Dear Prof. Frank,

We would like to thank you and the reviewers for the constructive and positive appraisal of our manuscript titled *"Rapid learning of predictive maps with STDP and theta phase precession*". We are extremely pleased that reviewer 1 was satisfied with the changes made in our initial revisions and that we also answered many of the points raised by reviewer 2. We have now addressed the final points outstanding from reviewer 2 which you helpfully summarised in the 'Editor's correspondence'.

Specifically, we have included further simulations to demonstrate the biological plausibility of the model and have added text to discuss the perceived biological limitations. We now present that manuscript for re-submission - new or changed text is coloured blue. The main changes made include:

1. **Spiking model**. A simulation showing that the STDP-successor weights learnt in a spiking model - where spiking activity in CA3 (as opposed to their rates) drives the spiking of CA1 neurons using a linear-nonlinear cascade model - is almost identical to the weight matrix learnt in our previous simulations (diagonal-aligned average across rows: R^2=0.99).

2. **CA3 → CA1 projections**. In the same fully-spiking simulation we show that when CA1 cells receive input from a large number of CA3 cells with partially overlapping fields, theta phase precession of spikes is CA1 is preserved and, when coupled with STDP, is almost identical to the STDP-successor matrix generated by our previous simulations (diagonal-aligned average across rows: R^2=0.99).

3. **Phase lags**. A simulation modelling the phase lag between phase precessing place cells in CA3 and CA1 using empirical data from Mizuseki et al., 2012. We show the model still approximates the TD-successor matrix when these realistic phase offsets are taken into account (diagonal-aligned average across rows: R^2=0.76)

4. **Updated Discussion** to account for the above points.

These additional simulations and discussion points are now covered in two new figure panels in Figure 2 figure supplements 2 & 4. Below we provide a detailed point-by-point response to the remaining points as they were presented in the Editor's correspondence. The Editor's points are indented and in italics, followed by our response, and excerpts from the updated manuscript in "quotations" and small font with changes coloured in blue.

> *1. Spiking model. We all agree with you that a full spiking model would be much too complex. However, since you already generate spikes using a Poisson process, it would be useful to see a simulation where the Poisson rate of CA1 cell is determined by the*

integration of the incoming CA3 spikes (perhaps with many incoming CA3 neurons). If this doesn't work, you should discuss why this is the case and what the implications are for the model.

Thank you for this suggestion. In order to address points 1 & 2 (see below), we have updated the manuscript to include a new simulation where the spiking activity in CA1 cells (N=50) is driven by a large number of spiking CA3 neurons (N=500) with overlapping fields that phase precess. To avoid the complexity of Hodgkin-Huxley / Leaky integrate-and-fire models, spiking activity in CA1 is determined by a linear-nonlinear cascade model, which we would like to thank Reviewer 2 for suggesting. In the simulation, which has been added as a panel in Figure 2 Supplement 2 (see below), we find that the resulting weights learnt via STDP in the spiking model are almost identical to those learnt by the standard STDP-successor learning rule used in most of our previous simulations (diagonal-aligned average across rows: R^2=0.99). Note that the resulting weight matrix is no longer square due to the x10 greater number of cells in CA3 vs. CA1.

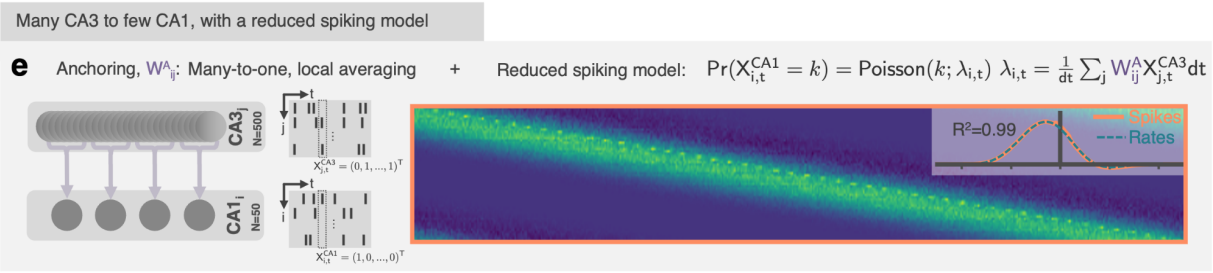

Figure 2 Supplement 2e: 500 CA3 neurons drive 50 CA1 neurons where each CA1 neuron is anchored to a Gaussian-weighted sum of the 10 closest CA3 cells. CA3 spikes now directly drive CA1 spikes according to a reduced spiking model. The inset shows the row-averages and a comparison to the result for an equivalent simulation with the rate-model used in the rest of the paper.

The results of the simulation are also now referred to in the Results:

"This effect is robust to variations in running speed (Fig. 2–supplement 1b) and field sizes (Fig. 2–supplement 1c), as well as scenarios where target CA1 cells have multiple firing fields (Fig. 2–supplement 2a) that are updated online during learning (Fig. 2–supplement 2b-d), or fully-driven by spikes in CA3 (Fig. 2–supplement 2e); see methods for more details."

and details of the simulations have been added to the Methods section 5.10.2 (equations for the spiking model are also summarised in the figure above).

2. CA3 => CA1 projections. CA1 cells still receive input from just one CA3 cell for each place field in the updated model (at least in the majority of simulations). This allows precise theta timing of the pre and post -synaptic neurons which appears to be critical for

the plasticity rule to function. For example, the mathematics of Geisler et al. 2007 shows that, if the CA1 cell would receive input from a set of phase precessing CA3 cells with spatially offset place field and a Gaussian weight profile (the most common way to model CA3-CA1 connections), then the CA1 cell would actually fire at the LFP theta frequency and wouldn't phase precess, and as a consequence the STDP mechanism would no longer learn the successor representation. This suggests strong constraints on the conditions under which the model can function which are currently not being adequately discussed. This should be investigated and discussed, and the constraints required for the model to function should be plainly laid out.

We agree that the one-to-one nature between CA3 => CA1 in the majority of simulations might suggest that this is a strict condition in order for the model to function. Rather, it is a condition we impose in order to simplify the model as well as to establish a clear connection between the model and successor feature theory (see Methods section 5.9). However it is important to note that this condition can easily be relaxed and that doing so still produces results that are extremely similar to true successor feature learning. In order to show this, we have updated the manuscript to include a new simulation, also outlined above, where the spiking activity in CA1 cells (N=50) is driven by a large number of spiking CA3 neurons (N=500) with overlapping fields that phase precess. We show that in this regime, even when the spiking output of each CA1 cell is determined by the spiking input of a large number of phase precessing CA3 cells, the resulting STDP synaptic weights closely resembles that of the STDP-successor matrix used in the majority of simulations (diagonal-aligned average across rows: R^2=0.99).

Additionally we also show a similar result in Figure 2 supplement 2b&c (added after the first round of review) where the synaptic weight matrix is updated online, during learning. In these models all 50 (not just one) CA3 cells are able to drive CA1 cells during learning and SR-like weight matrices still develop. In total, we believe these simulations and the new one performed here demonstrate that many-to-one projections do not pose a fundamental issue.

Regarding the effect on phase precession (or the lack thereof) when cells are driven by multiple neurons, the simulation described above, and the ones provided in response to the first round of reviews, show that this is not a substantial concern. Geisler et al (2010) raised the possibility that in such a situation phase precession would not emerge in CA1. However, our simulations show that it is possible for CA1 cells receiving input from *multiple* CA3 cells to phase precess as long as there is some spatial structure to the connections. If a CA1 cell is most strongly driven by a population of CA3 cells in a similar location on the track (and which therefore phase precess similarly) it too will phase precess. This spatial structure can be quite broad, for evidence of this please see Figure 2 supplement 2f included in our previous rebuttal, and Figure R1 below for an equivalent plot drawn from the spiking model simulation described above. In the figures we show the phase precession of CA1 when driven by the learnt synaptic weight matrix, W, which is significant over a large portion of the input CA3 neurons. This demonstrates that many-to-one connections are not incompatible with phase precession and therefore our proposed learning mechanism can still work.

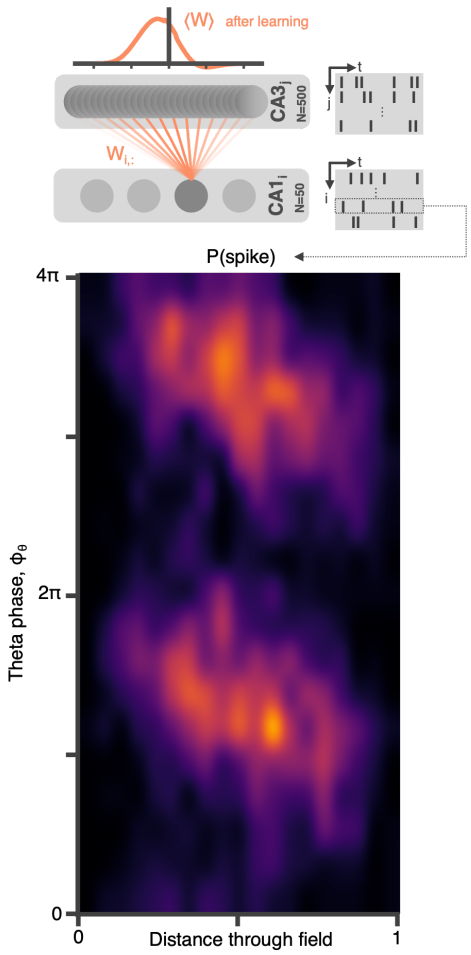

Figure R1. In a fully spiking model, CA1 neurons inherit phase precession from multiple upstream CA3 neurons. Top, model schematic - each CA1 neuron receives input from multiple CA3 neurons with contiguous place fields. Bottom, position vs phase plot for an indicative CA1 neuron, showing strong phase precession similar to that observed in the brain.

3. A similar concern holds with the phase offset between CA3 and CA1 found by Mizuseki et al. The theta+STDP mechanism learns the successor representation because the CA1 cells inherit their responses from a phase-precessing upstream CA3 cell, so the existence of a phase lag is troubling, because it suggests that CA1 cells are not driven causally by CA1 cells in the way the model requires. You may be right that, if some external force were to artificially impose a fixed lag between the CA3 and CA1 cell, the proposed learning mechanism would still function but now with a spatial offset. However, the Reviewer was concerned that the very existence of the phase lag challenges the basic spirit of the model, since CA1 cells are not driven by CA3 cells in the way that is required to learn causal relationships. At the very least, this needs to be addressed and discussed directly and openly in the discussion section, but it would be

better if the authors could implement a solution to the problem to show that the model can work when an additional mechanism is introduced to produce the phase lag (for example, a combination of EC and CA3 inputs at different theta phases?)

The reviewer is correct in that since there is a theta phase offset between CA3 and CA1, it is important to consider the possible impact on our model. Indeed, while Mizuseki et al., 2009 alludes to a fixed phase difference between CA3 and CA1 neurons, the consequences for phase precessing place cells are more nuanced. Importantly, in a later paper from 2012, Misuzeki et al. demonstrate this offset in phase between CA3 and CA1 place cells varies at different stages of the theta cycle. Thus as an animal first enters a place field and spikes are fired late in the theta cycle, CA1 spikes are emitted around 80°to 90° after spikes from CA3. However, as the the animal progresses through the field, spikes from both regions precess to earlier phases but the effect is more pronounced in CA1, meaning that by the time the animal exits a place fields the phase offset between the two regions is essentially 0° (the key figure from Mizuseki et al 2012 is shown below in Figure 2 Supplement 4g). Importantly this result fits with the work of Hasselmo et al (2002) and Colgin et al (2009) both of which point to there being enhanced CA3 > CA1 coupling at early theta phases - in other words CA3's influence on CA1 appears to be most pronounced in the latter half of place fields.

In response to this we have done two things. First, to simulate the effect of a variable phase offset, we ran the model as before but for offsets of 90°, 45°, and 0°, which correspond to late, mid and early theta phase. We then averaged the resulting STDP weight matrices to generate a single prediction for a system in which the CA3 to CA1 phase offset varies in a plausible fashion - the resulting matrix is still very similar to the TD successor matrix (diagonal-aligned average across rows: R^2=0.76), and clearly shows the SR-like asymmetry (positive band left of diagonal, negative band right) confirming that our model is robust to the observed phase offset. These simulations, including the weight matrices for offsets of 90°, 45°, and 0° have now been included in a new figure panel appended to Figure 2 Supplement 4:

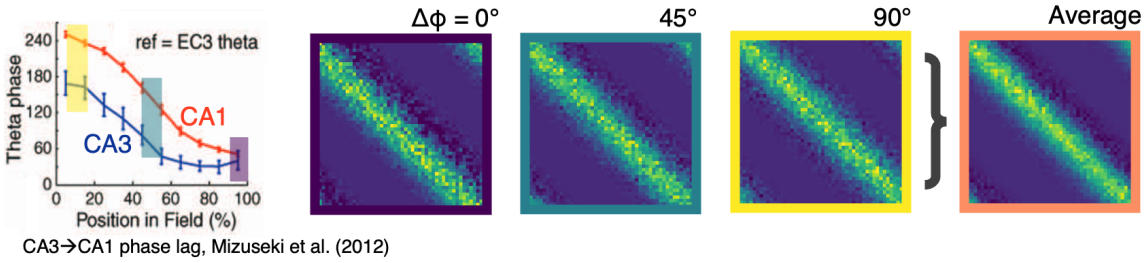

Figure 2 Supplement 4g: (Left) A decreasing phase shift is measured between CA3 and CA1, starting from 90° late in the theta cycle - the phase cells initially spike at as animals enter a field - and ending at 0° early in the cycle, panel adapted from Mizuseki et al. 2012. (Middle) Three phase shifts (0°, 45° and 90°) are simulated and the average of the resulting synaptic weight matrices is taken (right).

and are referred to in the results section:

“Additionally, we find these CA1 cells go on to inherit phase precession from the CA3 population even after learning when they are driven by multiple CA3 fields (Fig. 2–supplement 4f), and that this learning is robust to realistic phase offsets between the populations of CA3 and CA1 place cells (Fig. 2–supplement 4g).”

Secondly, we have also updated the discussion to cover these points in more detail and in particular have addressed the nuances suggested by the experimental results from Hasselmo et al (2002) and Colgin et al (2009). Specifically, we indicate that because CA3>CA1 coupling is most pronounced at early theta phases - when the phase offset between the regions is at its lowest - the effect of the offset is likely to be less important than might immediately be thought. Thus the simulation presented above, which still learns a good approximation of the TD SR matrix (diagonal-aligned average across rows: R^2=0.76), should be considered as a worst-case scenario.

We now expand upon these points in the Discussion:

“While the model is biologically plausible in several respects, there remain a number of aspects of the biology that we do not interface with, such as different cell types, interneurons and membrane dynamics. Further, we do not consider anything beyond the most simple model of phase precession, which directly results in theta sweeps in lieu of them developing and synchronising across place cells over time [60]. Rather, our philosophy is to reconsider the most pressing issues with the standard model of predictive map learning in the context of hippocampus (e.g., the absence of dopaminergic error signals in CA1 and the inadequacy of synaptic plasticity timescales). We believe this minimalism is helpful, both for interpreting the results presented here and providing a foundation on which further work may examine these biological intricacies, such as whether the model’s theta sweeps can alternately represent future routes [61] e.g. by the inclusion of attractor dynamics [62]. Still, we show this simple model is robust to the observed variation in phase offsets between phase precessing CA3 and CA1 place cells across different stages of the theta cycle [63]. In particular, this phase offset is most pronounced as animals enter a field (~90°) and is almost completely reduced by the time they leave it (~0°) (Fig 2 - figure supplement 4g). Essentially our model hypothesises that the majority of plasticity induced by STDP and theta phase precession will take place in the latter part of place fields, equating to earlier theta phases. Notably, this is in-keeping with experimental data showing enhanced coupling between CA3 and CA1 in these early theta phases [64, 65]. However, as our simulations show (figure 2 supplement 4 panel g), even if these assumptions do not hold true, the model is sufficiently robust to generate SR equivalent weight matrices for a range of possible phase offsets between CA3 and CA1 ”

with details of the simulations added to the Methods:

Phase shift between CA3 and CA1. In figure 2 supplement 4g we simulate the effect of a decreasing phase shift between CA3 and CA1. As observed by Mizuseki et al. (2012) [87] there is a phase shift between CA3 and CA1 neurons being maximally around 90 degrees at the end of each theta cycle, decreasing to 0 at the start. We simulate this by adding a temporal delay to all downstream CA1 spikes equivalent to the phase shifts of 0°, 45° and 90°. The average of the weight matrices learned over all three examples still displays clear SR-like structure.

4. - DV phase precession. The Reviewer would still like to see you introduce DV phase lags, which could be done with a simple modification of the existing simulations. At minimum, it is critical to remove/modify the sentence "A consequence of theta phase precession is that the cell with the smaller field will phase precess faster through the theta cycle than the other cell - initially it will fire later in the theta cycle than the cell with a larger field, but as the animal moves towards the end of the small basis field it will fire earlier." As R2 noted in their original review, this is not the case when DV phase lags are taken into account, as was shown by Leibold and Monsalve-Mercado (2017). Ideally, it would be best to update simulations updated to account for the DV phase lags and the discussion updated to account for their functional implications

Thank you for highlighting this. The sentence mentioned was actually intended to be a 'strawman' to motivate the subsequent analyses that show the different rates of phase precession induced by varied field sizes do not impair plasticity in a manner that is sufficient to segregate spatial scales (Figure 4). Note, to be clear we were referring to the fact that small place fields - found at the dorsal pole of the hippocampus - do phase precess more rapidly in *time.* The point being that phase precession is proportional to field size, so for a given distance travelled - say 20cm - a small field will exhibit a greater change in spiking phase than a large one. We apologise for presenting this information in a way that was not clear. We have now made the following changes to ensure the intention of this paragraph is clear:

"Hypothetically, consider a small basis feature cell with a receptive field entirely encompassed by that of a larger basis cell with no theta phase offset between the entry points of both fields. A potential consequence of theta phase precession is that the cell with the smaller field would phase precess faster through the theta cycle than the other cell - initially it would fire later in the theta cycle than the cell with a larger field, but as the animal moves towards the end of the small basis field it would fire earlier. These periods of potentiation and depression instigated by STDP could act against each other, and the extent to which they cancel each other out would depend on the relative placement of the two fields, their size difference, and the parameters of the learning rule."

Similarly, as outlined in the simulations above, graduated theta phase offsets of up to and including 90° are also insufficient to impair the plasticity induced by STDP and phase precession. Applying both of these findings in the context of theta as a travelling wave across the dorsal-ventral axis, our original conclusion that topographic organisation of place cells by size along the DV axis is necessary to prevent cross binding and preserve multiscale structure in the resulting successor features remains unchanged.

We now clarify these points in the discussion:

"The distribution of place cell receptive field size in hippocampus is not homogeneous. Instead, place field size grows smoothly along the longitudinal axis (from very small in dorsal regions to very large in ventral regions). Why this is the case is not clear – our model contributes by showing that, without this ordering, large and small place cells would all bind via STDP, essentially overwriting the short timescale successor representations learnt by small place cells with long timescale successor representations. Topographically organising place cells by size anatomically segregates place cells with fields of different sizes, preserving the multiscale successor representations. Further, our results exploring the effect of different phase

offsets on STDP-successor learning (Fig 2 - figure supplement 4g) suggest that the gradient of phase offsets observed along the dorso-ventral axis [79, 80] is insufficient to impair the plasticity induced by STDP and phase precession. The premise that such separation is needed to learn multiscale successor representations is compatible with other theoretical accounts for this ordering. Specifically Momennejad and Howard [39] showed that exploiting multiscale successor representations downstream, in order to recover information which is 'lost' in the process of compiling state transitions into a single successor representation, typically requires calculating the derivative of the successor representation with respect to the discount parameter. This derivative calculation is significantly easier if the cells – and therefore the successor representations – are ordered smoothly along the hippocampal axis."