

Integrating Spiking Neural Networks and Deep Learning Algorithms on the Neurorobotics Platform

Rachael Stentiford¹[0000-0002-8530-7209], Thomas C. Knowles¹[0000-0002-7750-2119], Benedikt Feldotto²[0000-0002-6070-913X], Deniz Ergene²[0000-0003-2520-7940], Fabrice O. Morin²[0000-0003-0185-7420], and Martin J. Pearson¹[0000-0002-8642-4845]

¹ Bristol Robotics Laboratory, University of the West of England, Bristol, UK
{rachael.stentiford,tom.knowles,martin.pearson}@uwe.ac.uk

² Department of Informatics, Technical University of Munich, 85748 Munich, Germany
{morinf, feldotto, ergene}@in.tum.de

Abstract. We present a neurobotic model that can associate self motion (odometry) with vision to correct for drift in a spiking neural network model of head direction based closely on known rodent neurophysiology. We use a deep predictive coding network to learn the generative model of representations of head direction from the spiking neural network to views of naturalistic scenery from a simulated mobile robot. This model has been deployed onto the Neurorobotics Platform of the Human Brain Project which allows full closed loop experiments with spiking neural network models simulated using NEST, a biomimetic robot platform called WhiskEye in Gazebo robot simulator, and a Deep Predictive Coding network implemented in Tensorflow.

Keywords: Neurorobotics Platform · Predictive coding · Spiking Neural Network · pyNEST · NRP · WhiskEye · Head Direction

1 Introduction

Neurorobotics is a discipline that works toward building embodied, bio-plausible models of animal neurology often with a view to improving our understanding of living neural systems. Modelling approaches include, but are not limited to, machine learning [13], bio- and neuromorphic hardware [9], functional sub-networks [18], layered control architectures [14] and spiking neural networks [1]. It is particularly well suited to addressing neuroscience questions for which conventional approaches are not ethically viable thus directly supporting one of the pillars of the NC3Rs: Replacement, by using advanced tools to address neuroscience questions without the use of animals [12]. Neurorobotics also allows us to perform experiments that alleviate the technical difficulties of animal behavioural experiments through the embodiment of models and their closed-loop interaction with either physical or real environmental stimuli.

The Neurorobotics Platform is a tool developed as part of the Human Brain Project for conducting embodied robotics experiments with embedded bio-inspired brain and control systems [4]. It provides synchronisation between Spiking Neural Network (SNN) simulators such as NEST, robot control and simulation tools such as ROS and Gazebo, as well as other popular libraries. Robot behaviours can be specified using the ROS framework and all run time data can be captured and exported after the experiment for analysis.

In this article we present the use of the Neurorobotics Platform (NRP) to integrate and coordinate an online model of the rodent head direction system. We model the rodent head direction cell system as an SNN which estimates the current head angle of a simulated robot based on self-motion (ideothetic) cues, and provides environmental (allothetic) information to the network using a Predictive Coding Network (PCN) implemented in Tensorflow. The model was developed offline in prior works [17] to explore the bidirectional learning that exists between self-motion cues serving as a scaffold for initial learning of change in pose, followed by corrective input from associated allocentrically anchored visual cues after learning. We had found that both discriminative Convolutional Neural Networks (CNN) and generative PCN approaches were appropriate for learning this association but that PCNs were more robust in applying corrective input to the SNN when visual allothetic cues were unreliable [17]. We explore this further in this article by correlating the prediction error generated by the PCN at its lowest layer as a proxy for the reliability of allothetic cues. We contend that this provides a biologically plausible signal which an animal may use to delineate trust in ideothetic or allothetic sensory cues for the update and representation of its estimate of head direction. The experiments presented to support this contention were made possible only through the closed loop integration provided by the NRP.

1.1 Background

Although the rodent head direction system works in the absence of vision, relying on self-motion (ideothetic) information to track head direction, the signal is subject to accumulated error (drift) [7, 8, 15, 19, 21, 22]. Rodents use allothetic information such as vision, to counter this drift in head direction estimate [19]. This requires forming learned associations between visual scenes and the current head angle, so that the estimated head angle can be corrected when this visual scene is experienced again. This visual control of head direction begins at the Lateral Mammillary Nuclei (LMN) [20], stabilising the head direction signal at its origin. The head direction is thought to originate in the reciprocal connections between the LMN and dorsal tegmental nuclei (DTN) [2, 3]. We model the head direction system as a continuous attractor network, exploiting the excitatory and inhibitory connections between these two regions.

To learn associations between heading and the visual scene we employ a generative paradigm that aims to learn to generate appropriate data samples like the training data in the appropriate contexts. The Predictive Coding Network (PCN) based on MultiPredNet [13], was re-purposed to receive head direction

and vision information rather than vision and tactile. The PCN works by outputting a prediction from its latent layer that passes through the nodes of the hidden layers (if any) to the input later of each modality. At each layer, the prediction from the layer above is compared to the activity at the current layer and the difference (error) calculated. Weights between layers are then updated locally according to their prediction errors. For a full description of both the PCN and the SNN see [17].

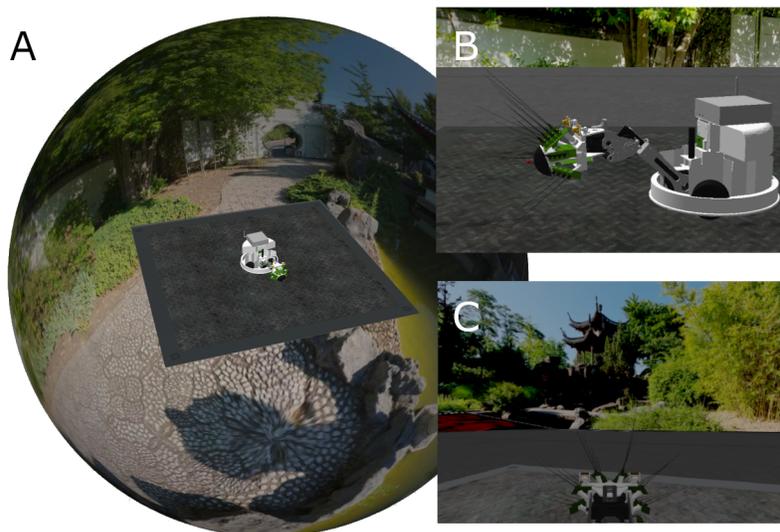


Fig. 1. Illustration showing the Chinese garden environment projected onto a sphere. (A) WhiskEye robot on a floating platform. (B) Image of the WhiskEye robot in the environment. (C) Image from a camera mounted on the WhiskEye robot.

2 Method

2.1 Neurorobotics Platform

We use a local copy of the NRP version 3.2.0. The NRP is a complete neuro-robotics experiment platform that builds upon many *de facto* standard open source software for robotics and physics simulation. Specific components used in this paper include ROS (Noetic)[16], Gazebo (11.3)[11] and NEST (2.18)[6]. Gazebo and its integrated ODE physics engine supports the use of robot model files describing joints, actuators and sensors, which can be imported alongside meshes and textures to simulate moving, sensing, 3D models of robots with rigid

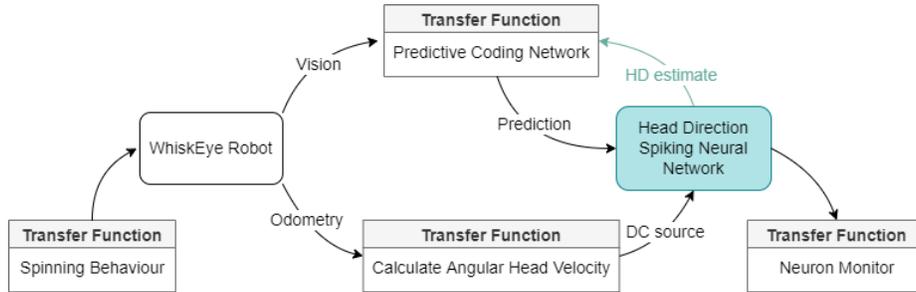


Fig. 2. Diagram showing the structure of the experiment in the NRP. The simulated WhiskEye robot turns on the spot as described by a transfer function. The pose of the robot is sent to a transfer function which calculates angular head velocity and converts it to current inputs to the conjunctive cells of the spiking neural network (SNN). The camera feed from the robot is passed to the trained Predictive Coding Network (PCN) in a transfer function and used to produce predictions which are converted to current input to each of the Head Direction cells in the SNN. The head direction estimate from the SNN, taken as the average of the active cells since the last prediction, is also passed to the PCN to measure prediction error.

body interactions inside a 3D virtual world. This often involves creating a virtual simulacrum of a real robot, such as the WhiskEye [10] robot, originally built physically at the BRL and later built in simulation for the NRP. This is then extended further by linking Gazebo to NEST and ensuring timely delivery of messages between them (such as sensory data from the virtual environment, or spiking output from NEST) via the Closed-Loop Engine (CLE). The CLE is the heart of the NRP, providing a framework for all supported tools to communicate to each other using ROS Messages and transfer functions in the form of Python scripts. Such congruence between real and simulated robots allows for experimental results to be compared and serves as a valid substitute for working with physical robots. This is all assisted by a web portal GUI that allows for the environment to be altered, robots to be added on the fly and transfer functions to be created, enabled, disabled and deleted as needed. A summary of NRP-specific terms are summarised in Table 1.

Integrating Tensorflow To avoid compatibility issues, Tensorflow 2.3 was installed into a separate Python virtual environment and imported via the Brain File. Though intended to set up NEST models, the Brain File also blocks simulation setup whilst it is running, making ideal for any heavy computations that need to be run once. Tensorflow code included defining the Predictive Coding model itself, creating a new session, and loading trained weights from file. Access to the Tensorflow model within transfer functions was achieved by assigning the model to a Brain File variable, which stay in memory for the duration of the simulation. Inference can then be done 'on the fly' as the robot explores the en-

Term	Meaning
Experiment	A collection of robot, brain and asset files for a particular simulation. This is described by an EXC file, which holds the settings of the simulation.
CLE	The Closed Loop Engine. This orchestrates the various components of the NRP with a unifying simulation clock, and enables communication between them via transfer functions.
Transfer Function	A Python script that is run as part of the CLE’s loop. They are decorated with @Robot2Neuron or @Neuron2Robot accordingly.
Brain File	A Python file that describes the NEST model to be loaded.

Table 1. List of terms used within the NRP

vironment. This allows us to close the loop, producing head direction predictions from virtual camera images as the simulation runs, and passing these predictions as current input to the brain model to track and correct drift in the estimated head angle in real time. The PCN can also receive live input from the NEST model, with each tick of the CLE retrieving the activity of the NEST devices that form its own ideothetic prediction of head direction. This provides both a target for causal inference for the PCN and allows error between the PCN and NEST estimates to be calculated, both which would be impossible without the synchronised execution of Tensorflow and NEST that the NRP supports. The flow diagram shown in figure 2 describes the interplay between major components of the experiment.

2.2 WhiskEye

WhiskEye in the NRP is a simulated version of a rat-inspired omnidrive robot with RGB cameras in place of eyes and active whisker-like tactile sensors reported previously (see Figure 1). WhiskEye was integrated into the Human Brain Project’s Neurorobotics Platform as part of prior work [10, 13]. ROS topics including body angle, neck position, camera feeds published related to the robot are subscribed to in transfer functions and these data passed to the PCN or SNN, see below. Camera feeds must be deserialised and converted into the flattened RGB format that the PCN expects; this is accomplished by the PCN Inference transfer function.

2.3 Environment

Within the NRP we situate a sphere mesh onto which a background image is projected, enclosing a central suspended platform with grey stone texture for WhiskEye to move on. The sphere is large enough that translation within the bounds of the platform lead to no perceptible changes to the visual scene. This

ensures the robot nominally observes the same scene when it returns to that head angle, analogous to rodents observing distal environmental cues.

2.4 Brain Model

The brain model is a pyNEST (2.18) SNN. Using a network structure supported by experimental observations, we define four populations each of 180 standard leaky integrate-and-fire neurons (`iaf_psc_alpha`) which use alpha-function shaped synaptic currents. These populations represent head direction cells in the Lateral Mammillary nuclei (LMN) and Dorsal Tegmental nuclei (DTN); and two conjunctive head direction by asymmetric angular head velocity populations [2]. LNM cells have constant current input of 450 pA that maintains spontaneous firing at a rate of 50 spikes per second prior to inhibitory input from the DTN.

Reciprocal connections between the LMN and DTN have been shown to be essential for generating the head direction signal [3]. A single stationary bump of activity is produced by providing inhibitory input to all cells surrounding the most active LMN cell. Excitatory connections from LMN neurons to the DTN with declining synaptic strength as a function of distance, inhibition is returned from DTN cells with synaptic strength decreasing as a function of distance offset by a constant (μ).

Conjunctive cells are connected one to one with LMN cells offset by one cell either clockwise or anticlockwise. Conjunctive cells require both AHV and HD input to fire, and shift the bump around the ring. AHV input is provided by a transfer function. For further description of the spiking neural network structure see [17].

2.5 Transfer Functions

In the NRP transfer functions are used to coordinate interactions between the various component simulators and libraries that make up an experiment. Table 2 summarises the transfer functions used in this study.

2.6 Predictive Coding Network

We use a multimodal predictive coding network, based on the previously proposed MultiPredNet [13], that attempts to reconstruct each pair of inputs - images of visual scenes and head direction - from a multimodal latent space. Prior work [13, 17] has shown this network’s effectiveness at predicting pose and head direction from natural scenes using its bio-plausible local learning rules. Having been trained in an offline setting on at least one full rotation of views from the same environment, these weights are loaded into the network during setup. This enables it to produce robust predictions of head direction based on the allothetic cues of the visual scene. These predictions are 180 elements in width at each update step, matching the number of HD cells. The prediction values are injected as current into the network via 180 `dc_source` devices connected one to one with the LMN population. To prevent negative current injections, any negative values in the prediction are set to 0.

Name	Function
AHV Input	The head velocity is cached using a Brain File variable, enabling Angular Head Velocity (AHV) to be calculated in each time step. Any changes in AHV are passed to the NEST model and injected into the conjunctive cell population.
Spinning	Head angle is varied continuously by sending z-axis angular velocity commands to WhiskEye.
PCN Inference	The PCN model is passed the current camera state as input, with influence from head direction input disabled. It then generates an inference of the current head direction based on its prediction of the camera state, which is then converted into a current value and passed to the NEST model.
Neuron Monitor	LMN head direction cell population spikes are collected using a <i>spike_recorder device</i> , and written to a CSV using a <i>CSVRecorder</i> . Recorded spikes are also displayed in the spike train window.

Table 2. List of transfer functions supporting the experiment

2.7 Spike Analysis

The most active cell in 40 ms bins is identified and assumed to be the current most active cell and the peak of the bump. Converting the cell number to a value between π and $-\pi$, these values are then compared to the ground truth head angle. The difference between the estimated and ground truth head angle indicates accumulation of drift over time, with total error measured as Root Mean Squared Error (RMSE).

3 Results

3.1 Head Angle Estimated by the SNN Follows Ground Truth

The bump of activity in the SNN is centred on the current estimate of head angle. Movement of the bump is driven by Direct Current (DC) input to the conjunctive layers which are connected one cell clockwise or anticlockwise around the ring. In the absence of corrective input from the PCN, the bump, driven by ideothetic input only, is subject to drift (Figure 3A blue), as the model for transferring AHV to current is not optimal. The difference between the ground truth head angle and the estimated head angle increases over time (Figure 3B blue), resulting in a total error of 2.42 radians (1.263 RMSE), after 3 minutes of simulation.

When the PCN transfer function is active, it produces predictions of the current head angle based on the current visual scene observed by the robot’s cameras. These predictions are converted to DC and injected into the head direction cells as corrective input. As seen in our previous work [17], this corrective

signal reduced the drift resulting in total error of 0.43 radians (82.23% reduction; 0.370 RMSE).

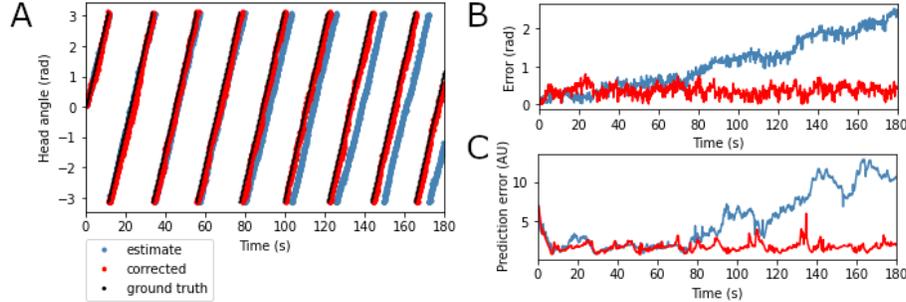


Fig. 3. Estimated head angle from the spiking neural network (SNN). (A) Plot showing estimated head angle with ideothetic drive only (blue), the corrected estimated head angle (red) which is updated using predictions from the predictive coding network (PCN), and ground truth head angle (black). (B) Error measured as the difference between the estimate and the ground truth is shown. Allothetic input from the PCN results in minimised drift and the corrected estimate and ground truth are almost indistinguishable. (C) Plot showing the prediction error at each time point when the head direction estimate is passed back to the PCN, when the predictions are used to correct for drift (red) or not (blue).

3.2 PCN Prediction Error Increases with Drift

Until this point we have measured the drift in the SNN by directly comparing the HD estimate to the ground truth. However, rodents do not have access to this ground truth information in order to evaluate the confidence of the HD estimate. As part of its inference process, the PCN produces a prediction error between the expected head angle and its reconstruction, which may be a suitable alternative. The NRP allows both the SNN and the PCN to run synchronously as the robot moves. This makes it possible to send continuous feedback between the two models. The predictions from the PCN can be used to update the SNN estimate, and in return the current HD estimate supplied to the PCN to be compared to the prediction generated based on the visual data. By passing the current head angle to the PCN, the prediction error between the presented head angle and the reconstruction is calculated. When the predictions are used to correct for error in the HD estimate, the prediction error remains low (Figure 3C red) with small peaks representing small inconsistencies that remain in the corrected head direction estimate. If we generate the predictions but do not feed these to the SNN, the head direction estimate drifts, resulting in an overall gradual increase in error (Figure 3C blue). The prediction error in the drifting case

correlates strongly with the calculated error (Figure 3B,C blue; 0.929 Pearson correlation coefficient). When the predictions are used to correct for error in the HD estimate, the prediction error reflects the reconstruction quality for each of the head angles (Figure 3C red). The oscillation in the prediction error indicates the prediction quality is not equal across all head angles.

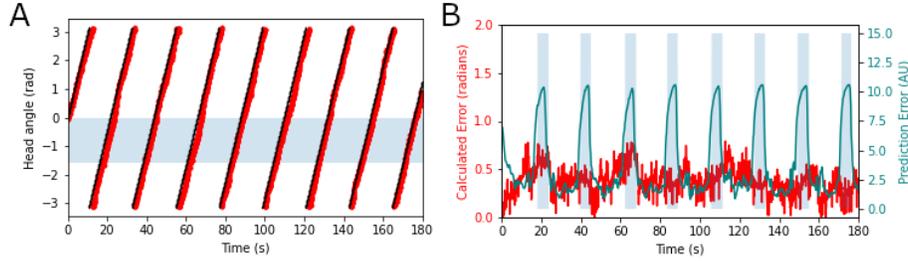


Fig. 4. Network response during periods of low visual information. (A) Plot showing the corrected estimated head angle (red) which is updated using predictions from the predictive coding network (PCN), and ground truth head angle (black). Visual information is obscured in the region shaded blue. (B) Plot showing the calculated error (red) between the ground truth and the estimate, and the prediction error from the PCN (red). Visual information is obscured in the regions shaded blue.

3.3 Ideothetic information drives network in periods of darkness

During periods of darkness or when distal visual information is not available because it is obscured by proximal objects, animals must rely upon ideothetic information to keep track of their head direction. To observe the response of the network to ambiguous visual information we obscured a 90 degree portion of the visual scene (Figure 4A blue shaded region). During this period the predictions produced by the PCN became close to a flat line, and very little current was injected into the network. Figure 4A shows ideothetic information driving the bump during the dark period with minimal drift. Figure 4B shows while the calculated error between the ground truth and the estimated head angle remains low (red), the prediction error (green) from the PCN shows peaks not visible in the calculated error which match up the the periods of darkness. This is a strong signal that the predictions are inaccurate, and in future experiments could be used to trigger learning after extended periods of poor prediction.

4 Discussion

This work demonstrates the use of using machine learning and spiking neural networks in a closed loop, embodied, and situated model on the Neurorobotics

platform (NRP). We have replicated the results of our previous work reducing drift in a Spiking Neural Network (SNN) model of the head direction system using head angle predictions inferred by a Predictive Coding Network (PCN). The NRP allowed us to close the loop, with the SNN and the PCN running synchronously as the robot moves, controlled by the Closed-Loop Engine of the NRP. Closing the loop has enabled new avenues for reciprocal information transfer between the PCN and the SNN to be explored that were not possible before NRP integration. As the prediction error produced by the PCN reflects drift in the network, it could serve as a parsimonious indicator of poor predictions. In novel environments or in the dark, where visual information is unreliable, we would expect this error to increase and the SNN to drift, as the PCN would be unable to reconstruct the head angle associated with a novel visual scene. This highlights, at least at a functional level of description, a biologically plausible error signal that could be used to drive learning as required. Furthermore, this error signal can act as a proxy for confidence in a particular modality, allowing the agent to dynamically apportion its reliance on allothetic vs ideothetic cues when estimating head direction.

Currently the simulation inside the NRP runs slower than real time. The NRP allows this by blocking processes, such as transfer functions (including the PCN) and the brain model, until the current process is complete. As the online NRP currently utilises CPUs only, we have restricted our work to CPUs, however Tensorflow is suited to using GPUs and as NEST is CPU intensive, a combination of the two will lead to faster simulation. Further, by working with SNN invites the future integration of SpiNNaker based neuromorphic hardware [5] via the NRP and ultimately leading toward real-time, embedded physical robot deployment. Using these methods and tools enable new neuroscience questions to be explored which in turn inspires new questions for experimental studies to be generated. We anticipate similar experiments will encourage collaboration between experimental and computational neuroscientists or roboticists, working towards the principles of Replacement and Reduction of animal research set out by the NC3Rs.

The use of the NRP presented in this paper centres around rat navigation. Though the hippocampal formation has been clearly linked to navigation in mammals, there are competing models of how the various sub-problems of navigation - integrating position and head direction, incorporating external sensory cues, dealing with unreliable sensors - are solved in the brain. This paper's solution is consistent with available neuroscience data and theory; a ring attractor network formed by reciprocal LMN-DTN connections, driven by angular head velocity input from the supragenual nucleus, with a deep predictive coding model of sensory cortices feeding corrective inputs via the postsubiculum. Our contribution in this space being that one of the intrinsic network components of the predictive coding approach to learning can be used as a useful proxy for determining confidence in associating sensory stimuli to head directions. This is turn could be used to drive further learning or initiate changes in behaviour to accommodate the sensory ambiguity.

Acknowledgements

This research has received funding from the European Union’s Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3). Chinese Garden in Stuttgart panorama taken and released under CC0 license by Andreas Mitchok.

References

1. Aljalbout, E., Walter, F., Röhrbein, F., Knoll, A.: Task-independent spiking central pattern generator: A learning-based approach. *Neural Processing Letters* **51**(3), 2751–2764 (2020)
2. Bassett, J.P., Taube, J.S.: Neural correlates for angular head velocity in the rat dorsal tegmental nucleus. *The Journal of Neuroscience* **21**, 5740–51 (8 2001)
3. Blair, H.T., Cho, J., Sharp, P.E.: The anterior thalamic head-direction signal is abolished by bilateral but not unilateral lesions of the lateral mammillary nucleus. *The Journal of neuroscience* **19**, 6673–83 (8 1999)
4. Falotico, E., Vannucci, L., Ambrosano, A., Albanese, U., Ulbrich, S., Vasquez Tieck, J.C., Hinkel, G., Kaiser, J., Peric, I., Denninger, O., Cauli, N., Kirtay, M., Roennau, A., Klinker, G., Von Arnim, A., Guyot, L., Peppicelli, D., Martínez-Cañada, P., Ros, E., Maier, P., Weber, S., Huber, M., Plecher, D., Röhrbein, F., Deser, S., Roitberg, A., van der Smagt, P., Dillman, R., Levi, P., Laschi, C., Knoll, A.C., Gewaltig, M.O.: Connecting artificial brains to robots in a comprehensive simulation framework: The neurorobotics platform. *Frontiers in Neurobotics* **11**, 2 (2017). <https://doi.org/10.3389/fnbot.2017.00002>
5. Furber, S., Galluppi, F., Temple, S., Plana, L.: The spinaker project. *Proceedings of the IEEE* **102**, 652–665 (05 2014). <https://doi.org/10.1109/JPROC.2014.2304638>
6. Gewaltig, M.O., Diesmann, M.: Nest (neural simulation tool). *Scholarpedia* **2**(4), 1430 (2007)
7. Goodridge, J.P., Taube, J.S.: Preferential use of the landmark navigational system by head direction cells in rats. *Behavioral neuroscience* **109**, 49–61 (1995). <https://doi.org/10.1037/0735-7044.109.1.49>
8. Goodridge, J.P., Dudchenko, P.A., Worboys, K.A., Golob, E.J., Taube, J.S.: Cue control and head direction cells. *Behavioral neuroscience* **112**, 749–761 (1998). <https://doi.org/10.1037/0735-7044.112.4.749>
9. Hunt, A., Szczecinski, N., Quinn, R.: Development and training of a neural controller for hind leg walking in a dog robot. *Frontiers in Neurobotics* **11** (2017)
10. Knowles, T.C., Stentiford, R., Pearson, M.J.: Whiskey: A biomimetic model of multisensory spatial memory based on sensory reconstruction. In: *Lecture Notes in Artificial Intelligence* 13054. pp. 408–418. Springer (9 2021). https://doi.org/10.1007/978-3-030-89177-0_43
11. Koenig, N., Howard, A.: Design and use paradigms for gazebo, an open-source multi-robot simulator. In: *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE Cat. No.04CH37566). vol. 3, pp. 2149–2154 vol.3 (2004). <https://doi.org/10.1109/IROS.2004.1389727>
12. NC3Rs: The 3rs., <https://www.nc3rs.org.uk/who-we-are/3rs>

13. Pearson, M.J., Dora, S., Struckmeier, O., Knowles, T.C., Mitchinson, B., Tiwari, K., Kyrki, V., Bohte, S., Pennartz, C.M.: Multimodal representation learning for place recognition using deep hebbian predictive coding. *Frontiers in Robotics and AI* **8** (12 2021). <https://doi.org/10.3389/frobt.2021.732023>
14. Prescott, T.J., Redgrave, P., Gurney, K.: Layered control architectures in robots and vertebrates. *Adaptive Behavior* **7**(1), 99–127 (1999)
15. Stackman, R.W., Golob, E.J., Bassett, J.P., Taube, J.S.: Passive transport disrupts directional path integration by rat head direction cells. *Journal of neurophysiology* **90**, 2862–2874 (11 2003). <https://doi.org/10.1152/JN.00346.2003>
16. Stanford Artificial Intelligence Laboratory et al.: Robotic operating system, <https://www.ros.org>
17. Stentiford, R., Knowles, T.C., Pearson, M.J.: A spiking neural network model of rodent head direction calibrated with landmark free learning. *Frontiers in Neurorobotics* **16** (2022). <https://doi.org/10.3389/fnbot.2022.867019>, <https://www.frontiersin.org/article/10.3389/fnbot.2022.867019>
18. Szczecinski, N., Hunt, A., Quinn, R.: Design process and tools for dynamic neuromechanical models and robot controllers. *Biological Cybernetics* **111**, 105–127 (2017)
19. Taube, J.S., Burton, H.L.: Head direction cell activity monitored in a novel environment and during a cue conflict situation. *Journal of neurophysiology* **74**, 1953–1971 (1995). <https://doi.org/10.1152/JN.1995.74.5.1953>
20. Yoder, R.M., Peck, J.R., Taube, J.S.: Visual landmark information gains control of the head direction signal at the lateral mammillary nuclei. *Journal of Neuroscience* **35**, 1354–1367 (2015). <https://doi.org/10.1523/JNEUROSCI.1418-14.2015>
21. Yoder, R.M., Taube, J.S.: Head direction cell activity in mice: robust directional signal depends on intact otolith organs. *The Journal of neuroscience* **29**, 1061–1076 (1 2009). <https://doi.org/10.1523/JNEUROSCI.1679-08.2009>
22. Yoder, R.M., Taube, J.S.: The vestibular contribution to the head direction signal and navigation. *Frontiers in Integrative Neuroscience* **8**, 32 (4 2014). <https://doi.org/10.3389/fnint.2014.00032>