Neuroevolution of a Recurrent Neural Network for Spatial and Working Memory in a Simulated Robotic Environment

Xinyun Zou
University of California, Irvine
xinyunz5@uci.edu

Eric Scott
George Mason University

Alexander Johnson
University of California, San Diego

Kexin Chen
University of California, Irvine

Douglas Nitz
University of California, San Diego

Kenneth De Jong
George Mason University

Jeffrey Krichmar
University of California, Irvine

ABSTRACT

We evolved weights in a recurrent neural network (RNN) to replicate the behavior and neural activity observed in rats during a spatial and working memory task. The rat was simulated using a robot simulator to navigate a virtual maze. After evolving weights from sensory inputs to the RNN, within the RNN, and from the RNN to the robot's motors, the robot successfully navigated the space to reach four reward arms with minimal repeats before the timeout. Our current findings suggest that it is the RNN dynamics that are key to performance, and that performance is not dependent on any one sensory type, which suggests that neurons in the RNN are performing mixed selectivity and conjunctive coding. The RNN activity resembles spatial information and trajectory-dependent coding observed in the hippocampus. The evolved RNN exhibits navigation skills, spatial memory, and working memory.

CCS CONCEPTS

• **Networks** → **Network algorithms**; • **Computing methodologies** → **Modeling and simulation**;

KEYWORDS

evolutionary robotics, recurrent neural networks

ACM Reference Format:
Xinyun Zou, Eric Scott, Alexander Johnson, Kexin Chen, Douglas Nitz, Kenneth De Jong, and Jeffrey Krichmar. 2021. Neuroevolution of a Recurrent Neural Network for Spatial and Working Memory in a Simulated Robotic Environment. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3459565

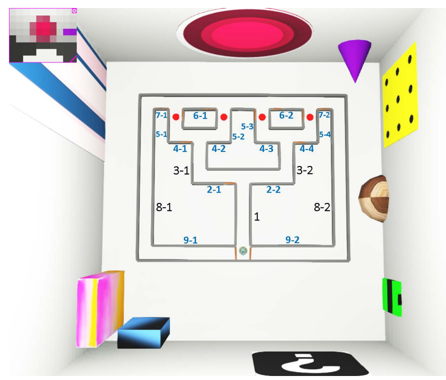

Figure 1: The maze visualization in Webots.

1 INTRODUCTION

We used the Webots robot simulation environment [1] to investigate cognitive map behavior observed in rats during a spatial and working memory task, known as the triple T-maze [2]. In this task, the rat or the robot must take one of four paths to receive a reward. If it repeats a path, there is no additional reward. It should eventually learn to quickly reach each of the four rewards with minimal repeats. This requires knowledge of where it is now, where it has been, and where it should go next.

Our results show that the evolved RNN was capable of guiding the robot through the triple-T maze with similar behavior to that observed in the rat. Our analysis of the RNN activity indicates that the behavior was not dependent on any one sensory projection type but rather relied on the evolved RNN dynamics. Furthermore, the population of neurons in the RNN were not only sufficient to predict the robot's current location but also carried a predictive code of future intended reward paths. The present method for evolving neural networks for robot controllers may also be applicable to other memory tasks. More details describing the model and results are available in Zou et al. [4] .

2 METHODS

Figure 1 shows the maze simulation environment. The red circles, which denote the location of the rewards, were not observable by the robot and are only included in the figure for illustrative purposes. The agent was an e-puck robot which has an accelerometer, a front camera, 8-direction proximity sensors, several LEDs, and 2 wheel motors. The e-puck needed to learn by neuroevolution to find four rewards (and return home after each reward visit) with minimal repeats before the timeout.

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07...$15.00
https://doi.org/10.1145/3449726.3459565

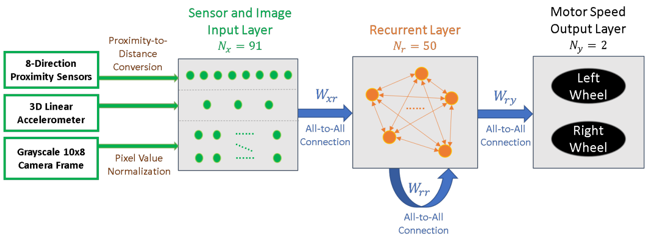

Figure 2: Neural network architecture for controlling the e-puck robot in Webots.

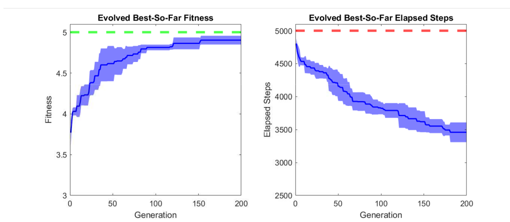

Figure 3: Evolved performance for the best-so-far agent.

Table 1: Ablation performance (mean ± 95% CI).

	Fitness	Elapsed Steps
No Ablation	3.65±0.14	4644±169
Proximity Sensors	3.29±0.16	4774±128
Linear accelerometer	3.49±0.16	4640±131
Grayscale Vision	3.21±0.17	4785±67
Input Weights	2.93±0.085	**4981±16∗**
Recurrent Weights	**2.95±0.11∗**	**4946±43∗**
Output Weights	**2.84±0.10∗**	**4978±22∗**

The neural network architecture received inputs from the e-puck's 8-direction proximity sensors, 3D linear accelerometer values, and normalized pixel values from its 10×8 grayscale camera frame (Figure 2). These 91 input neurons were fully connected to 50 recurrent neurons, which were fully connected with one another. This recurrent layer was then fully connected with the two neurons in the output layer that controlled the rotational speed of the two wheel motors separately.

3 RESULTS

Figure 3 shows the best-so-far evolutionary performance and the number of elapsed steps for five runs. Each run lasted 200 generations, with 50 genotypes per generation. In each generation, the fitness value of a genotype was averaged over 5 trials.

We carried out a set of ablation simulations to test whether performance was dependent on any sensory projection type or just evolved weights in the neural network (see Table 1). We either randomly shuffled different sensor input values or shuffled the RNN input (W_{xr},) recurrent (W_{rr}), or output (W_{ry}) weights. Interestingly, none of the sensory projection ablations had a significant impact on performance. However, ablating the evolved weights (input, recurrent, and output) all had a significant impact. This suggests that it was the RNN dynamics that were key to performance.

Table 2: Prospective path prediction for segments.

	Seg1	Seg3-1	Seg3-2	Seg8-1	Seg8-2
correctness	41%	77%	69%	47%	55%
bins off	0.9	0.2	0.2	2	2

Borrowing techniques from neuroscience [2, 3], we further tested whether the RNN contained spatial information with a population code. The location prediction for each bin in all 25 test trials shows that the RNN activity was sufficient to predict the robot's position in the maze. The robot's position was predicted with perfect accuracy on 58% of the bins, and the predicted error had an average distance of 3.1 bins (i.e., 0.25 meters).

Solving the triple-T maze task requires the agent to remember which path it has already taken and which path to take next. Table 2 shows how well the RNN activity could predict the robot's path. The probability of correct path prediction on Segment 1, where the robot could take one of 4 paths, was well above chance level (t-test; $p < 0.0001$). The correctness on Segment 3-1 or 3-2, where the robot could take one of two paths, was also well above chance level (t-test; $p < 0.0001$ for Segment 3-1 and $p < 0.005$ for Segment 3-2). This suggests that the RNN carried a prospective code of where the robot intended to go next. The probabilities of correct path prediction on Segments 8-1 and 8-2 were not significant. These results suggest that the evolved RNN had prospective information of whether the robot intended to turn left or right, but we could not yet observe retrospective information of where it had already visited.

4 CONCLUSIONS

We introduced a recurrent neural network (RNN) model that linked the robot sensor values to its motor speed output. The evolved network architecture achieved the goal of successfully performing a cognitive task that required spatial and working memory. The RNN population carried spatial information sufficient to localize robot in the triple T-maze. It also carried predictive information of which path robot intended on taking. Robot behavior was dependent on RNN dynamics rather than a sensor-to-motor mapping. Our method shows that complex robot behavior, similar to which being observed in animal models, can be evolved and realized in RNNs.

ACKNOWLEDGMENTS

This material is based upon work supported by Air Force Office of Scientific Research (AFOSR) under Contract No. FA9550-19-1-0306.

REFERENCES

[1] Olivier Michel. 2004. Cyberbotics Ltd. Webots™: professional mobile robot simulation. *International Journal of Advanced Robotic Systems* 1, 1 (2004), 39–42.
[2] Jacob M. Olson, Kanyanat Tongprasearth, and Douglas A. Nitz. 2017. Subiculum neurons map the current axis of travel. *Nature Neuroscience* 20, 2 (2017), 170–172.
[3] MA Wilson and BL McNaughton. 1993. Dynamics of the hippocampal ensemble code for space. *Science* 261, 5124 (1993), 1055–1058.
[4] Xinyun Zou, Eric O. Scott, Alexander B. Johnson, Kexin Chen, Douglas A. Nitz, Kenneth A. De Jong, and Jeffrey L. Krichmar. 2021. Neuroevolution of a Recurrent Neural Network for Spatial and Working Memory in a Simulated Robotic Environment. (2021). arXiv:cs.NE/2102.12638