (2019) 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."

I am in the camp that believes that achieving general animal intelligence is a necessary, if not sufficient, path to AGI. While imbuing machines with abstract thought is a worthy goal, in humans there is a massive amount of ancient neurology that underlies this talent.

Hans Moravec put it thusly (1988): "Encoded in the large, highly evolved sensory and motor portions of the human brain is a billion years of experience about the nature of the world and how to survive in it. The deliberate process we call reasoning is, I believe, the thinnest veneer of human thought, effective only because it is supported by this much older and much more powerful, though usually unconscious, sensorimotor knowledge. We are all prodigious Olympians in perceptual and motor areas, so good that we make the difficult look easy. Abstract thought, though, is a new trick, perhaps less than 100 thousand years old. We have not yet mastered it. It is not all that intrinsically difficult; it just seems so when we do it."

So how should we proceed? Emulating organisms at the level of neurons (whole-brain emulation) is a possible approach to understanding animal intelligence. However, efforts to do this with the human brain have met with little success (Yong, 2019). Scaling down to mice is an option. The human brain dwarfs the mouse brain, but even mouse brains are daunting: a cubic milliliter of mouse cortex contains 900,000 neurons and 700,000,000 synapses (Braitenberg and Schüz, 1998). At much a simpler scale, years have been spent studying the relationship between the connectome of the nematode C. elegans (Wood, 1988), with only 302 neurons, and its behaviors, but even this creature continues to surprise and elude full definition. Nevertheless, some researchers believe that it is now feasible for the whole-brain approach to be applied to insects such as the fruit fly, with its 135,000 neurons (Collins, 2019). Partial brain analysis is also an option. For example, the navigation skills of honey bees are of value to drone technology. Fortunately, it appears that the modular nature of the honey bee brain can be leveraged to replicate this skill (Nott, 2018).

An issue with emulation, besides the complexity, is the difficulty of mapping the relationship between specific neural structures and behaviors (Krakauer et al., 2016; Yong, 2017). For AI, this is a key aspect, as behavior is the goal. Artifacts and quirks left over by evolution introduce unnecessary complexity in neurological systems. Moreover, even a comprehensive emulation might defy description. Nature, always a blind tinkerer, has been known to work this way. For

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

example, despite the enthusiasm following the mapping of the human genome, the mechanisms by which genes express proteins, and thus phenotypes, is not as modular as hoped for. Rather, it is extraordinarily complex (Wade, 2001).

The field of artificial life (Alife) offers another path to AGI. This path starts with simulating life, and letting evolution optimize artificial organisms to achieve intelligence as a fitness criteria. Schöneburg's (2019) "alternative path to AGI", sees intelligence emerging from *holobionts*, which form cooperating collectives of artificial agents.

Another approach, one taken here, is to simulate AI at the behavioral, or functional, level. AI is no more constrained to adhere to biology than aeronautics is confined to bird flight. Considering the vastly different "clay" that biological and computing systems are built with, cells vs. transistors and software, behavioral simulation offers another possible path to artificial intelligence.

Morphognosis comprises an artificial neural network (ANN) enhanced with a framework for organizing sensory events into hierarchical spatial and temporal contexts. Nature has hard-wired knowledge of space and time into the brain as way for it to effectively interact with the environment (Bellmund et al., 2018; Hainmüller and Bartos, 2018; Lieff, 2015; Vorhees and Williams, 2014). These capabilities are modeled by Morphognosis. Interestingly, grid cells also appear in humans to be capable of representing not only spatial relationships, but non-spatial multidimensional ones, such as the relationships between members of a group of people (Bruner et al., 2018; Tarvaras et al., 2015).

The bee dancing behavior, as a sequential process, has temporal components. For example a bee must remember a past event, the existence of surplus nectar in a flower, and use that information to perform a dance that indicates both direction and distance to the nectar. In addition, bees that observe a dance must internally persist the distance signal and use it to measure how far to fly.

Sequential processes are type of task that recurrent artificial neural networks (RNNs) have been successfully applied to (Elman, 1990; Hochreiter and Schmidhuber, 1997). However, RNNs do not inherently also support spatial information. RNNs maintain internal feedback that allow them to retain state information within the network over time. This contrasts with Morphognosis, where the input itself contains temporal state information.

Morphognosis was partly inspired by some what-if speculation. In simpler animals, the "old" brain (amygdala, hypothalamus, hippocampus, etc.) deals more directly with an unfiltered here-and-now version of the environment. Considering nature's penchant for repurposing existing capabilities, might it be that in more complex animals a purpose of the neocortex, sitting atop the old brain and filtering incoming sensory information, is to track events from distant reaches

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

of space and time and render them, as though near and present, to the old brain whose primal functions have changed little over time?

The author has previously conducted research explorations into a number of issues that differentiate conventional AI from natural intelligence. These include context, motivation, plasticity, modularity, instinct, and surprise (Portegys, 2007, 2010, 2013, and 2015). Morphognosis, in particular, has been previously applied to the task of nest-building by a species of pufferfish (Portegys, 2019).

To date, including the honey bee project, Morphognosis has been implemented as a cellular automaton (Toffoli and Margolus, 1987; Wolfram, 2002), as the rules that it develops while learning are ideally captured in a grid structure. Conceptually, however, Morphognosis is not tied to the cellular automaton scheme.

DESCRIPTION

This section first briefly describes the Morphognosis model. The honey bee implementation is described next.

MORPHOGNOSIS OVERVIEW

Morphognosis (*morpho* = shape and *gnosis* = knowledge) aims to be a general method of capturing contextual information that can enhance the power of an artificial neural network (ANN). It provides a framework for organizing spatial and temporal sensory events and motor responses into a tractable format suitable for ANN training and usage.

Introduced with several prototype tasks (Portegys, 2017), Morphognosis has also modeled the locomotion and foraging of the C. elegans nematode worm (Portegys, 2018) and the nest-building behavior of a pufferfish (Portegys, 2019). Morphognosis is a temporal extension of a spatial model of morphogenesis (Portegys, et al., 2017).

MORPHOGNOSTICS

The basic structure of Morphognosis is a cone of sensory event recordings called a *morphognostic*, shown in Figure 1. At the apex of the cone are the most recent and nearby events. Receding from the apex are less recent and possibly more distant events. A morphognostic can thus be viewed as a structure of progressively larger nested chunks of space-time knowledge that form a hierarchy of contexts. A set of morphognostics forms long-term memories that are learned by exposure to the environment. Scaling can be accomplished by aggregating event information. This means that more recent and nearby events are recorded in greater precision than events more distant in space and time.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

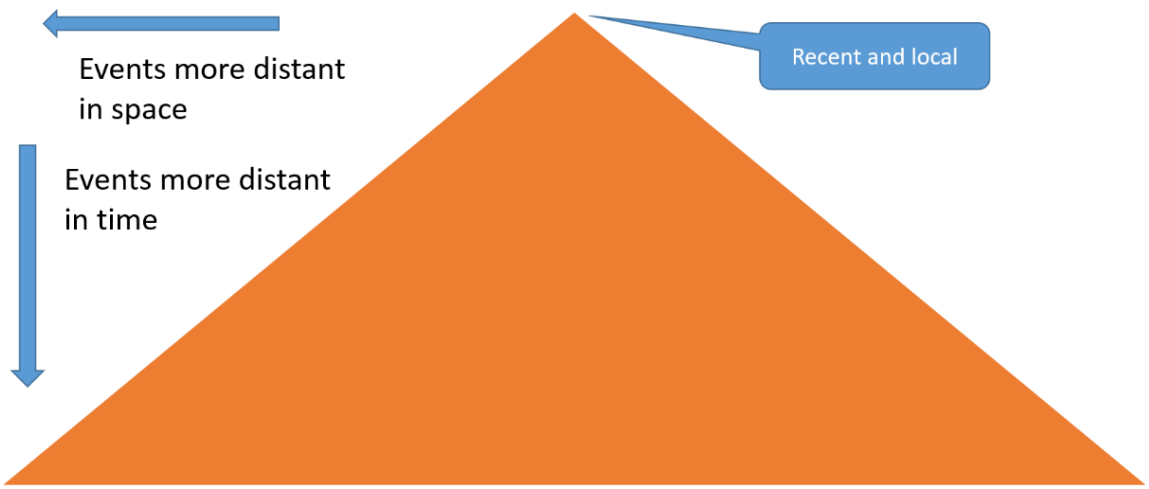

Figure 1 - Morphognostic event cone.

The following are possible definitions of the spatial and temporal morphognostic neighborhoods. The software is parameterized to permit many variations of these definitions.

Morphognostic spatial neighborhoods

A cell defines an elementary neighborhood:

$$neighborhood_0 = cell \quad (1)$$

A non-elementary neighborhood consists of an *NxN* set of *sectors* surrounding a lower level neighborhood:

$$neighborhood_i = NxN(neighborhood_{i-1}) \quad (2)$$

Where *N* is an odd positive number.

The value of a sector is a vector representing a histogram of the cell type densities contained within it:

$$value(sector) = (density(cell\text{-}type_0), density(cell\text{-}type_1), \ldots density(cell\text{-}type_n)) \quad (3)$$

Where the number of cells contributing to the density histogram of a sector of $neighborhood_i = N^{i-1}xN^{i-1}$ (4)

Morphognostic temporal neighborhoods

A neighborhood contains events that occur within a *duration*, which is a time window between the present and some time in the past. Here is a possible method for calculating the duration of each neighborhood algorithmically:

$$t1_0 = 0$$

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

$t2_0 = 1$

$t1_i = t2_{i-1}$

$t2_i = (t2_{i-1} * 3) + 1$

$duration_i = t2_i - t1_i$

Morphognostic example

Figure 2 is an example of a morphognostic implemented in a cellular automaton as a nested set of 3x3 neighborhoods and aggregated histograms of cell state value densities. On the left is the cellular automata grid that contains various cell state values. Moving right is a 3x3 neighborhood surrounding one of the cells of interest. This neighborhood is sensed in the immediate present. Moving right to the 9x9 neighborhood, each sector is a 3x3 neighborhood aggregated in space and time. Thus the densities are variable. The rightmost panel continues this theme to a 27x27 neighborhood.

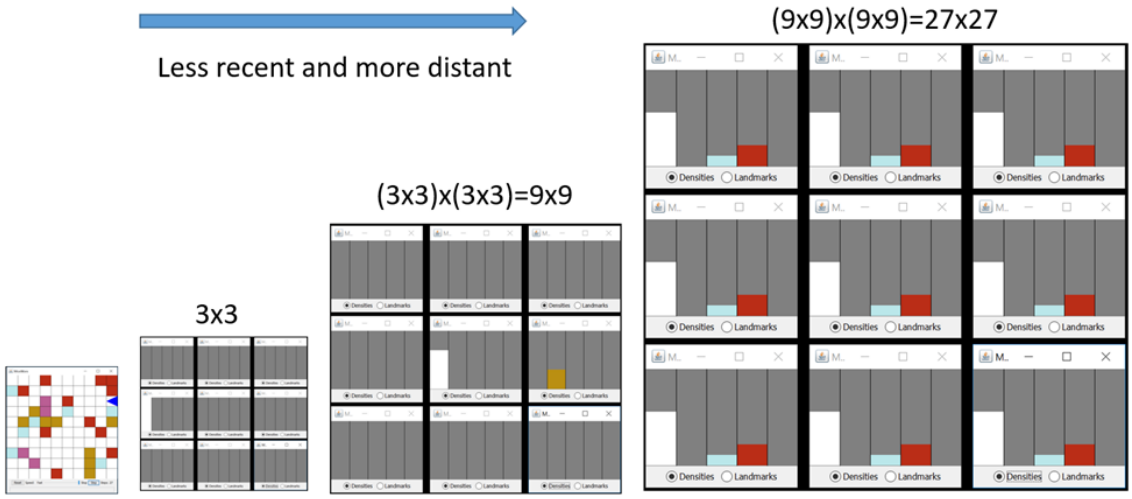

Figure 2 – Cellular automaton implementation of Morphognosis.

Metamorphs

In order to navigate and manipulate the environment, it is necessary for an agent to be able to respond to the environment. A *metamorph* embodies a morphognostic→response rule. A set of metamorphs can be learned from a manual or programmed sequence of responses within a world.

Metamorphs are used to train an ANN, as shown in Figure 3, to learn responses associated with morphognostic inputs. During operation, the current morphognostic, representing the state of the environment, is input to the ANN to produce a learned response.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

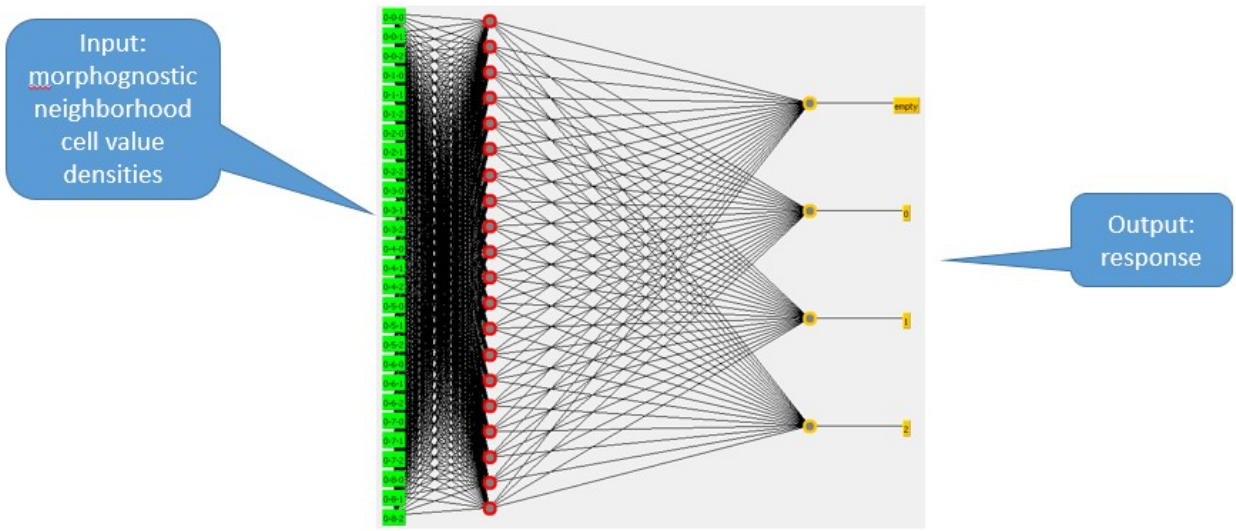

Figure 3 – Metamorph artificial neural network.

HONEY BEES

BEHAVIOR

A brief explanatory video is available on YouTube:
https://www.youtube.com/watch?v=kUAv2QO7qYM

SENSORY AND RESPONSE CAPABILITIES

Senses:

External state:

- Hive presence.
- Nectar presence.
- In-hive bee nectar signal: Orientation and distance to nectar.

Internal state:

- Orientation.
- Carrying nectar.

Responses:

- Wait
- Move forward
- Turn in compass directions: N, NE, E, SE, S, SW, W, NW
- Extract nectar
- Deposit nectar

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

- Display nectar distance.

World

Figure 4 shows a graphical view that shows a hive (central yellow area), three bees, and three flowers. The topmost flower contains a drop of nectar to which the topmost bee, as best it can in a cellular grid, is indicating the direction and an approximate distance to, as indicated by the orientation of the bee and the length of the dotted line, respectively.

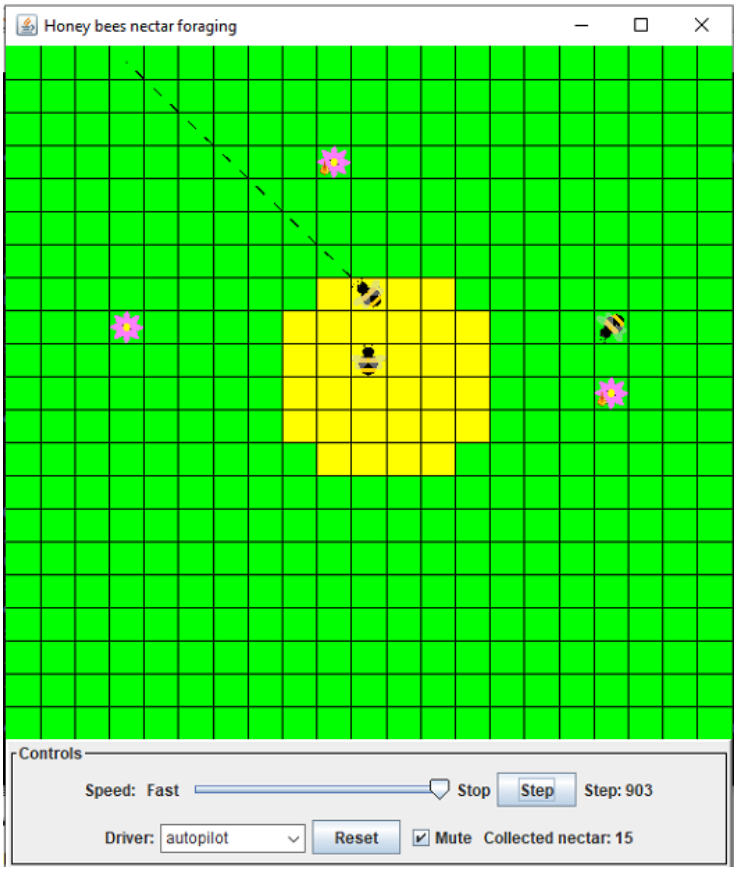

Figure 4 – Graphical view.

Flowers

Flowers are initialized with nectar, which after being extracted by a bee, will probabilistically either replenish after a specific time or immediately replenish. In the latter case, the bee will sense the presence of surplus nectar and will perform a dance to indicate its direction and distance once it returns to the hive.

Foraging

The bees forage in two phases. In phase one, the nectar discovery phase, a bee flies about semi-randomly until it encounters a flower with nectar. Phase two is a deterministic process that deals with known nectar. Phase two is described below.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Once discovered, the bee extracts the nectar from the flower, flies directly to the hive and deposits the nectar in the hive. If the bee, after depositing the nectar, remembers that the flower contained "surplus" nectar, meaning more nectar than the bee could carry, it will commence a dance to indicate the direction and distance to the nectar to other bees in the hive, including itself. The direction is indicated by orienting toward the nectar. The directions are confined to the eight compass points. The distance is indicated by displaying a value for short or long distance. Both direction and distance can be sensed by bees in the hive. The graphical view draws a short or long dotted line as a visual representation.

Once a bee completes the dance, it and any other bees in the hive that sensed the dance will proceed in the direction of the nectar for the distance exhibited by the dance. If any of these bees encounters nectar en route, it will switch over to extracting the nectar and returning with it to the hive, possibly performing a dance there. If no nectar is encountered en route after traveling the indicated distance, the bee resumes phase one foraging.

If no surplus nectar was sensed after extracting the nectar, the bee will switch to phase one foraging immediately after depositing the nectar.

SCENARIO

Figures 5 through 11 present a graphical nectar foraging scenario.

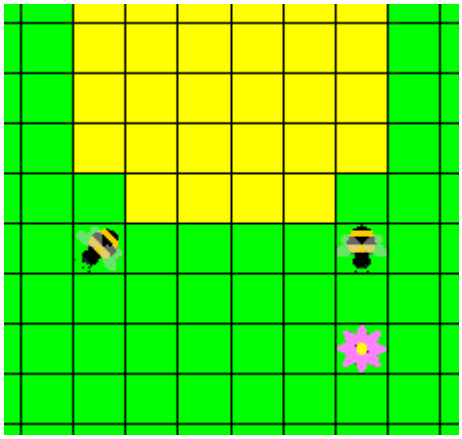

Figure 5 - Bee on right is moving down and is about to light on flower.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

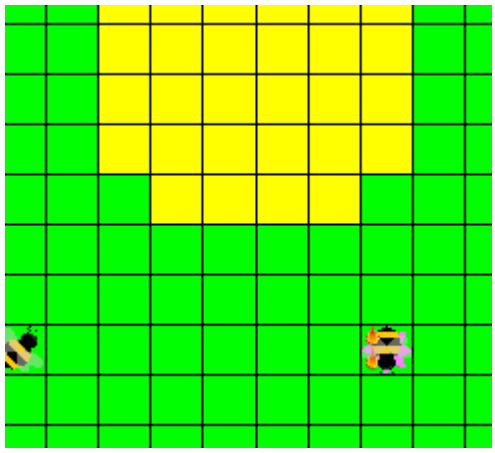

Figure 6 – Bee has extracted nectar from flower.

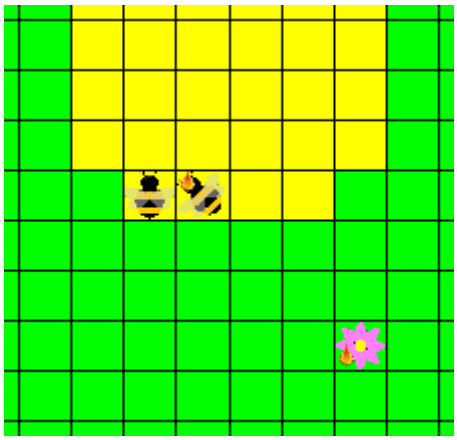

Figure 7 – Bee with nectar returns directly to the hive to deposit nectar. It is also aware of surplus nectar remaining in the flower. The other bee is incidentally also in the hive.

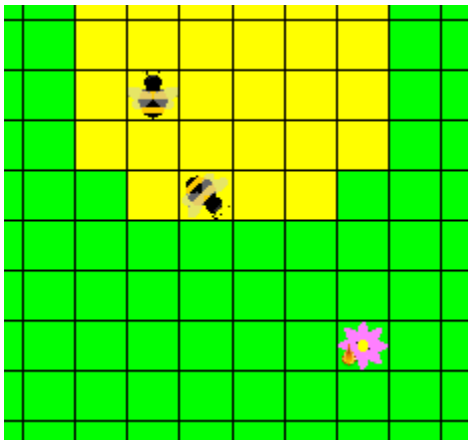

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Figure 8 – Bee has deposited nectar in the hive. Since the bee knows there is surplus nectar, the bee performs the first part of dance: orient toward nectar. If there was no surplus nectar the bee would resume foraging. The other bee is moving about the hive.

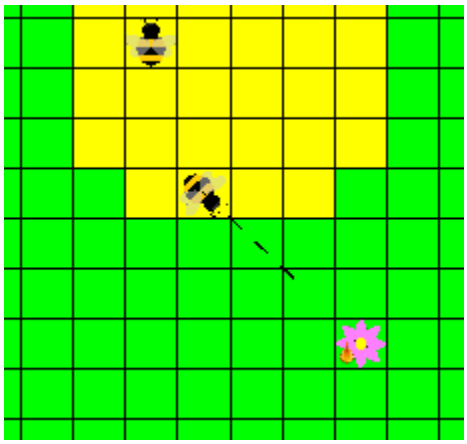

Figure 9 – The second part of dance: indicate a short distance to nectar, as shown by the dotted line. The other bee has become aware of the direction and distance to the nectar.

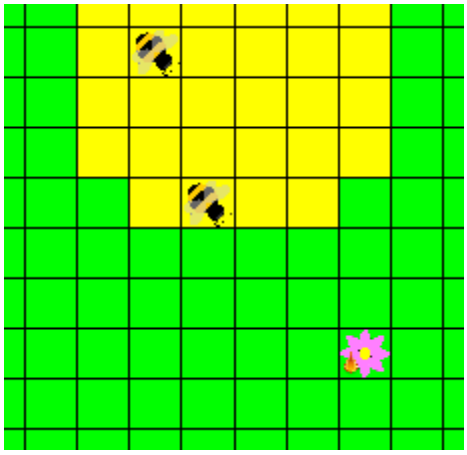

Figure 10 - Both bees respond to dance by orienting toward nectar.

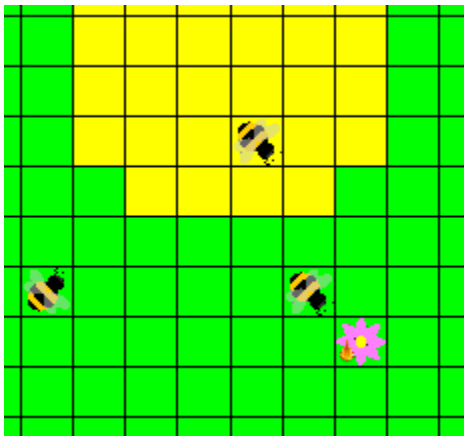

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Figure 11 - Both bees move toward nectar.

Implementation

Modes

In *autopilot* mode, the bees forage programmatically. Autopilot mode generates metamorphs that are used to train the neural network, as shown in Figure 12. Since phase one foraging consists of semi-random movements, metamorphs are only generated in phase two, dealing with known nectar. Once trained, the bees can be switched to *metamorphNN* mode, in which the neural network drives phase two behavior. Phase one behavior remains programmatic in metamorphNN mode. While in metamorphNN mode, new metamorphs are not accumulated.

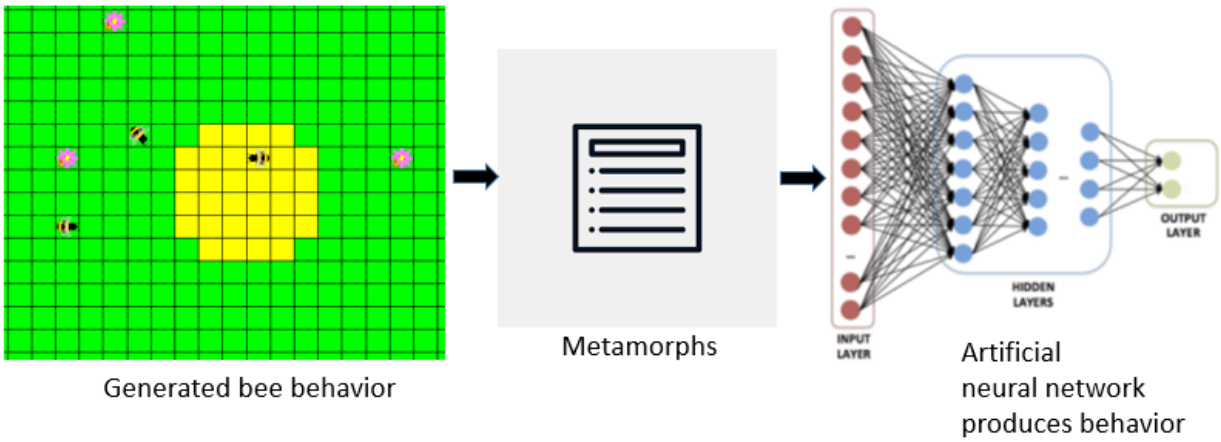

Figure 12 – Generating metamorphs to train the neural network.

Morphognostic

Each bee contains a morphognostic that maps its sensory inputs as spatial and temporal events that maintain its state in the environment.

Events

There are 22 binary event variables:

0. hive presence
1. nectar presence
2. surplus nectar presence
3. nectar dance direction north
4. nectar dance direction northeast
5. nectar dance direction east
6. nectar dance direction southeast
7. nectar dance direction south

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

8. nectar dance direction southwest
9. nectar dance direction west
10. nectar dance direction northwest
11. nectar dance distance long
12. nectar dance distance short
13. orientation north
14. orientation northeast
15. orientation east
16. orientation southeast
17. orientation south
18. orientation southwest
19. orientation west
20. orientation northwest
21. nectar carry

Neighborhoods

The morphognostic contains 4 3x3 neighborhoods, with durations and event mappings shown in Table 1.

Neighborhood	Duration	Events
0	1	All except: *surplus nectar presence* *nectar dance distance long* *nectar dance distance short*
1	7	*hive presence* *surplus nectar presence* *nectar dance short distance*
2	10	*hive presence* *surplus nectar presence* *nectar dance long distance*
3	75	*hive presence*

Table 1 – Morphognostic neighborhoods.

Neighborhood 0 maps “immediate” events, such as orientation, that are of use only in the present, as denoted by the duration of 1.

Neighborhood 1 has a duration, 7, that allows a bee to retain knowledge of the presence of surplus nectar and/or observation of a dance indicating a short distance. The *nectar dance short distance event*, for example, allows the bee to “count” steps towards surplus nectar. When the event expires due to the duration of the neighborhood it no longer affects the bee’s behavior.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Neighborhood 2 serves the same purpose as neighborhood 1, except for *nectar dance long distance* event, for which the duration, and thus steps, is greater than for the *nectar dance short distance* event.

Neighborhood 3, as well as all the other neighborhoods, track the presence of the hive as it is recorded in its 3x3 sectors for a long duration of 75. This allows the bee to locate the hive after possibly lengthy foraging and return with nectar. After 75 steps, without returning to the hive, its location will be lost and the bee will have to return "blindly".

Morphognostic neighborhoods can be configured to either keep a density/average value of event values over its duration, or an on/off event value, meaning the event value is 1 if the event occurs at any time within the neighborhood's duration window. Although surrendering information, the on/off configuration is chosen for the honey bees to improve training time while retaining acceptable performance.

Example

Figures 13a and 13b show the state of the bee selected by the red square for neighborhood 2 of its morphognostic.

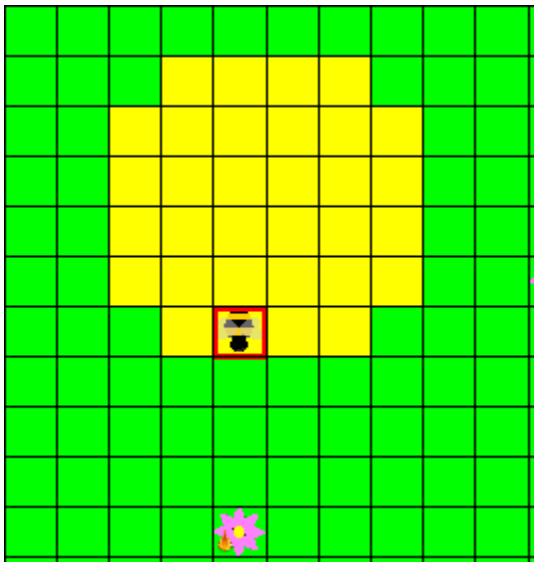

Figure 13a – Bee after dance indicating surplus nectar. The next step is to proceed toward nectar.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

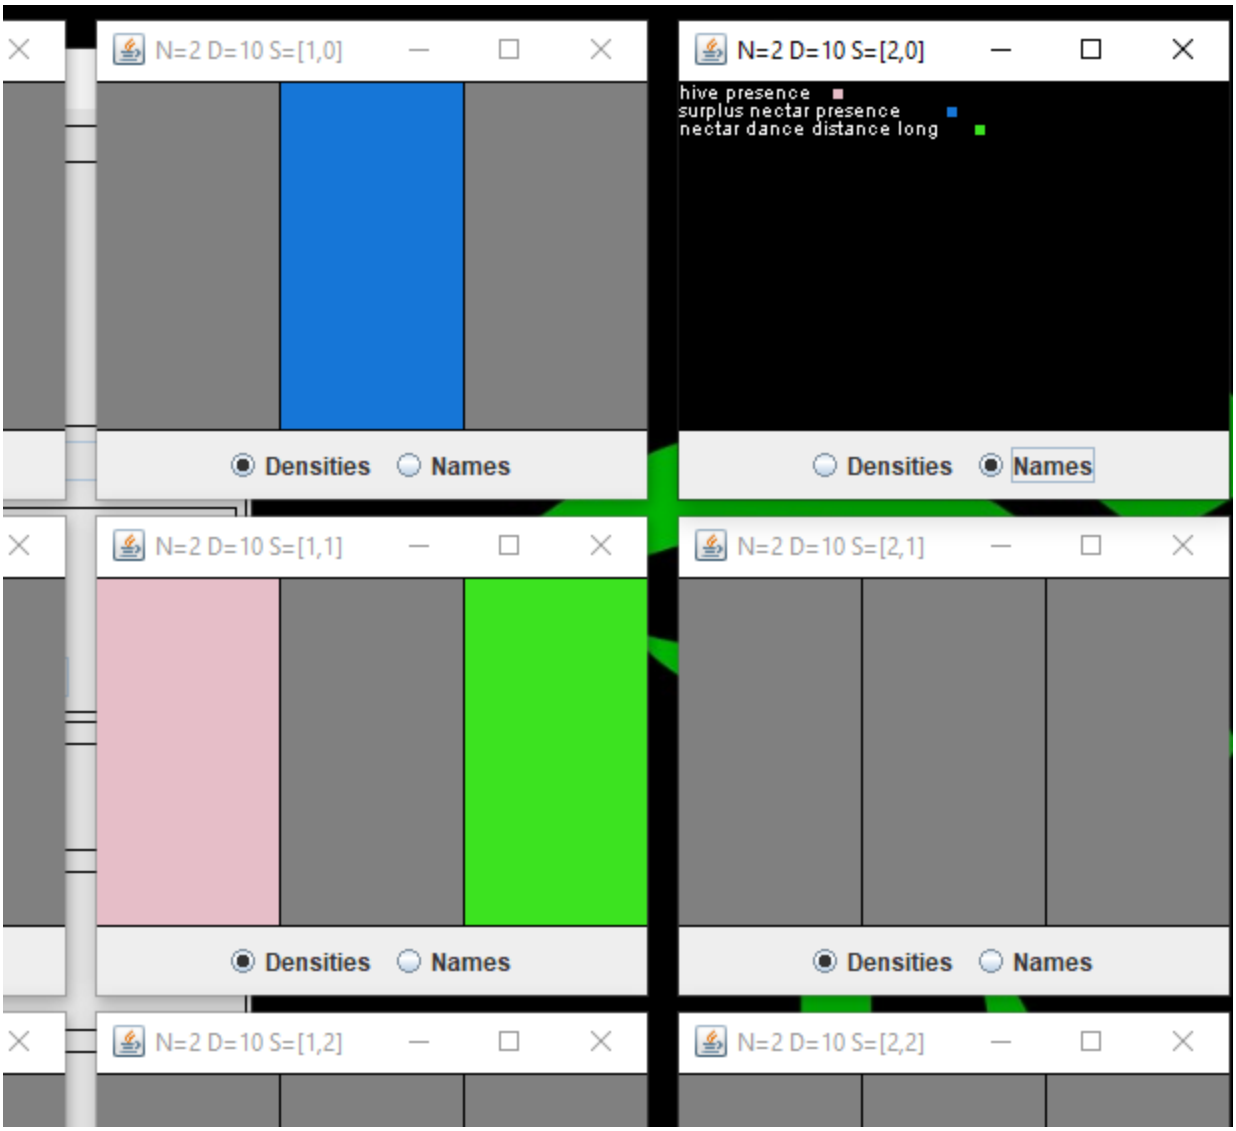

Figure 13b – Morphognostic neighborhood 2. At the center sector [1, 1] the *hive presence* and *nectar dance distance long* events are recorded. The location of the surplus nectar is recorded in sector [1, 0] and was used to orient toward the surplus nectar as part of the dance.

Code

The Java code is available on GitHub: https://github.com/morphognosis/HoneyBees

Results

Artificial neural network

The artificial neural network used was the MultiLayerPerceptron class in the Weka 3.8.3 machine learning library (https://www.cs.waikato.ac.nz/ml/weka/).

These parameters were used:

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

- learning rate = 0.1
- momentum = 0.2
- training epochs = 5000

The morphognostic configured as previously described, four 3x3 neighborhoods, produces 234 binary inputs to the network. There are 14 outputs representing the honey bee responses.

Base level testing

Neither a randomly generated responses nor an untrained network resulted in any nectar collected over 20,000 steps in a 3 flower and 3 bee configuration.

Test flower and bee quantities

In order to determine how the system scales up, three variations of flowers and bees were tested: 3 flowers and bees, 5 flowers and bees, and 7 flowers and bees. The amount of nectar collected was used as a success metric.

The world was set at 21x21 cells, and the hive at radius 3. The network was configured with 50 hidden neurons. Running the world for 20,000 steps on autopilot generated a metamorph dataset to train the neural network on. Datasets were generated for 10 trials.

Table 2 shows the average training dataset size and training accuracy. Of note is the increase in the number of metamorphs as the world become more complex with additional flowers.

Flowers/Bees	Metamorphs	Accuracy
3	972.4	99%
5	2271.7	99%
7	3852.3	99%

Table 2 – Number of metamorphs and training accuracy by varying flower and bee quantities.

Figure 14 shows the results of running programmatically (Autopilot) vs. with the trained network (Morphognosis). The network performs comparably.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

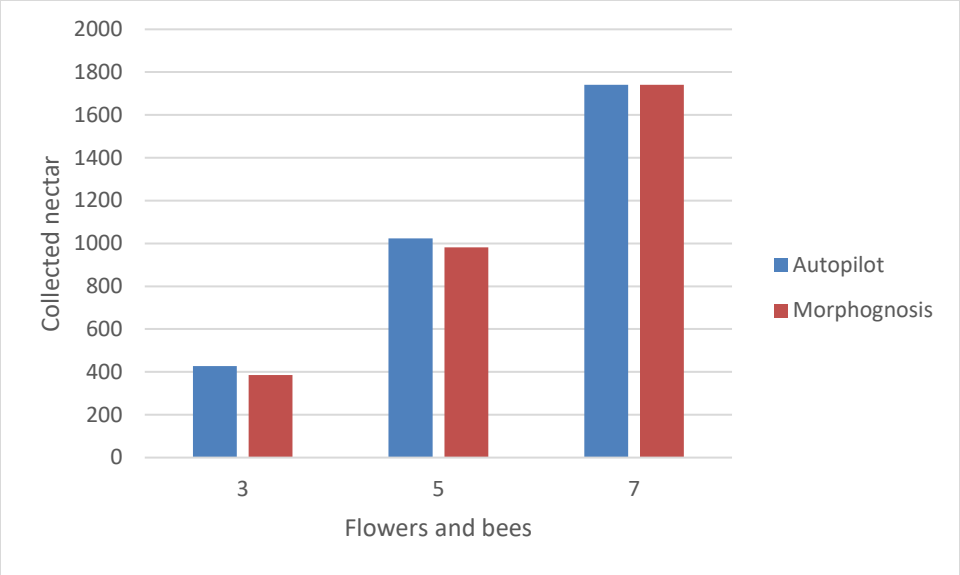

Figure 14 – Collected nectar for variations of flowers/bees.

TEST HIDDEN NEURONS

In order to observe how the system is affected by the neural network size, three variation of hidden neuron quantities were tested: 25, 50, and 100.

Table 3 shows the average training dataset size and training accuracy.

Hidden neurons	Metamorphs	Accuracy
25	913.9	99%
50	972.4	99%
100	887.4	99%

Table 3 – Number of metamorphs and training accuracy by varying hidden neurons.

Figure 15 shows the results, indicating that fewer hidden neurons are sufficient to achieve comparable performance.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

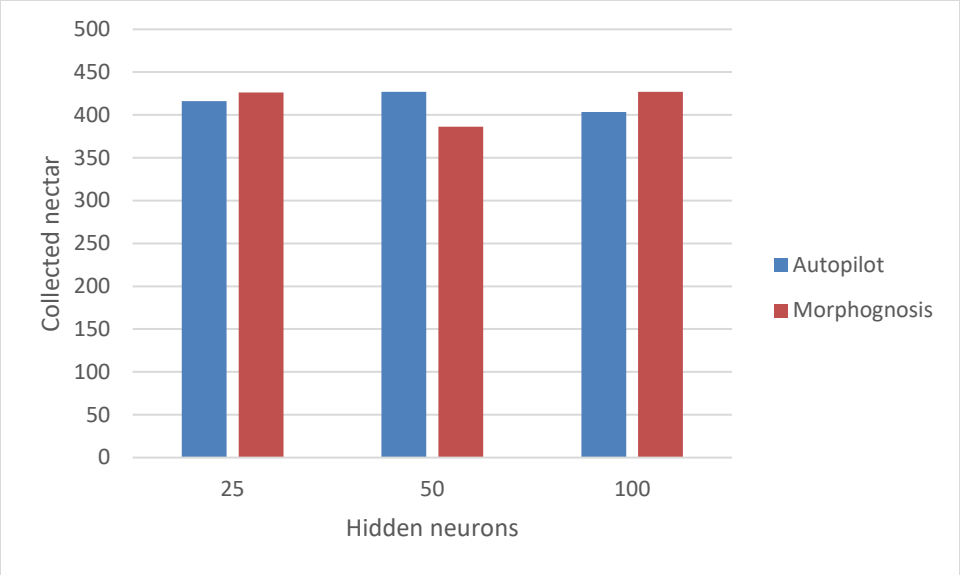

Figure 15 – Collected nectar for variations of hidden neurons.

TEST HIVE RADIUS

In order to observe how the system is affected by the hive size, two variation of hive sizes were tested: radii of 2 and 3.

Table 4 shows the average training dataset size and training accuracy. Of note is the reduction in metamorphs with a smaller hive. This is likely due to fewer "trajectories" to and from the hive.

Hive radius	Metamorphs	Accuracy
2	622.4	99%
3	972.4	99%

Table 4 – Number of metamorphs and training accuracy by varying hive radius.

Figure 16 shows the results, indicating that a smaller hive reduces the amount of nectar collected. A possible contributing factor for this is congestion due to bee collisions.

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

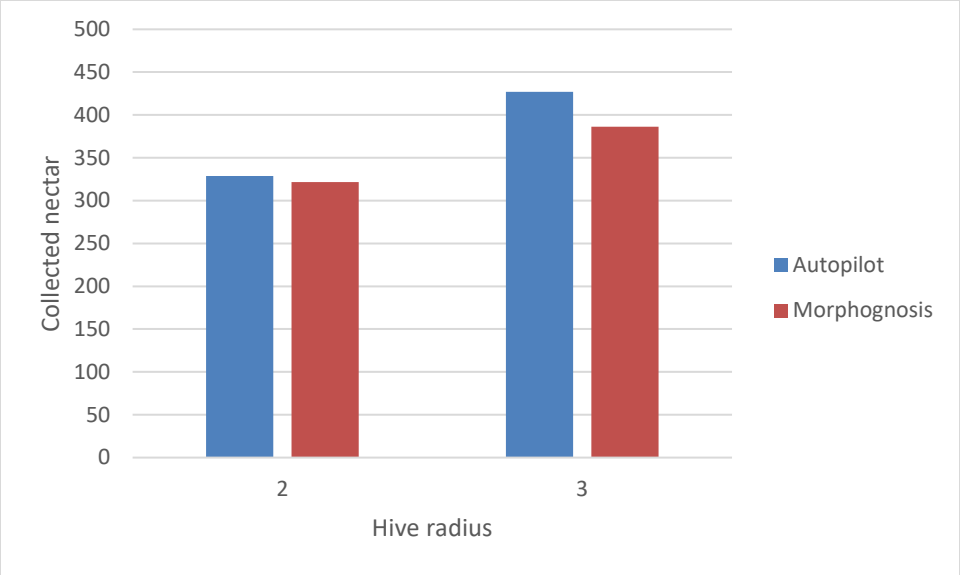

Figure 16 – Collected nectar for variations of hive radius.

LSTM COMPARISON

A key ability of a honey bee is to be able to track the location of the hive as it forages. This allows it to return to the hive with nectar. As a check of the ability of an unaugmented recurrent neural network (RNN) to perform this task, a Long Short Term Memory (LSTM) recurrent network (Hochreiter and Schmidhuber, 1997) was trained given sequences between 5 and 15 steps consisting of random orientation changes and forward movements probabilistically identical to those used by the honey bees. Despite variations in the network capacity, the training accuracy averaged approximately 30%, which was about the same as a random guess.

The LSTM network used is in the JANN 0.10 machine learning library:
https://github.com/JANNLab/JANNLab

CONCLUSION

The brain, a complex structure resulting from millions of years of evolution, can be viewed as a solution to problems posed by an environment existing in space and time. Internal spatial and temporal representations allow an organism to navigate and manipulate the environment. Following nature's lead, Morphognosis comprises an artificial neural network enhanced with a framework for organizing sensory events into hierarchical spatial and temporal contexts.

The successful simulation of honey bee foraging behavior suggests future projects are worth undertaking:

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

- The metamorph structure bears a close resemblance to deep reinforcement learning training elements (Francois-Lavet et al., 2018), suggesting the possibility of applying such learning to implement goal-seeking behavior.
- The aggregation scheme that supports scalability is a simple histogram-like method for dimensionality reduction.
    - The use of ANN dimensionality reduction techniques, such as autoencoding, might scale with higher information content.
    - The value of each neighborhood sector essentially represents a single centroid of sensory event values that have occurred in its space-time cube. An extension of this would be to retain multiple centroids within a sector, possibly weighted by frequency, increasing in number for higher level neighborhoods which encompass greater extents of space-time. This might increase the richness of behavioral variability while limiting information overload.
- The model is currently implemented in a cellular automaton spatial grid of cells. However, it is not inherently tethered to this platform and in fact may benefit from extending beyond it.
- The configuration of the morphognostic is vital to successful performance. For the honey bee task, this was a manual design. This process should be amenable to optimization/evolution methods.

REFERENCES

Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*. doi:10.1126/science.aat6766

Betti, M., LeClair, J., Wahl, L. M., Zamir, M. (2017). Bee++: An Object-Oriented, Agent-Based Simulator for Honey Bee Colonies. *Insects*, *8*(1), 31; https://doi.org/10.3390/insects8010031

Braitenberg, V., Schüz, A. (1998). Statistics and Geometry of Neuronal Connectivity, Second Edition (Berlin: Springer-Verlag).

Bruner, I. K., Moscovitch, M., Barense, M. D. (2018). Boundaries Shape Cognitive Representations of Spaces and Events. *Trends in Cognitive Sciences*. Volume 22, Issue 7, P637-650. https://doi.org/10.1016/j.tics.2018.03.013

Chittka, L., Wilson, C. (2018). Bee-brained. Are insects 'philosophical zombies' with no inner life? Close attention to their behaviours and moods suggests otherwise. *Aeon*. https://aeon.co/essays/inside-the-mind-of-a-bee-is-a-hive-of-sensory-activity

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Collins, L. (2019). The case for emulating insect brains using anatomical "wiring diagrams" equipped with biophysical models of neuronal activity. *Biological Cybernetics* https://doi.org/10.1007/s00422-019-00810-z

Cope, A. J., Richmond, P., Marshall, J., Allerton, D. (2013). Creating and simulating neural networks in the honeybee brain using a graphical toolchain. http://greenbrain.group.shef.ac.uk/wp-content/uploads/2013/11/SFN_2013_GB.pdf

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, Volume 14, Issue 2, Pages 179-211.

Fernando, S., Kumarasinghe, N. (2015). Modeling a Honeybee using Spiking Neural Network to Simulate Nectar Reporting Behavior. *International Journal of Computer Applications* 130(8):32-39. DOI: 10.5120/ijca2015907078.

Fox, A. (2019). Bees 'get' addition and subtraction, new study suggests. *Science Magazine*. https://www.sciencemag.org/news/2019/02/bees-get-addition-and-subtraction-new-study-suggests

Francois-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., Pineau, J. (2018). An Introduction to Deep Reinforcement Learning. https://arxiv.org/abs/1811.12560

Hainmüller, T., Bartos, M. (2018): Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature*. doi: 10.1038/s41586-018-0191-2

Hochreiter, S., Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.

Karaboga, D., Akay, B. (2009). A survey: algorithms simulating bee swarm intelligence. *Artificial Intelligence Review* volume 31, Article number: 61.

Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*. https://doi.org/10.1016/j.neuron.2016.12.041

Lieff, J. (2015). Time Cells Organize Memory. http://jonlieffmd.com/blog/time-cells-organize-memory

MaBouDi, H., Shimazaki, H., Giurfa, M., Chittka, L. (2017). Olfactory learning without the mushroom bodies: Spiking neural network models of the honeybee lateral antennal lobe tract reveal its capacities in odour memory tasks of varied complexities. *PLOS Computational Biology*; 13 (6): e1005551 DOI: 10.1371/journal.pcbi.1005551

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Moravec, H. (1988). Mind Children: The Future of Robot and Human Intelligence. (Harvard University Press).

Nosowitz, D. (2016). I Asked Leading Entomologists: 'What's The Smartest Bug In The World?' *Atlas Obscura*. https://getpocket.com/explore/item/i-asked-leading-entomologists-what-s-the-smartest-bug-in-the-world?utm_source=pocket-newtab

Nott, G. (2018). How a brain the size of a sesame seed could change AI forever. *ComputerWorld*. https://www.computerworld.com.au/article/647401/how-brain-size-sesame-seed-could-change-ai-forever/

Oldroyd, B. P. (2007). *"What's Killing American Honey Bees?". PLoS Biology. 5 (6): e168. doi:10.1371/journal.pbio.0050168. PMC 1892840. PMID 17564497*

Portegys, T. (1986). GIL - An Experiment in Goal-Directed Inductive Learning. Ph.D. dissertation, Northwestern University, Evanston, Illinois, 1986.

Portegys, T. (2007). Learning Environmental Contexts in a Goal-Seeking Neural Network. *Journal of Intelligent Systems*, Vol. 16, No. 2.

Portegys, T. (2010). A Maze Learning Comparison of Elman, Long Short-Term Memory, and Mona Neural Networks. *Neural Networks*.

Portegys, T. (2013). Discrimination Learning Guided By Instinct. *International Journal of Hybrid Intelligent Systems*, 10, 129–136.

Portegys, T. (2015). Training Artificial Neural Networks to Learn a Nondeterministic Game. *ICAI'15: The 2015 International Conference on Artificial Intelligence*.

Portegys, T., Pascualy, G., Gordon, R., McGrew, S., Alicea, B., (2017). Morphozoic: cellular automata with nested neighborhoods as a metamorphic representation of morphogenesis. In Multi-Agent Based Simulations Applied to Biological and Environmental Systems, ISBN: 978-1-5225-1756-6.

Portegys, T. (2017). Morphognosis: the shape of knowledge in space and time. *The 28th Modern Artificial Intelligence and Cognitive Science Conference (MAICS)*, Fort Wayne Indiana, USA.

Portegys, T. (2018). Learning C. elegans locomotion and foraging with a hierarchical space-time cellular automaton. *Neuroinformatics 2018 Montreal*. *F1000Research* 2018, **7**:1192 (doi: 10.7490/f1000research.1115884.1)

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Portegys, T. (2019). Generating an artificial nest building pufferfish in a cellular automaton through behavior decomposition. *International Journal of Artificial Intelligence and Machine Learning (IJAIML)* 9(1) DOI: 10.4018/IJAIML.2019010101.

Roper, M., Fernando, C., Chittka, L. (2017). Insect Bio-inspired Neural Network Provides New Evidence on How Simple Feature Detectors Can Enable Complex Visual Generalization and Stimulus Location Invariance in the Miniature Brain of Honeybees. *PLoS Comput Biol*. Feb; 13(2): e1005333. doi: 10.1371/journal.pcbi.1005333

Schöneburg, E. (2019). Alternative AI (AAI) – An alternative path to AGI. Keynote: Artificial Life. https://www.youtube.com/watch?v=OeZM1y-AK5U&feature=share&fbclid=IwAR3D_WsLIstB0VaYmPLNzPiuOSE0HVKzXd9GDS3jzHQtFwKIfvRlA8OWozM

Schürch, R., Zwirner, K., Yambrick, B., Pirault, T., Wilson, J. M., Couvillon, M. J. (2019). Dismantling Babel: creation of a universal calibration for honey bee waggle dance decoding, *Animal Behaviour*. DOI: 10.1016/j.anbehav.2019.01.016

Tavares, R. M., Mendelsohn, A., Grossman, Y., Williams, C. H., Shapiro, M., Trope, Y., Schiller, D. (2015). A Map for Social Navigation in the Human Brain. *Neuron*. Volume 87, Issue 1, P231-243. DOI:https://doi.org/10.1016/j.neuron.2015.06.011

Toffoli, T., Margolus, N. (1987). Cellular Automata Machines: A New Environment for Modeling. MIT Press. p. 27. ISBN 9780262200608.

Von Frisch, K. (1967). The Dance Language and Orientation of Bees. *Harvard University Press*. ISBN 9780674418776.

Vorhees, C. V., and Williams, M. T. (2014). Assessing Spatial Learning and Memory in Rodents. *ILAR Journal*. 55(2), 310–332. http://doi.org/10.1093/ilar/ilu013.

Wade, N. (2001). Genome's Riddle: Few Genes, Much Complexity. *The New York Times*. https://www.nytimes.com/2001/02/13/health/genomes-riddle-few-genes-much-complexity.html

Wolfram, S. (2002). A New Kind of Science. Wolfram Media. ISBN-10: 1579550088.

Wood, W. B. editor. (1988). The Nematode Caenorhabditis elegans. *Cold Spring Harbor Monograph Series*. ISBN 978-087969433-3.

Yong, E. (2019). The Human Brain Project Hasn't Lived Up to Its Promise. *The Atlantic.* https://www.theatlantic.com/science/archive/2019/07/ten-years-human-brain-project-simulation-markram-ted-talk/594493/

bioRxiv preprint doi: https://doi.org/10.1101/2020.03.14.992263; this version posted March 17, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Yong, E. (2017). How Brain Scientists Forgot That Brains Have Owners. *The Atlantic*. https://getpocket.com/explore/item/how-brain-scientists-forgot-that-brains-have-owners?utm_source=pocket-newtab

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