
Proceedings Paper:
Whelan, M.T., Prescott, T.J. orcid.org/0000-0003-4927-5390 and Vasilaki, E. orcid.org/0000-0003-3705-7070 (2020) Fast reverse replays of recent spatiotemporal trajectories in a robotic hippocampal model. In: Vouloutsi, V., Mura, A., Tauber, F., Speck, T., Prescott, T.J. and Verschure, P.F.M.J., (eds.) Biomimetic and Biohybrid Systems : 9th International Conference, Living Machines 2020, Freiburg, Germany, July 28–30, 2020, Proceedings. Living Machines 2020, 28-30 Jul 2020, Freiburg, Germany. Lecture Notes in Computer Science (12413). Springer, Cham , pp. 390-401. ISBN 9783030643126

https://doi.org/10.1007/978-3-030-64313-3_37

Fast Reverse Replays of Recent Spatiotemporal Trajectories in a Robotic Hippocampal Model

Matthew T. Whelan[1,2], Tony J. Prescott[1,2], and Eleni Vasilaki[1]

[1] Department of Computer Science, University of Sheffield, Sheffield, UK, S10 2TN
[2] Sheffield Robotics, University of Sheffield, Sheffield, UK, S1 3JD
{mwhelan3,e.vasilaki,t.j.prescott}@sheffield.ac.uk

Abstract. A number of computational models have recently emerged in an attempt to understand the dynamics of hippocampal replay, but there has been little progress in testing and implementing these models in real-world robotics settings. Presented here is a hippocampal CA3 network model, that runs in real-time to produce reverse replays of recent spatiotemporal sequences in a robotic spatial navigation task. For the sake of computational efficiency, the model is composed of continuous-rate based neurons, but incorporates two biophysical properties that have recently been hypothesised to play an important role in the generation of reverse replays: intrinsic plasticity and short-term plasticity. As this model only replays recently active trajectories, it does not directly address the functional properties of reverse replay, for instance in robotic learning tasks. But it does support further investigations into how reverse replays could contribute to functional improvements, since a biophysical robotic model of reverse replay emergence is a necessary condition towards this goal.

Keywords: Robotics · Hippocampal replay · Computational modelling.

1 Introduction

How the nervous system represents, stores and retrieves memories is an ongoing research problem, but an interesting hypothesis now gaining strong experimental support is that hippocampal replay plays an important role [9, 14, 7, 24]. Place cell activities in the CA region of the hippocampus, which are cells that respond preferentially when a rodent is positioned in the place cell's spatial receptive field [23, 22], are often invoked during hippocampal replay events and are therefore a useful concept for understanding hippocampal replay. In its simplest form, hippocampal replay is the temporally preserved reactivation of recently active place cells during sleep [37, 31, 18] and during periods of awake immobility or quiescence [4, 28], and have been shown to occur during brief periods of hippocampal sharp-wave ripple events [6]. Replays can reinstate the temporal ordering of the place cells in either the forward direction [31, 18] or the reverse direction [8], termed *forward replay* and *reverse replay*, respectively.

There are a number of important differences in both the behavioural and neural states between reverse replays and forward replays. For instance, whilst

forward replays show to occur during both awake and sleep states, and can reinstate either remotely (i.e. reinstating experiences from locations that are spatially distant from the original experience) or locally [4], reverse replays tend to occur almost exclusively during awake replay events, reflecting the immediate local past experience [8]. And unlike awake forward replays, reverse replays are strongly modulated by rewards, such that the frequency of reverse replays initiating at reward sites increases in the presence of increased rewards and reduces in response to decreased rewards [1]. In addition, reverse replays are capable of emergence following only a single trial, and it has therefore been proposed that reverse replays may be a consequence of place cell excitability changes not relying on traditional synaptic plasticity, such as lingering place cell activities [6], or more recently long-term potentiation of intrinsic excitability [25]. Reverse replays may then induce synaptic changes that provide the mechanism for replays during states where the transient activity has faded, such as during sleep states [2]. These point towards an intricate interplay between transient excitability changes and synaptic plasticities in the emergence of the various forms of replay under different behavioural states.

Presented here is a biophysical, continuous rate-based network model of reverse replay implemented on a simulated version of the biomimetic robot MiRo [20], and is based on two recent models of hippocampal replay dynamics in recurrent CA3 networks [11, 25]. Reverse replays in this model occur as a consequence of two modes of transient neural states. The first is due to the implementation of a time decaying model of *intrinsic plasticity*. Intrinsic plasticity is the ability of a cell to increase heterosynaptic long-term potentiation of post-synaptic potentials following recent activity [38, 13], and has recently been proposed as a potential mechanism for the occurrence of reverse replays [25]. The second transient neural state implementation is in short-term plasticity, which acts to ensure unidirectional, stable replays [11]. This is due to short-term depression suppressing synaptic currents after a given amount of continuous firing, thus preventing unbounded synaptic transmissions.

2 Methods

2.1 Network Architecture

The network consists of 100 rate-based neurons representing place cells, arranged in a grid of size 10×10, each of which has its place fields spread evenly across an open circular environment. Each cell forms a bidirectional and symmetric synaptic connection to its 8 nearest neighbours, with all weights fixed at a value of 1. Figure 1 gives an example of the network architecture for a subset of cells.

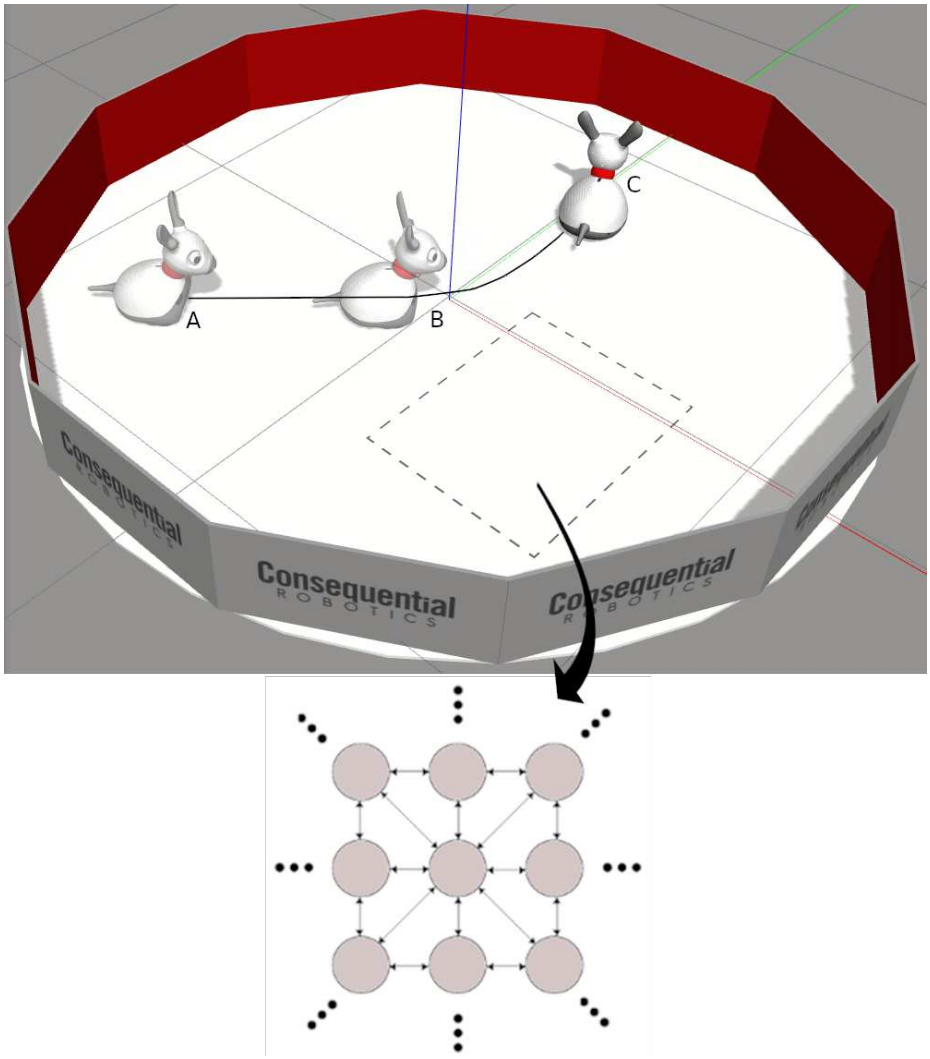

Fig. 1. The simulated environment used to test the model with the MiRo robot. The network architecture consists of a 10x10 array of place cells with place fields uniformly covering the environment. Bidirectional symmetric connections exist between each cell's eight nearest neighbours in space, as shown for a small patch of the environment here. An example trajectory is shown here, in which MiRo begins at the start position in A, passes through location B and ends in the goal location at C.

2.2 Network Dynamics

The activity of a single neuron, i, is described by a first-order decaying differential, with its activity increasing according to incoming recurrent synaptic inputs and place specific inputs, and reducing in response to a global inhibitory term

$$\tau_I \frac{d}{dt} I_i = -I_i + \sigma_i I_i^{syn} + I_i^{place} - I^{inh} \tag{1}$$

with $\tau_I = 0.05s$. The activity is passed through a linear rectifier with a lower bound of 0Hz and an upper bound of 100Hz to give the final rate, r_i

$$r_i^{'} = \alpha \left(I_i - \epsilon \right)$$

$$r_i = \begin{cases} 0 & \text{if } r_i^{'} < 0 \\ 100 & \text{if } r_i^{'} > 100 \\ r_i^{'} & \text{otherwise} \end{cases} \tag{2}$$

with $\alpha = 1$ and $\epsilon = 2$.

I_i^{place} is the place specific input, for which each neuron has associated with it a place field in the environment (in this instance a 2m×2m environment, Figure 1). The place fields are spread uniformly across the environment, each having a centre point, (x_i, y_i). The place-specific input for neuron i is then given as an exponential of the distance the robot is from the place field's centre point,

$$I_i^{place} = I_{max}^{p} \exp \left[-\frac{(x - x_i)^2 + (y - y_i)^2}{2d^2} \right] \tag{3}$$

with $I_{max}^{p} = 50Hz$ and $d = 0.1m$.

I_i^{syn} represents the synaptic input and is given as the sum of the incoming synaptic connections from the cell's 8 nearest neighbours

$$I_i^{syn} = \lambda \sum_{j=1}^{8} w_{ij} r_j D_j F_j \tag{4}$$

where w_{ij} represents the weight from neuron j onto neuron i. λ takes on a value of 0 or 1, depending on whether the robot is exploring or resting at the reward (see Section 2.3), respectively. D_j and F_j are short-term plasticity terms representing short-term depression and short-term facilitation, respectively, and are described by (as in [11], but see [32])

$$\frac{d}{dt} D_j = \frac{1 - D_j}{\tau_{STD}} - r_j D_j F_j \tag{5}$$

$$\frac{d}{dt} F_j = \frac{U - F_j}{\tau_{STF}} + U \left(1 - F_j \right) r_j \tag{6}$$

with $\tau_{STD} = 1.5s$, $\tau_{STF} = 1s$ and $U = 0.6$. If a cell fires continuously for a given amount of time, eventually D_j drops to 0 thus preventing that cell from any further synaptic transmissions.

The inhibitory input, I_i^{inh}, is a global term given as a summation of the whole network's activity

$$\frac{d}{dt}I^{inh} = -\frac{I^{inh}}{\tau^{inh}} + w_{inh}\sum_j r_j D_j F_j \tag{7}$$

with $\tau^{inh} = 0.05s$ and $w_{inh} = 0.1$, and acts to prevent too many cells being active at once.

The σ_i term in Equation 1 is specific to each cell, representing the intrinsic plasticity for that cell. It acts to scale the incoming synaptic inputs and is described by

$$\frac{d}{dt}\sigma_i = \frac{\sigma_{ss} - \sigma_i}{\tau_\sigma} + \frac{\sigma_{max} - 1}{1 + \exp\left[-\beta(r_i - r_\sigma\right]} \tag{8}$$

with $\tau_\sigma = 10s$, $\sigma_{ss} = 0.1$, $\sigma_{max} = 4$, $r_\sigma = 10Hz$ and $\beta = 1$. The second term is a sigmoid, and follows the modelling approach taken by Pang and Fairhall [25], but with the addition here of time decaying dynamics to model extinction effects. When the cell no longer fires, σ_i decays to a steady state value of σ_{ss}. If $\sigma_i > \sigma_{max}$, then σ_i is set to σ_{max}.

2.3 Two-Stage Dynamics

Implemented here are two different behavioural states that define two different sets of network dynamics, similar to a previous two-stage modelling approach of CA3 replay dynamics [29].

The first behavioural state is defined as *active exploration*: the state in which MiRo is actively searching for the hidden reward. Under this state, it is assumed there is little to no synaptic transmission across the network due to the effects of high acetylcholine levels [17], which has been shown experimentally to inhibit the recurrent post-synaptic inputs in the hippocampal CA3 region [12]. To capture this effect, λ is set to 0 in Equation 4 thus preventing post-synaptic transmission.

The second behavioural state is defined as *quiescent reward*, in which MiRo remains awake yet quiescent whilst it is at the reward point, and is the state under which reverse replays occur. It is assumed here that acetylcholine levels have dropped, similar to that found during slow-wave sleep states [17], thus permitting synaptic transmission in the CA3 network. To model this effect, λ is set to 1 in Equation 4.

3 Results

3.1 Searching for a Hidden Reward

The model is run on the MiRo robot in a simulated open arena environment, having a diameter of 2m (Figure 1) and with simulation time steps of 10ms. From a random start location, MiRo is left to freely explore its environment via a basic implementation of a random walk, with the goal of finding a hidden reward. This is the *active exploration* phase, and during this phase the network rates are driven solely by the place specific inputs with no recurrent synaptic transmissions. There is no synaptic plasticity implemented in this experiment, and so all weights, w_{ij}, are fixed at a value of 1. Figures 2A and 2B show the activity of the network during active exploration. Due to the distribution of the place-specific input, no more than 4 cells are active at any one time, though most often this amounts to no more than 2 or 3 cells being simultaneously active. This sparse representation during exploration provides a neural representation of space. Neurons that become active due to the place specific input then undergo increases in intrinsic plasticity, decaying exponentially (according to Equation 8) when activity in the neuron drops.

Upon reaching the hidden reward location, MiRo pauses and enters the *quiescent reward* phase. Place specific inputs are computed using Equation 3 and are input into the network via pulses of 0.1s-ON and 1.9s-OFF. Recall that during this phase, recurrent synaptic conductances are allowed. Due to the increase in synaptic recurrent conductance and post-synaptic activity being scaled by the intrinsic plasticity, activity propagates quickly through the network, reinstating the most recently active cells in a temporally reversed order to that seen during exploration. Figure 2C shows the activity of the network midway through a replay event. Notice the trace in the intrinsic plasticity plots, which transiently stores the most recent sequence of activity in the network and provides the mechanism for faithful replays of the recent trajectory. In this instance, many more cells are found to be simultaneously active, but their time points for peak activity retain the temporal ordering seen during exploration (Figure 3).

In order to provide a more detailed comparison of the network's activity during the exploration phase versus the quiescent phase in which reverse replays occur, Figure 3 displays a time course plot of the rates for the 14 cells that were active during exploration in Figure 2A. It is clear in Figure 3 that the temporal ordering of cell firing during a reverse replay event is preserved in comparison to the ordering during exploration.

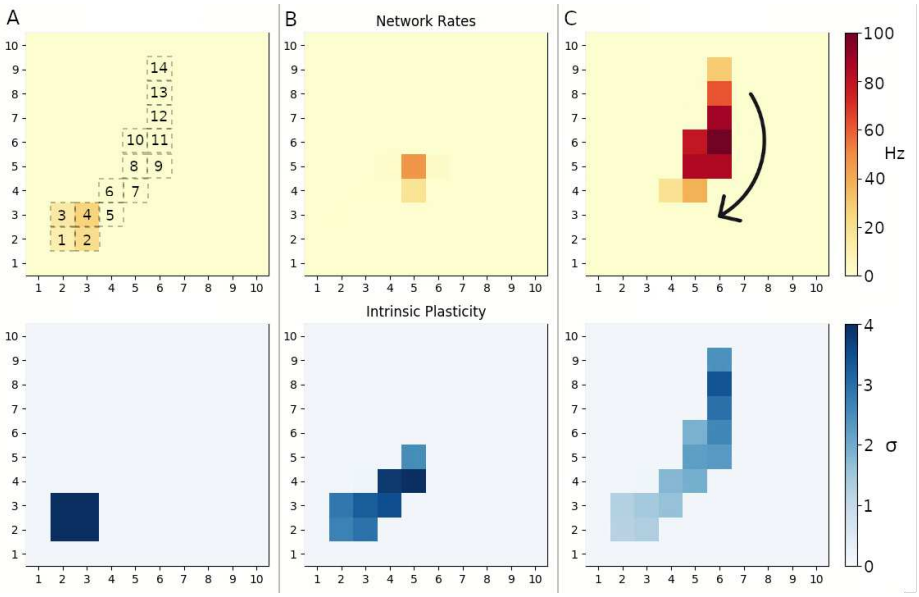

Fig. 2. Rates (top plots) and intrinsic plasticities (bottom plots) for the 10x10 network are shown here for the locations marked in the trajectory of Figure 1. These are: A) MiRo is at the start location. The numbered boxes ranging from 1 to 14 here represent all cells that were active during the exploration phase and the temporal order in which they fired during exploration (i.e. cell 1 fired first, cell 14 last). Note however that at the start point, only the first 4 cells were active. B) MiRo is exploring the environment. C) MiRo has reached the reward and reverse replays are being initiated. The arrow indicates the temporal order of firing during this replay event.

3.2 Removing intrinsic plasticity

To show the effects of removing intrinsic plasticity from the model, σ_i is set to 1 for all cells and the model is run once more on a similar trajectory (Figure 4). In this instance, rather than a direct replay of the recent trajectory, the activity in the network displays a divergent replay event across the whole network from the point of initiation. This effect was similarly seen in the model of Haga and Fukai [11], who assumed a similar network architecture to this one but did not model intrinsic plasticity. This shows that the intrinsic plasticity is important for restricting the replay event to the previously experienced trajectory only. However, divergent replays could have potential benefits in the learning of goal-oriented paths (see Discussion).

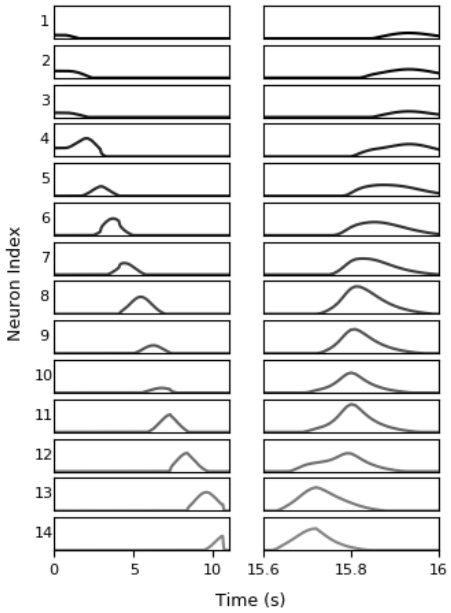

Fig. 3. A time course plot of the cell rates for the cells indexed in Figure 2A. The lower and upper limits in each box plot is 0Hz and 100Hz. Plots on the left show the activities during exploration, occurring over a time period of approximately 12s. The plots on the right show the activity during a reverse replay event. Note that Figures 2A, B and C are snapshots of the network's activity at times 0s, 5s and 15.8s, respectively.

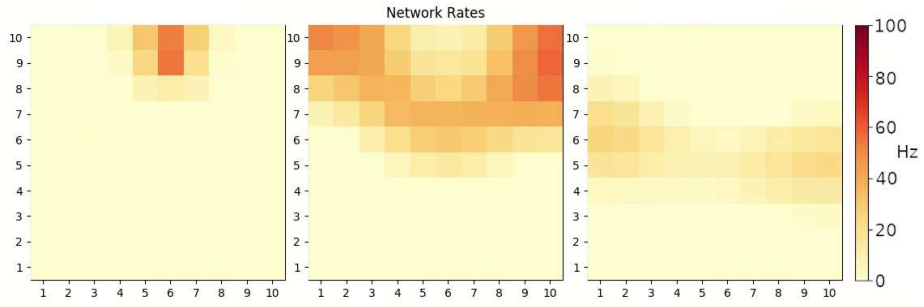

Fig. 4. Example of a replay event without intrinsic plasticity, where $\sigma_i = 1$ for all neurons. A similar trajectory as in Figures 1/2 is taken here, with reverse replay events initiated at the same location. The heat maps, from left to right, show the temporal ordering of network activity during a replay event. As intrinsic plasticity is homogeneous across the network, there is no preferential trajectory for the sequence of cell activities to follow. As such a divergent wave propagates across the whole network from the point of initiation.

4 Discussion

We have presented here a biophysical model of a CA3 hippocampal network that produces fast reverse replays of recently active place cell trajectories. Whilst the network connectivity remains static and symmetric, the implementation of intrinsic plasticity produces asymmetries in the network that amplifies incoming synaptic currents, enabling activity to travel through the network along a trajectory determined by levels of intrinsic plasticity. Intrinsic plasticity was first introduced as a potential mechanism for hippocampal replays by Pang and Fairhall [25], but as we are running the model on the MiRo robot, for which it can very quickly cover a whole area, time decaying dynamics have had to be included so that the whole network does not become intrinsically potentiated. Given only a subset of the network becomes potentiated by intrinsic plasticity (i.e. those cells most recently active), this creates a certain level of sparsity in the network, and is interesting to compare with a previous computational model of replay dynamics by Chenkov et al [5] who showed that sparsity in their network was important for generating effective and controlled replays. Yet, whilst they achieved sparsity by changing the number of synaptic connections, here it is achieved through intrinsic plasticity changes. These results nonetheless point towards a level of sparsity that is important for specific and controlled replays.

Another important component in this model for generating stable propagations of replay sequences is short-term plasticity effects. This mechanism was first shown by Haga and Fukai [11] in a reverse replay model. It is perhaps a useful analogy to consider short term plasticity in this instance having the effect of a 'refractory period' for activity propagation, in that it prevents further transmission of activity after a given amount of continuous activity. Refractory periods have been shown in previous models to ensure stable, unidirectional replays [15, 25]. However, implementing refractory periods requires a model of spiking neurons, and so modelling short-term plasticity lends itself to rate-based implementations of replay. This is of course particularly useful in real-time robotic applications where spiking neuron models may be computationally inefficient. But short-term plasticity could have a more interesting property during reverse replays. Haga and Fukai [11] showed that short-term plasticity could generate reversed synaptic weight changes. This enables reverse replays to strengthen synaptic traces in the forwards direction, despite the replay event occurring in the reverse. Thus, whilst their model produced divergent replay events similar to that seen here when intrinsic plasticity is removed, the reversed synaptic potentiations proved useful in generating synaptic traces towards a goal location, even if particular place cells had not been active during exploration. These could prove useful if, for instance, the network connectivity provides a neural map of the environment. Replays could then provide a means to explore trajectories towards goal locations even for trajectories that have never been physically explored.

A third component of the model that was necessary for appropriately timed replays was the implementation of a two-stage dynamic, which prevented the network from transmitting recurrent synaptic currents during the *exploration phase*, but allowed synaptic transmission during the *quiescent reward* phase (where

MiRo sat quietly at the reward location). This was based on findings that suggest different levels of acetylcholine during active exploration and sleep states [17], which alters CA3 synaptic conductances [12] – higher levels of acetylcholine inhibit synaptic conductance. However, what is not clear is that acetylcholine levels drop significantly enough during the quiescent reward state for which reverse replays occur, given it follows immediately after exploration [8]. Whilst levels of acetylcholine have been found to change quickly on the time scale of a few seconds, at least in the prefrontal cortex [26], it is unclear as to whether this occurs in the hippocampal CA3 region. What is perhaps interesting to note, however, is that cholinergic stimulation, which leads to an increase in acetylcholine, has been shown to suppress hippocampal sharp-wave ripples yet promote theta oscillations [33]. Given theta activity is found to co-occur with exploratory states [34], whilst replays occur usually during sharp-wave ripple events [6], this suggests that for reverse replays to arise, acetylcholine levels must phasically drop during a quiescent reward state to enable sharp-wave ripples.

4.1 Scope for Future Research

We argued previously that, though there are a number of computational models attempting to explain the dynamics of hippocampal replay, there had been little in terms of real-world robotic applications of these models [36]. And whilst there does exist reinforcement learning models that attempt to capture some of the functional properties of replay [21, 3, 19], they are not biophysical models, nor do they adopt continuous state-action spaces which are likely necessary for robotic applications. Thus, though these models may perhaps help answer the question of why replays are functionally useful, they do not answer how replays emerge. A complete model of hippocampal replay should ideally answer both these questions. This work attempts to understand the problem of how hippocampal replay emerges by utilising robotics to test the models in real-world settings, and in so doing, could help bridge the gap between those computational models that seek to understand how replays emerge with those that utilise replays for improving learning in artificial systems. As noted, reverse replays are particularly thought to be involved in reinforcement learning, given their reward-modulated occurrences [1] and coordinated activity with neurons in the striatum [27, 10]. As such, an investigation into whether reverse replays could improve learning in biophysical models of reinforcement learning, such as [35], should be conducted to further ground these hypotheses in theory.

Finally, in order to generate cell activity that was place specific in this model, the global x-y coordinates for the robot had to be used which is a biologically unrealistic property of place cell emergence. Questions remain therefore around how other models of the hippocampus, ones attempting to understand the emergence of place cells, grid cells, head direction cells, etc., in this instance particularly robotic implementations [30, 16], could be consolidated with hippocampal replay models. For consolidating these dynamics is surely a necessary condition for a full working model of the hippocampus.

References

1. Ambrose, R.E., Pfeiffer, B.E., Foster, D.J.: Reverse replay of hippocampal place cells is uniquely modulated by changing reward. Neuron **91**(5), 1124–1136 (2016)
2. Atherton, L.A., Dupret, D., Mellor, J.R.: Memory trace replay: the shaping of memory consolidation by neuromodulation. Trends in neurosciences **38**(9), 560–570 (2015)
3. Aubin, L., Khamassi, M., Girard, B.: Prioritized sweeping neural dynaq with multiple predecessors, and hippocampal replays. In: Conference on Biomimetic and Biohybrid Systems. pp. 16–27. Springer (2018)
4. Carr, M.F., Jadhav, S.P., Frank, L.M.: Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. Nature neuroscience **14**(2), 147 (2011)
5. Chenkov, N., Sprekeler, H., Kempter, R.: Memory replay in balanced recurrent networks. PLoS computational biology **13**(1), e1005359 (2017)
6. Diba, K., Buzsáki, G.: Forward and reverse hippocampal place-cell sequences during ripples. Nature neuroscience **10**(10), 1241 (2007)
7. Foster, D.J.: Replay comes of age. Annual review of neuroscience **40**, 581–602 (2017)
8. Foster, D.J., Wilson, M.A.: Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature **440**(7084), 680 (2006)
9. Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., Zugaro, M.B.: Selective suppression of hippocampal ripples impairs spatial memory. Nature neuroscience **12**(10), 1222 (2009)
10. Gomperts, S.N., Kloosterman, F., Wilson, M.A.: Vta neurons coordinate with the hippocampal reactivation of spatial experience. Elife **4**, e05360 (2015)
11. Haga, T., Fukai, T.: Recurrent network model for learning goal-directed sequences through reverse replay. Elife **7**, e34171 (2018)
12. Hasselmo, M.E., Schnell, E., Barkai, E.: Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region ca3. Journal of Neuroscience **15**(7), 5249–5262 (1995)
13. Hyun, J.H., Eom, K., Lee, K.H., Bae, J.Y., Bae, Y.C., Kim, M.H., Kim, S., Ho, W.K., Lee, S.H.: Kv1. 2 mediates heterosynaptic modulation of direct cortical synaptic inputs in ca3 pyramidal cells. The Journal of physiology **593**(16), 3617–3643 (2015)
14. Jadhav, S.P., Kemere, C., German, P.W., Frank, L.M.: Awake hippocampal sharp-wave ripples support spatial memory. Science **336**(6087), 1454–1458 (2012)
15. Jahnke, S., Timme, M., Memmesheimer, R.M.: A unified dynamic model for learning, replay, and sharp-wave/ripples. Journal of Neuroscience **35**(49), 16236–16258 (2015)
16. Jauffret, A., Cuperlier, N., Gaussier, P.: From grid cells and visual place cells to multimodal place cell: a new robotic architecture. Frontiers in neurorobotics **9**, 1 (2015)
17. Kametani, H., Kawamura, H.: Alterations in acetylcholine release in the rat hippocampus during sleep-wakefulness detected by intracerebral dialysis. Life sciences **47**(5), 421–426 (1990)
18. Lee, A.K., Wilson, M.A.: Memory of sequential experience in the hippocampus during slow wave sleep. Neuron **36**(6), 1183–1194 (2002)
19. Mattar, M.G., Daw, N.D.: Prioritized memory access explains planning and hippocampal replay. Nature Neuroscience **21**(11), 1609 (2018)

20. Mitchinson, B., Prescott, T.J.: Miro: a robot "mammal" with a biomimetic brain-based control system. In: Conference on Biomimetic and Biohybrid Systems. pp. 179–191. Springer (2016)
21. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. Nature **518**(7540), 529 (2015)
22. O'Keefe, J.: Place units in the hippocampus of the freely moving rat. Experimental neurology **51**(1), 78–109 (1976)
23. O'Keefe, J., Dostrovsky, J.: The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. Brain research (1971)
24. Ólafsdóttir, H.F., Bush, D., Barry, C.: The role of hippocampal replay in memory and planning. Current Biology **28**(1), R37–R50 (2018)
25. Pang, R., Fairhall, A.L.: Fast and flexible sequence induction in spiking neural networks via rapid excitability changes. eLife **8**, e44324 (2019)
26. Parikh, V., Kozak, R., Martinez, V., Sarter, M.: Prefrontal acetylcholine release controls cue detection on multiple timescales. Neuron **56**(1), 141–154 (2007)
27. Pennartz, C., Lee, E., Verheul, J., Lipa, P., Barnes, C.A., McNaughton, B.: The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples. Journal of Neuroscience **24**(29), 6446–6456 (2004)
28. Pfeiffer, B.E., Foster, D.J.: Hippocampal place-cell sequences depict future paths to remembered goals. Nature **497**(7447), 74–79 (2013)
29. Saravanan, V., Arabali, D., Jochems, A., Cui, A.X., Gootjes-Dreesbach, L., Cutsuridis, V., Yoshida, M.: Transition between encoding and consolidation/replay dynamics via cholinergic modulation of can current: a modeling study. Hippocampus **25**(9), 1052–1070 (2015)
30. Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arleo, A., Gerstner, W.: Is there a geometric module for spatial orientation? insights from a rodent navigation model. Psychological review **116**(3), 540 (2009)
31. Skaggs, W.E., McNaughton, B.L.: Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. Science **271**(5257), 1870–1873 (1996)
32. Tsodyks, M., Pawelzik, K., Markram, H.: Neural networks with dynamic synapses. Neural computation **10**(4), 821–835 (1998)
33. Vandecasteele, M., Varga, V., Berényi, A., Papp, E., Barthó, P., Venance, L., Freund, T.F., Buzsáki, G.: Optogenetic activation of septal cholinergic neurons suppresses sharp wave ripples and enhances theta oscillations in the hippocampus. Proceedings of the National Academy of Sciences **111**(37), 13535–13540 (2014)
34. Vanderwolf, C.H.: Hippocampal electrical activity and voluntary movement in the rat. Electroencephalography and clinical neurophysiology **26**(4), 407–418 (1969)
35. Vasilaki, E., Frémaux, N., Urbanczik, R., Senn, W., Gerstner, W.: Spike-based reinforcement learning in continuous state and action space: when policy gradient methods fail. PLoS computational biology **5**(12), e1000586 (2009)
36. Whelan, M.T., Vasilaki, E., Prescott, T.J.: Robots that imagine – can hippocampal replay be utilized for robotic mnemonics? In: Conference on Biomimetic and Biohybrid Systems. pp. 277–286. Springer International Publishing (2019)
37. Wilson, M.A., McNaughton, B.L.: Reactivation of hippocampal ensemble memories during sleep. Science **265**(5172), 676–679 (1994)
38. Zhang, W., Linden, D.J.: The other side of the engram: experience-driven changes in neuronal intrinsic excitability. Nature Reviews Neuroscience **4**(11), 885–900 (2003)