Dear all, thank you in earnest for the detail provided in each of the reviews. We have addressed in as much detail as possible what seem to be the major points. Where we have not replied to a specific point, please assume that we agree and will correct in the final manuscript.

1. Motivation for work We **agree that a more detailed discussion/broader citations of existing work** will strengthen the case for our proposed model (R1-3). As noted (R2), there exists much excellent work modelling spatial processing. Specifically, work implementing SLAM in a biological context (Milford et al., 2004; 2008), posing GCs as an eigendecomposition of the transition structure of the environment (Stachenfeld et al., 2014; 2017; i.e. a HPC to GC mapping (R2); **we will also include relevant experimental citations**) and proposing a method for correction of GC activity by sensory inputs (e.g. Fuhs and Touretzky, 2006) we consider particularly relevant. However, **none of these works discuss explicit probabilistic processing or representations**, which is clearly important to performing robust inference under uncertainty. Ours is also the **first model of distortions** to the grid pattern, a subject of intense current interest to the field (Krupic et al., 2018; Hagglund et al., 2019). If accurate, our model would be a significant advance in the understanding of how the brain encodes space.

Neither do these models link to replay. Of existing models, **none account for coordinated mEC-HPC activity (R2)**, only one gives normative insight (Mattar and Daw, 2018) and almost all focus on reward processing. We believe the latter point emphasizes an important contribution of our work; recent work casts some doubt on the strength of the specific predictions of RL based hypotheses (e.g. Stella et al., 2019). Not only does our theory pose an important alternative, but it makes testable predictions. As correctly noted (R1), we are explicitly making the point that rewarded locations are also likely to be locationally informative; an apple is highly indicative of location in an otherwise featureless maze. We are not claiming that reward is an explicit sensory stimulus; **we agree that this point is clumsily made in the text**. Lastly, that replay is observed in behavioural states such as sleep, immobility or at choice points in decision making tasks, **implies that animals must rely on another mechanism for online localization (R1)** - as we propose.

2. Contribution to AI More importantly perhaps, we believe that our work is relevant to the AI community as a technical "bridging" (R3) document and **aligns with several strands of active research**, in particular the graph network (GN) community. It has been suggested that grid cells represent and eigendecomposition of the Laplacian of a graph-like representation of environmental states (Stachenfeld et al., 2014). Take together with our work, this suggests that the brain might perform hierarchical asynchronous message passing on low-dimensional embeddings of such a graph. Whereas there have been several recent paper on learnable attention modules (Velickovic et al., 2018) and other methods for reducing the complexity of graph convolutions (e.g. Bruna et al., 2014; Kipf and Welling, 2017), to our knowledge, there is no existing work on asynchronous processing in GNs (despite the success of the mechanism employed in our paper, although originally proposed by Elidan et al., 2012, which **we will make more clear where relevant (R3)**). We would hope to stimulate serious discussion, given evidence for these processes in the brain and **supporting theory** (our work and others), as to whether these approaches are worth investing more energy in investigating, especially given current trends in neuromorphic hardware and the increasing need for low-power solutions in e.g. mobile devices. Notwithstanding the above, and although we do not believe these are the primary contributions of this paper, to our knowledge it is the first to detail mathematically the process of performing belief propagation (BP) in **periodic manifolds with hexagonal** symmetry (R3; there are clear cases where this may be better than using square bases, c.f. Hoogenboom et al. 2018), or to use principled message scheduling to selectively update the map in a SLAM context.

3. Outline of contributions Given the above, we will improve the introduction to **a. State more clearly what results are obtained and what analyses are performed (R1/3)** and **b. Outline their relevance the AI/Neuro fields (R3)**.

4. Mathematical notation We thank R2 for their detailed comments. **We will correct what we agree is abusive use** of *mathcal{N}* and apologise if it caused difficulty in interpreting the methods. Likewise regarding the mixed usage of scalar functions and vectors. Initially, we felt it important to abstract as much as possible the technical details from their mechanistic implementation. On re-reading however, we appreciate that it causes the notation to become unnecessarily unwieldy. **As suggested, we will specify the mapping and continue therein with linear algebra notation**.

5. Structure of methods We thank R2 for their suggestion and **will restructure if it will improve the clarity of the work**. To be clear, we do not consider how the metric space (i.e. the GC-GC connectivity) is formed (although c.f. Stachenfeld et al., 2014; Whittington et al. 2018). Rather, the metric space is fixed, but the associative map (PC-PC connections) and the associative-metric mapping (GC-PC) are learned. Our hypothesis is that, given no sensory drive to the grid cells, the agent has some weak prior notion of 'path integration' that would be based on self-motion cues, which is used by the *online* system to form priors over the encoding of encountered landmarks. The role of the *offline* update is then to correct this encoding, given the joint distribution over all landmarks. In this way, separated clusters of associated landmarks are still embedded proximally in the metric space by virtue of their encoding via path integration. The distortions arise from the notion that the self-motion estimate of 'distance travelled' through metric space is only approximate; in fact, we believe that it is affected by information flow, which is captured in the pairwise landmark distances (PC-PC connections) and ultimately induces distortions in the readout.