Explaining heterogeneity in medial entorhinal cortex with task-driven neural networks

Anonymous Author(s)
Affiliation
Address
email

Abstract

Medial entorhinal cortex (MEC) supports a wide range of navigational and memory related behaviors. A series of well-known experimental results have revealed the existence of specialized cell types within MEC — e.g. hexagonal grid cells, boundary cells, head-direction cells — whose highly stereotypical response profiles are suggestive of the role they might play in supporting MEC functionality. However, the majority of MEC neurons do not exhibit such cleanly stereotypical firing patterns, and have so far defied categorization within a succinct tuning curve classification scheme. It is an open question how to describe the response profiles of these more "heterogeneous" cells, and whether or how they contribute to behavior. In this work, we took a computational modeling approach to addressing these questions. We first performed a statistical analysis that shows that heterogeneous MEC cells are just as consistent in their distinct response patterns as the more stereotypical grid cells, suggesting that that they have a coherent functional role, and are meaningful targets for explanation. Next, we evaluated a spectrum of candidate models in terms of their ability to describe the response profiles of both stereotypical and heterogeneous MEC cells. We found that recently developed neural network models optimized to solve path integration tasks are dramatically better than traditional grid cell-centric models at matching most MEC neuronal response profiles, despite not being explicitly trained for this purpose. Specific choices of network architecture (such as gated nonlinearities and an explicit intermediate place cell representation) have an important effect on the ability of the model to generalize to novel scenarios, with the best of these models closely approaching the reliability ceiling of the data itself. We then performed a virtual knock-out experiment for multiple cell populations within this model and evaluated path integration performance, finding that the heterogeneous cells are at least as important for task competence as grid cells, if not more so. Finally, inspired by recent data showing that, going beyond their spatial response selectivity, MEC cells are also responsive to non-spatial rewards, we introduce a new MEC model that performs reward-modulated foraging, in the context of the path-integration task. We find that this unified task-optimized model matches neural responses across all reward and spatial conditions. Taken together, our results suggest that a process of biological performance optimization directly shaped the neural mechanisms in MEC as a whole, and provides a conceptually fruitful program to guide the way for future experimental efforts.

1 Introduction

From exploring new areas, planning shortcuts, and returning to remembered locations, the ability to self-localize within an environment subserves a range of navigational behaviors that are essential

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

for survival. The hippocampus (HPC) and medial entorhinal cortex (MEC) are known to contain cells that encode the position of an animal by displaying firing fields with strikingly regular response patterns, influenced by self-motion [O'Keefe and Dostrovsky, 1971, Hafting et al., 2005, Kropff et al., 2015, Solstad et al., 2008, Sargolini et al., 2006]. For example, MEC grid cells possess characteristically symmetric and periodic tuning curves, yielding hexagonal arrays of neural activity over physical space. Additionally, border cells that fire maximally near environmental boundaries and head direction cells that fire only when an animal faces a particular direction are among the other MEC cell types with interpretable tuning curves.

However, a large fraction of the MEC population have unconventional and heterogeneous tuning to navigational variables and are less obviously well-described in terms of simple, stereotypical tuning patterns [Hinman et al., 2016, Hardcastle et al., 2017]. How can we characterize what these other, more heterogeneous, populations of cells do? Are they critical to the abilities of the animal in real-world memory tasks? And if so, how is their role distinct from the more stereotypical grid-like cells? Here, we take a quantitative computational modeling approach to answering these questions.

To start with, we address the question of whether there is a phenomenon to model in the first place, and how one would measure model accuracy quantitatively. To this end, we identify the *similarity transform* between neural populations in different animals — that is, a mapping which takes neuronal population vectors in a "source" animal and maps it to corresponding neuronal population vectors in a "target" animal. To the extent that robust similarity transforms can be found that match up MEC population responses across multiple animals measured in multiple experimental conditions, then a reliable pattern of neural response behaviors (across conditions) has been identified. This identification strategy is well-defined even when there is no known *a priori* taxonomy of functional response types. We find here that with the right mapping class (linear regression with ridge regularization), the inter-animal consistency of MEC neuronal responses is in absolute terms very high (> 0.9), and in relative terms just as high as for heterogeneous cells as for more stereotypical grid cells.

Using this same similarity transform, we then evaluate the ability of each of multiple computational models to explain response variance of MEC neurons, treating each candidate model as a potential "source animal" and measuring how well it maps to each target real animal. We look to the recent literature to identify potential candidate models. Over the past several decades, collaborations between experimental and computational neuroscientists have led to the formulation of dynamical-systems models of grid cell formation, and helped illustrate possible functional roles for grid cells in supporting path integration-based hippocampal place cells [Skaggs et al., 1992, Zhang, 1996, Fuhs and Touretzky, 2006, Burak and Fiete, 2009]. These powerful models make a number of non-obvious predictions about MEC neural properties, some of which have been confirmed in subsequent experimental work [Ocko et al., 2018, Campbell et al., 2018]. Despite their success, such models are limited in their explanatory scope, hand-designed to capture the properties of one stereotypical cell-type class (e.g. grid or border cells) at a time, or combinations of several cell types via multiple dedicated type-specific modules [Couey et al., 2013, Yoon et al., 2013].

A potential solution to this problem arises out of recent work creating learned neural networks that achieve path integration [Cueva and Wei, 2018, Banino et al., 2018]. Intriguingly, these models have been found to contain internal units that resemble grid cells, suggesting that such stereotypical cells embody a computational solution to path integration that naturally arises from satisfying an end-to-end functional constraint. Recent work has demonstrated that the underlying mathematical reason for this fact is due to pattern forming dynamics under a nonnegativity constraint [Sorscher et al., 2019, 2020]. Intriguingly, in addition to having units that resemble stereotypical grid or border cells, these learned neural networks also naturally possess a wide variety of other less easily described unit types, raising the possibility that these artificial "heterogeneous" units might be somehow resemble the actual heterogeneous cells making up the majority of real MEC populations.

Motivated by these ideas, we generate a wide variety of candidate model networks by varying architectural structure and end-to-end optimization objectives, each expressing a different hypothesis for MEC circuit structure and function. Architecturally, we formulate variants based on using different types of nonlinearities and different types of local recurrent memory circuits (e.g. RNNs [Elman, 1990] and LSTMs [Hochreiter and Schmidhuber, 1997]). From a task point of view, we test both simple dimensionality reduction [Stachenfeld et al., 2014, Dordek et al., 2016] as well as place cell mediated vector path-integration [Banino et al., 2018] and direct position estimation [Cueva and Wei, 2018].

Our core result is that there is substantial variation in the models' abilities to explain MEC responses, especially those of the heterogeneous cells, as a function of a model architecture and task. Some models, such as the classic "Grid Cell" model based on low-rank decomposition of place cell fields, do a reasonable job explaining grid cell responses but are quite poor at explaining most other neurons. Task-optimized learned models typically do better, especially those optimized for place cell-mediated vector path-integration. The best model – with memory-gated rectified nonlinearities, introduced in Sorscher et al. [2020] – essentially *solves* the neurons, capturing nearly 100% of the noise ceiling of the data. This is of substantial interest, given that the nonlinear components of this model are not directly optimized to match neural data (just to solve the task), and given that a series of strong control models capture much less of the MEC neural variability. This same model also best generalizes to a variety of novel experimental conditions, and has the best match to the empirical data on a grid score distribution metric.

With this predictive model in hand, we then begin to address our second core question: what is the functional role of heterogeneous neurons? We generate several results suggesting that heterogeneous neurons are important for path integration, including a response-based virtual knockout experiment, in which we compare performance degradation when deleting grid cells as compared to heterogeneous units: strikingly, the heterogeneous knockout suffers more than the grid cell knockout at comparable knockout strengths, a prediction ripe for experimental corroboration.

Finally, building on the above results, and inspired by recent results showing that MEC cells modulate in response to extrinsic reward, we introduce a new MEC model that performs reward-modulated foraging, in the context of the path-integration task. We find that this unified task-optimized model matches neural responses across all reward and spatial conditions. Taken together, our results suggest that a process of biological performance optimization directly shaped the neural mechanisms in MEC as a whole, and provides a conceptually fruitful program to guide the way for future experimental efforts.

2 Reliability of heterogeneous cell response profiles

What firing patterns of MEC cell populations are common across multiple animals, and thus worthy of computational explanation? This question is comparatively straightforward for stereotypical MEC cells, because the very presence of these stereotypical features (e.g. hexagonal grids of a given orientation and spatial frequency) allows the definition and measurement of observables that arise reliably across trials and animals (e.g. the grid score distribution). But how can this be done generally for populations in which one does not have a prior characterization of what each cell encodes? We take inspiration from methods that have proven useful in modeling visual, auditory, and motor cortex [Yamins and DiCarlo, 2016, Kell et al., 2018, Michaels et al., 2020]. Specifically, we aim to identify the *narrowest* class of similarity transforms needed to map the firing patterns of one animal's MEC population to that of another (Fig. 1a). As with other cortical areas, this transform class likely cannot be so strict as to require fixed neuron-to-neuron mappings between MEC cells, since even within the same animal at different times, MEC and HPC populations can undergo remapping that shift cell responses across the population [Farhoodi et al., 2020, Low et al., 2020]. However, the transform class for MEC also cannot be so loose as to allow a completely unconstrained linear mapping, since the highly structured response patterns of stereotypical MEC cell types such as grid cells may not be guaranteed to be preserved under arbitrary linear transforms.

We explore a variety of mapping transform classes between the population rate maps of MEC neurons in mice (12 animals; $n = 620$ cells) undergoing open foraging for randomly scattered crushed cereal in a $100cm^2$ 2D arena [Mallory et al., 2021]. For all methods, mapping fit is performed on a subset of spatial positions, and evaluated shown for a held-out subset of positions (Fig. 1(b), see supplement for more details). For the strictest one-to-one mapping transform, each target unit is mapped to the single most correlated unit in the source animal. One-to-one mapping yielded the lowest inter-animal consistency among the maps considered, capturing barely 50% of target neural response variability. Linear regression with sparseness priors, such as Lasso and ElasticNet, also proved to be too strict, yielding no improvement over the one-to-one mapping. However, Ridge (L2-regularized) regression is highly effective at inter-animal mapping, recovering nearly all target variability for most neurons.

Different target cells required different numbers of source units to achieve effective mapping (Fig. 1(c)). Moreover, as might be expected from the intuitive notion of heterogeneity, simpler

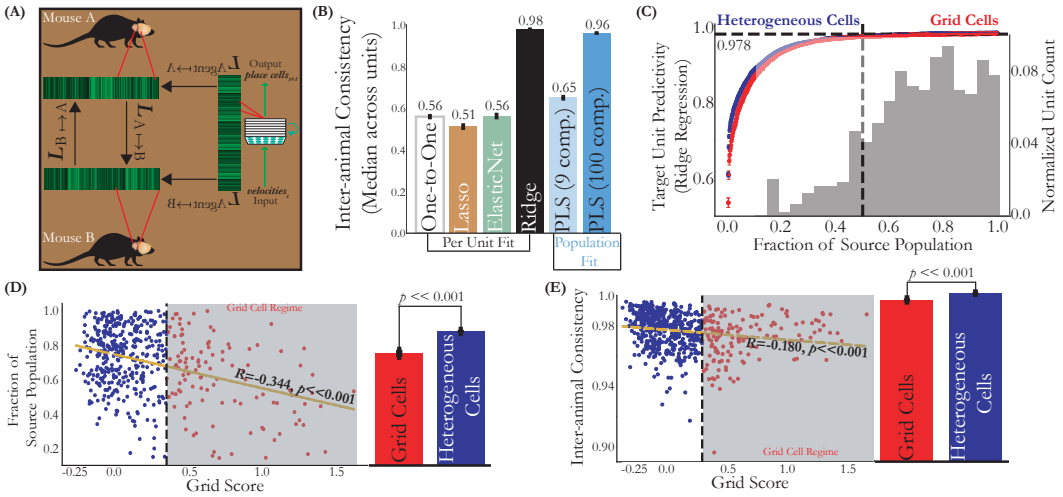

Figure 1: **Heterogeneous cells are just as reliable as grid cells across 2D environments and task conditions. (A)** Inter-animal consistency mapping procedure, where L denotes the use of a transform which is at most linear in functional complexity. The agent (right) is then mapped to each animal's MEC population separately. The agent is always evaluated in the same arena the animal was in, though its weights can be optimized in a different environment. **(B)** Inter-animal consistency across different mapping transform classes. Median and s.e.m. across 620 cells. **(C)** Per unit predictivity under ridge regression as a function of the size of the source population. The horizontal dashed line indicates the median inter-animal consistency under ridge regression computed in (B). Overlaid is the density of the minimum number of source units needed to reach the inter-animal consistency for a given cell. **(D)** The minimum number of source units per target cell plotted against its grid score. Red denotes grid cells with grid score > 0.3 and blue denotes heterogeneous cells with grid score ≤ 0.3. The bar plot on the right is the median and s.e.m. of this quantity for the identified grid and heterogeneous cell populations. **(E)** Inter-animal consistency of each unit plotted against its grid score. The bar plot on the right is the median and s.e.m. of this quantity for the identified grid and heterogeneous cell populations.

cells with more stereotypically grid-like response patterns typically require fewer source cells to capture (Fig. 1(d)). However, heterogeneous cells did not have lower inter-animal consistency than grid cells (Fig. 1(e)). Taken together these results show that the heterogeneous non-stereotyped cell populations are reasonably similar across animals – at least up to (ridge-regularized) linear transform – establishing that there is a reliable target pattern to study in the first place.

3 Task-optimized models of MEC spatial response variation

Evaluating a spectrum of candidate models. We evaluated models of several basic types. First, we considered learnable neural networks, all of which accept a stream of two-dimensional velocity input vectors, have a single layer of hidden neurons identified as the putative MEC population, and which are optimized to perform some form of path integration readout on simulated motion paths (Erdem and Hasselmo [2012]) in a fixed-sized arena. The networks varied according to the nature of their local recurrent cell structure (ungated RNN vs. LSTM), activation functions (Linear vs. Tanh vs. ReLU), and output objective function (explicitly constructing a set of simulated place field neurons [Banino et al., 2018] or directly performing two-dimensional position integration [Cueva and Wei, 2018]). We ensured that comparisons were fair by equalizing the size of the hidden layer across models, and included an architecture-only control with untrained filter weights.

We also implemented a class of models describing MEC activations as Non-negative Matrix Factorization (NMF) operating on simulated place cells. The lowest-rank version acts a "Grid Cell" model, inspired by recent work positing MEC as a low-dimensional embedding of hippocampal place fields [Stachenfeld et al., 2014, Dordek et al., 2016], and further validated by the mathematical theory developed in Sorscher et al. [2019, 2020] to explain the emergence of grid cells. Higher-rank NMF models titrate between the low-rank Grid Cell model and a full-rank control that measures how well MEC cells can be explained as a linear projection of their putative place cell outputs.

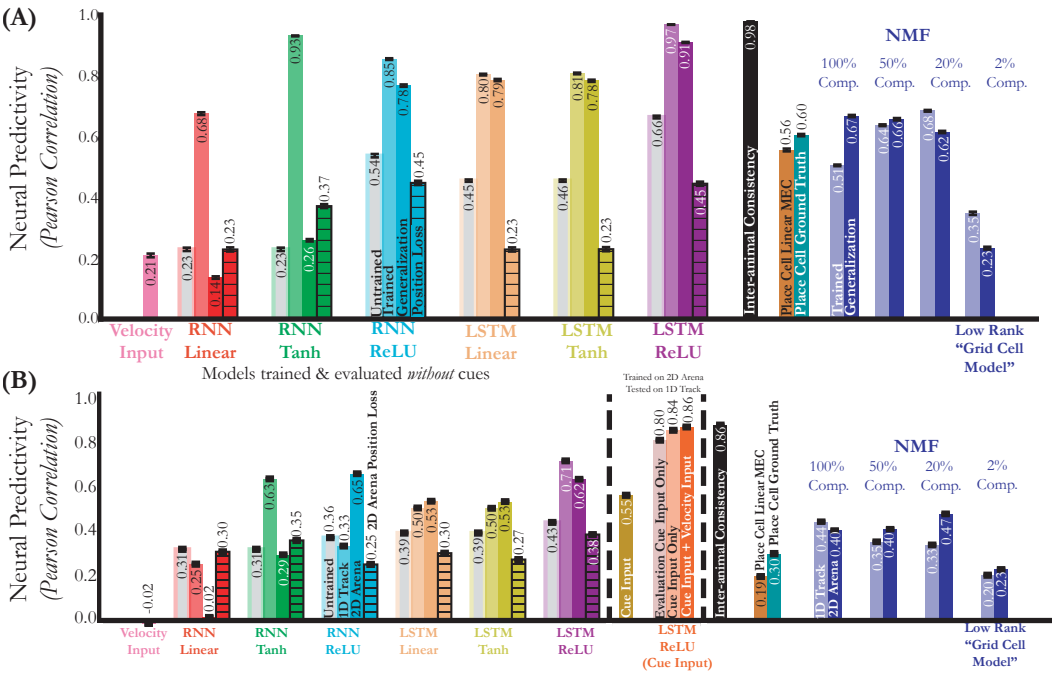

Figure 2: **Task-optimized navigational models best predict the entire MEC population.** **(A)** Neural predictivity of the model MEC units to the real MEC responses of 12 animals in a $100cm^2$ 2D open field under the ridge regression transform. Median and s.e.m. across 620 total units. **(B)** Same as (A), but now the models are evaluated against responses of MEC units while the animal is traversing a $400cm$ 1D track. "1D Track" refers to models trained to path integrate on this same 1D track. "2D Arena" refers to models trained to path integrate on the $2.2m^2$ arena and evaluated on the 1D track. Median and s.e.m. across 2861 units from 8 different animals.

For evaluation, we map each model's proposed MEC activations to empirically measured firing patterns of the real neurons, using the same type of cross-validated ridge-regression transform used to measure inter-animal similarity as in Section 2. The predicted neural responses under the mapping are then compared to real target neural responses on a neuron-by-neuron basis, and the median of accuracy of these predictions taken over target neurons. In addition to evaluating model-data match when constructed on arenas of the same size in which the neural data were collected ($100cm^2$), we also evaluated models trained on arenas of a larger size ($2.2m^2$) but tested on the $100cm^2$ arena.

Neural prediction results. The results of our model evaluation, shown in Fig. 2(a), support several inferences:

- Different models are substantially different in their ability to predict neural responses. MEC electrophysiology data collected during 2D open-field foraging is thus a strong model target that effectively separates candidate models from each other.
- Task-optimized models are reliably better than their untrained controls, across all architecture types and objective functions.
- Rectification is substantially better than Linear or Tanh activation, especially for promoting generalization to new arena sizes.
- The gated LSTM nonlinearity improves model fits compared to the simple ungated alternative, across activation types and objective functions.
- The explicit place cell construction task leads to substantially better fits than the direct two-dimensional path integration task ("Position Loss"), even though both were reported in the literature to create grid cell-like units. Enforcing a place field representation at the output of the network thus appears to be an important constraint in order to recapitulate responses in MEC.
- The Grid Cell (low-rank NMF) model is a poor fit to the MEC population overall.
- Full-rank NMF and place cell linear MEC output-based controls capture approximately half the explainable variance of MEC neurons, significantly more than the simple velocity linear input control, but substantially less than any of the input-driven ReLU and LSTM-based RNNs.

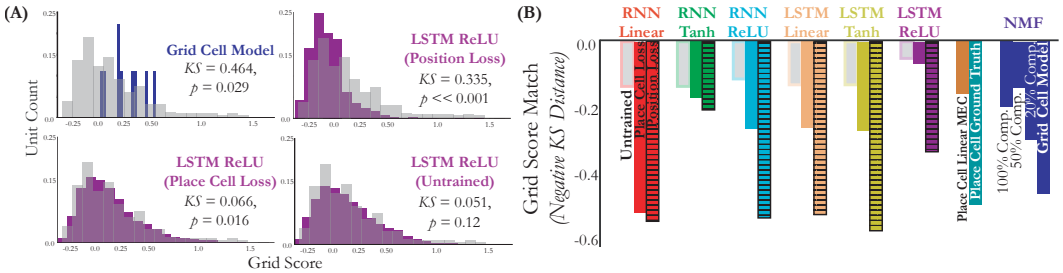

Figure 3: **Relationships to grid score distribution.** **(A)** Example distributions of grid scores from selected models against the ground truth distribution in the neural data (yellow). Negative Kolmogorov-Smirnov distance between the distribution of grid scores of the model units to those of the 620 units in the data. The models have all been trained in a different $2.2m^2$ arena, but are evaluated in a $100cm^2$ environment that the animals are in. **(B)** Quantification across all models.

- The best model is the LSTM ReLU model, first introduced in [Sorscher et al., 2019]. This model captures nearly 100% of the explainable variance of the MEC neural population response when trained for place cell construction on the arenas of the same size as that on which neural data was collected.

Generalization to novel experimental conditions. If a model is truly correct, then once trained, it should capture neural responses in any new tested condition. We tested generalization in two ways. First, as shown in the second from the right bars of Fig. 2a for each model architecture class, we performed a 2D arena size generalization test by constructing models on one arena size, testing against neural data on another. We found this generalization test gives essentially the same rank-order comparison as in the same-arena-size test, with one striking exception: the RNN Tanh model performs well within arena size, but fails to generalize. The LSTM ReLU model again performs the best, capturing approximately 93% explained variance of the neural responses on the novel arena size.

Second, we also evaluated the same 2D-pretrained models by running them on 1D track, comparing models to neural data collected from mice in a 1D virtual reality setup, using the same Ridge-regression mapping procedure as in the 2D comparisons (see supplement for more experimental details). We found (Fig. 2b) that model results were in similar rank order as for the original 2D results, with both ReLU-based 2D-pretrained models capturing approximately 72% of the explainable variance, and the LSTM ReLU model achieving the best match (83%) when trained for place-field construction directly in the 1D environment. We hypothesized that the larger gap between the best model and the inter-animal noise ceiling, as compared to the 2D case, was due to the fact that during the 1D experimental data collection, mice were also also presented with visual cues. MEC neurons are known to respond to visual input, but as none of our evaluated models had a visual front-end, the could not respond accordingly. To test this hypothesis we trained the LSTM ReLU model in 2D, but with phantom visual cue locations added. We then evaluated this model against 1D neural prediction (with the real visual cue locations as input), and found that including cues rescued model performance back to essentially the noise ceiling, as in 2D (Fig. 2b, orange bars).

Assessing grid score distribution match. The *grid score* is a metric of how stereotypically grid-like the response pattern of a given unit is (see supplementary material for specific definitions), with "grid cell" typically defined as having a grid score of greater than 0.3. The ground-truth distribution of grid scores for cells in an unbiased experimental population (gray bars in each subpanel of Fig. 3a) characterize the extent to which real MEC populations are non-stereotypical in their responses. To further assess model accuracy, we also compared the distribution of grid scores within each model to that of the real data, using the (negative) Kolmogorov-Smirnov (KS) distance of the empirical and model grid score distributions as a quantitative metric. This metric is both stronger and weaker than the mapping accuracy metric used above — stronger in the sense that, since there is no parameter fitting in the metric, to fit it well a model has to have the correct distribution in its raw feature output; and weaker in that it only assesses cells on one component of their profile (i.e. "grid-ness").

We observed that the low-rank Grid Cell model has poor fit on this metric, essentially because it contains *too many* grid cells (see Fig. 3a, upper left). In contrast, the same model that achieves best fits on the neural-fit metric (LSTM ReLU) also achieves the best on the grid cell distribution match metric, and is the only architecture that cannot be distinguished from the ground truth at the

$KS = 0.1$ significance level. Across models architectures, KS-distance was generally in line with neural regression fit metrics (see Fig. 3b and Fig. S2). One key difference between the metrics, however, is that for most architectures, the untrained filters provided *better* grid score distribution matches as models with task-trained parameters — except for the LSTM ReLU architecture, where the trained place cell construction model is as good as its untrained counterpart — and models trained for direct path integration were especially poor on this metric. These results suggest the two metrics in Figs. 2 and 3 are complementary which aspects of model correctness they address, and together help zero in on the most effective models overall.

4 Predicting the functional relevance of heterogeneous cells

Correlating task performance and neural predictivity. To begin to address the question of the functional relevance of heterogeneous cells, we first looked at the overall correlation between model task performance and neural fit (Fig. 4a). Though the correlation is imperfect (and somewhat hard to assess given the small number of datapoints), the empirical relationship is comparatively strong, suggesting that improved task performance may be causally related to the ability capture neuron response patterns across the population. In contrast, there is no clear relationship between grid score distribution match and model performance (Fig. 4b), illustrating the importance in matching cell properties other than "grid-ness" in predicting task performance.

Relative predictivity gain for heterogeneous cells. The above result is put into greater perspective by comparing neural predictivity differential between a task-trained model and the Grid Cell model, on a per-neuron basis, as a function of grid score (Fig. 4c). The task-trained model has improved neural predictivity relative to the Grid Cell model for both grid cells (grid score > 0.3) and heterogeneous cells, but the improvements on the latter are substantially larger than the former. This again suggests that the heterogeneous cells are playing a substantial role in allowing the trained model to achieve improved performance.

Virtual knockout experiment. To address the question most directly, we performed a cell-type-targeted virtual knockout comparison experiment. In doing this, we used the LSTM ReLU-place cell model shown in the previous section to match data in a generalizable way across multiple complementary neural metrics. We identified the units in the trained model with high grid score (> 0.3), and ablated varying fractions of these units ranging from near-complete knockout to near-complete preservation, while measuring task generalization performance (see supplement for details of this knockout procedure). We similarly ablated matched numbers of heterogeneous units (grid score ≤ 0.3), again measuring model performance. The result of this comparison (Fig. 4d) shows, perhaps unexpectedly, that the heterogeneous knockout is more injurious to model performance than grid cell knockout, across the entire range of knockout strengths. This result held not only in the case of generalization to a new arena size as show in Fig. 4d, but also in the same $2.2m^2$ arena the model was trained in (Fig. S4). It would be of substantial future interest to confirm or reject this model prediction with a real targeted knockout experiment *in vivo*.

5 Modeling reward-driven remapping in MEC

Recent work in both rodents and humans has uncovered that MEC and HPC neurons represent not only literal space, but also capture spatialized layouts in a more abstract sense in modalities other than spatial position [Constantinescu et al., 2016, Aronov et al., 2017]. It has also been seen [Butler et al., 2019, Boccara et al., 2019] that non-spatial rewards can influence the shape of MEC response profiles, restructuring them to incorporate the location of the learned reward. These intriguing phenomena represent a natural direction for modeling, but are not captured by any of the neural network models described in Section 3, as they do not have reward-modulated inputs. (Note that while Banino et al. [2018] used their pretrained network – the LSTM Tanh (place cell loss) model evaluated above – to perform downstream reinforcement-based navigation tasks, this does not address the neural modeling question raised here, since [Banino et al., 2018] does not input reward state to their MEC model or propagate reward outcomes back into MEC representation.)

To tackle this problem, we obtained neural data (from [Butler et al., 2019]) for animals in conditions both with reward (**reward+**) and without it (**reward-**), whereby the animal navigates to a $20cm^2$ reward zone within a $150cm^2$ arena to receive 0.5-1 units of cereal. We next performed inter-animal

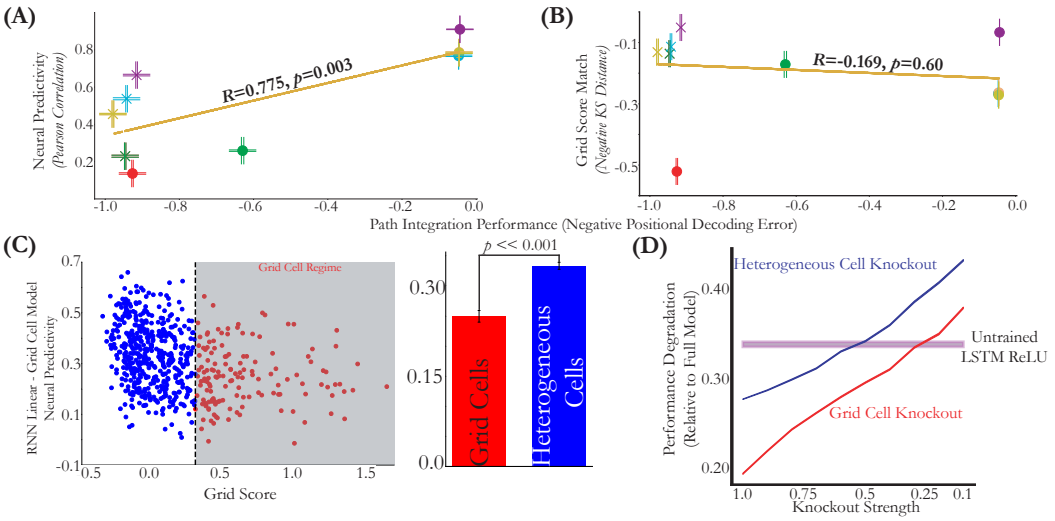

Figure 4: **Heterogeneous cells are relevant to navigation.** **(A)** Neural predictivity (median and s.e.m. across 620 units) versus path integration performance (mean and s.e.m. across 100 episodes), measured by negative positional decoding error, of the neural network models. The models are either untrained ("X") or trained with the place cell loss ("O") in the $2.2m^2$ arena. Positional decoding error is measured by taking the top 3 most active place cell outputs in each model, evaluated in the $100cm^2$ arena. **(B)** Same as (A), but with grid score distribution instead of neural predictivity. **(C)** (*Left*) Per unit neural predictivity difference between the RNN Linear and Grid Cell Model trained and evaluated on the same $100cm^2$, plotted against that unit's grid score. (*Right*) Quantification of this difference aggregated across the grid cell and heterogeneous cell populations, respectively (mean and s.e.m). **(D)** Normalized performance degradation as evaluated on the $100cm^2$ arena, relative to the full LSTM ReLU trained on the $2.2m^2$ arena with the place cell loss. We identify 895 units in this trained model with grid score > 0.3 and knockout the same number of heterogeneous cells (randomly sampled 100 times), to yield the "heterogeneous cell knockout". The x-axis denotes "knockout strength", namely, the fraction of active units for each knockout type (equalized to be the same amount in terms of absolute unit count). Mean and s.e.m. over 100 evaluation episodes.

consistency checks, using the same method as we did with the purely spatial-condition data in Section 2 above, finding that inter-animal consistency for neural responses is high both within reward condition (e.g. just spatial modulation in **reward+** and **reward-** separately), and across both spatial and reward-modulated conditions overall (Fig. S5).

We first evaluated the models (trained in the $2.2m^2$ arena) described in the previous sections. While the best of these models was able to capture neural responses in the **reward+** and **reward-** separately (Fig. 5b, hatched bars), neither this model nor any of the others described above achieved effective neural response prediction when asked to predict response patterns across both reward conditions – unsurprisingly, since the models do not accept reward-modulating input. Simply augmenting the network input to receive an additional binary reward (Figure 5b,"Reward Input"), but without using that input to specifically modulate the task, did not substantially improve neural predictivity, showing that reward must somehow be explicitly used during model training. We thus sought to build new network models that respond to the existence of an extrinsic reward in a behaviorally-meaningful fashion. One natural approach would be to build a full reinforcement-learning (RL) setup in which reward state is which reward-modulating environmental features (e.g. visual presence of food units), together with reward state, are integrated back into the MEC network and allowed to modulate neural responses. It is quite reasonable that training a model to optimally forage in this end-to-end way could be successful in creating a representation that matches MEC reward-response modulation.

However, we took a simpler approach: simply change the behavior of the agent in response to reward in the way RL training would be expected to do if successful. In our case, this simply meant modifying the foraging trajectories of the animal from the typical random foraging path used in earlier models (a kind of pure exploration condition) to either purely exploit a known reward resource (e.g. by directly moving to the reward location when it is present), or to bias it toward the reward while still maintaining some exploratory behavior. In modeling these scenarios, we sought to roughly

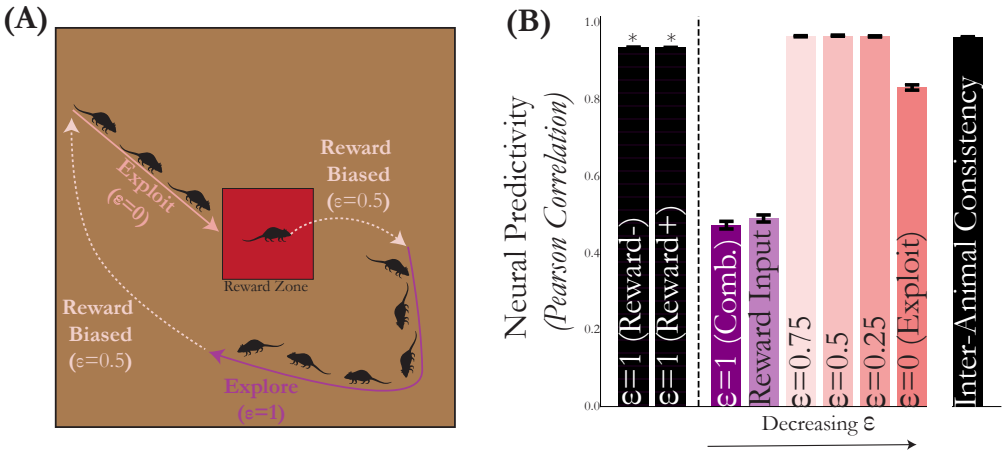

Figure 5: **Reward biased path integration captures remapping of responses in the presence of a reward.** **(A)** Schematic of the three conditions. "Explore" is the original set of random bout trajectories used to train the agent as before. "Exploit" involves the agent navigating directly to the reward zone (red box) within a fixed number of timesteps per episode (7), and spending the remaining 13 timesteps only path integrating in the reward zone. "Reward Biased" refers to training with each of the above two trajectories on $0 < \varepsilon < 1$ of the training episodes. **(B)** Neural predictivity of the LSTM ReLU (place cell loss) model, evaluated in the $150cm^2$ arena. Median and s.e.m. of 598 cells from 7 animals. $*$ denotes comparison of the model to each condition separately.

mimic experimental observations showing animals often take rapid and direct paths to the reward zone [Butler et al., 2019].

We trained the LSTM ReLU (place cell loss) model with a variety of foraging policies titrating between pure exploration and exploitation (see Fig. 5a, and supplement for details of foraging policy implementation). Throughout, the actual training task of the network remained the same as above — accuracy in path integration as measured by place cell construction — but the network now had to use its reward-presence input to signal the MEC-layer representation to modulate its response to maintain positional knowledge in the face of reward-based velocity input modulation.

We found that the agent trained with a purely exploitative policy does substantially better in matching neural responses across reward conditions than original random-walk model, while an optimal mixture of exploration and exploitation leads to networks with even larger neural match improvements, reaching the noise ceiling of the data itself (Fig. 5b). Overall, we found the results robust to a mixture of exploration and exploitation, so long as there was a non-zero amount of exploration, which performed best. We found that a positive differential between neural predictivity on heterogeneous cell vs. grid cell population responses than compared to the random walk model (Fig. S6a). Moreover, there was a *greater* performance degradation when we ablated heterogeneous cells as compared to grid cells in these reward-modulated networks (Fig. S6b) , suggesting that these heterogeneous cells may play a comparatively greater role in reward-driven navigation than grid cells.

6 Discussion

Our specific findings suggest that task-driven RNNs with certain architectural motifs effectively explain neural responses in medial entorhinal cortex (MEC), and that heterogeneous cells may play an important role in MEC-mediated animal behaviors. More generally, our results suggest that using computational models to go beyond traditional "easy to visualize" neural response profiles may be a fruitful approach for studying navigation and memory in the MEC, and beyond.

Going forward, on the experimental design front, our results suggest that moving beyond landmark-free 2D arenas to cue-rich, high-variation 2D (rather than 1D) environments will further push model development and evaluation. It will be of special interest to measure neural responses in environments and conditions where animal behavior exceeds today's best neural navigation and memory networks [Eichenbaum et al., 1999], where large gaps in model-data match would likely help both further our future efforts to understand the brain, and focus efforts to improve neural network design.

References

D. Aronov, R. Nevers, and D. W. Tank. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, 543(7647):719–722, 2017.

A. Banino, C. Barry, B. Uria, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick, T. Degris, J. Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557 (7705):429–433, 2018.

C. N. Boccara, M. Nardin, F. Stella, J. O'Neill, and J. Csicsvari. The entorhinal cognitive map is attracted to goals. *Science*, 363(6434):1443–1447, 2019.

Y. Burak and I. R. Fiete. Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol*, 5(2):e1000291, 2009.

W. N. Butler, K. Hardcastle, and L. M. Giocomo. Remembered reward locations restructure entorhinal spatial maps. *Science*, 363(6434):1447–1452, 2019.

M. G. Campbell, S. A. Ocko, C. S. Mallory, I. I. Low, S. Ganguli, and L. M. Giocomo. Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. *Nature neuroscience*, 21(8):1096–1106, 2018.

A. O. Constantinescu, J. X. O'Reilly, and T. E. Behrens. Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292):1464–1468, 2016.

J. J. Couey, A. Witoelar, S.-J. Zhang, K. Zheng, J. Ye, B. Dunn, R. Czajkowski, M.-B. Moser, E. I. Moser, Y. Roudi, et al. Recurrent inhibitory circuitry as a mechanism for grid formation. *Nature neuroscience*, 16(3): 318–324, 2013.

C. J. Cueva and X.-X. Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. *ICLR 2018, arXiv:1803.07770*, 2018.

Y. Dordek, D. Soudry, R. Meir, and D. Derdikman. Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *Elife*, 5:e10094, 2016.

H. Eichenbaum, P. Dudchenko, E. Wood, M. Shapiro, and H. Tanila. The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron*, 23(2):209–226, 1999.

J. L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.

U. M. Erdem and M. Hasselmo. A goal-directed spatial navigation model using forward trajectory planning based on grid cells. *European Journal of Neuroscience*, 35(6):916–931, 2012.

S. Farhoodi, M. Plitt, L. Giocomo, and U. Eden. Estimating fluctuations in neural representations of uncertain environments. *Advances in Neural Information Processing Systems*, 33, 2020.

M. C. Fuhs and D. S. Touretzky. A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience*, 26(16):4266–4276, 2006.

T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005.

K. Hardcastle, N. Maheswaranathan, S. Ganguli, and L. M. Giocomo. A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. *Neuron*, 94(2):375–387, 2017.

J. R. Hinman, M. P. Brandon, J. R. Climer, G. W. Chapman, and M. E. Hasselmo. Multiple running speed signals in medial entorhinal cortex. *Neuron*, 91(3):666–679, 2016.

S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

A. J. Kell, D. L. Yamins, E. N. Shook, S. V. Norman-Haignere, and J. H. McDermott. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*, 98(3):630–644, 2018.

E. Kropff, J. E. Carmichael, M.-B. Moser, and E. I. Moser. Speed cells in the medial entorhinal cortex. *Nature*, 523(7561):419–424, 2015.

I. I. Low, A. H. Williams, M. G. Campbell, S. W. Linderman, and L. M. Giocomo. Dynamic and reversible remapping of network representations in an unchanging environment. *bioRxiv*, 2020.

C. S. Mallory, K. Hardcastle, M. G. Campbell, A. Attinger, I. I. Low, J. L. Raymond, and L. M. Giocomo. Mouse entorhinal cortex encodes a diverse repertoire of self-motion signals. *Nature communications*, 12(1):1–20, 2021.

J. A. Michaels, S. Schaffelhofer, A. Agudelo-Toro, and H. Scherberger. A goal-driven modular neural network predicts parietofrontal neural dynamics during grasping. *Proceedings of the National Academy of Sciences*, 117(50):32124–32135, 2020.

S. A. Ocko, K. Hardcastle, L. M. Giocomo, and S. Ganguli. Emergent elasticity in the neural code for space. *Proceedings of the National Academy of Sciences*, 115(50):E11798–E11806, 2018.

J. O'Keefe and J. Dostrovsky. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 1971.

F. Sargolini, M. Fyhn, T. Hafting, B. L. McNaughton, M. P. Witter, M.-B. Moser, and E. I. Moser. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762, 2006.

W. E. Skaggs, B. L. McNaughton, K. M. Gothard, and E. J. Markus. An information-theoretic approach to deciphering the hippocampal code. In *Proceedings of the 5th International Conference on Neural Information Processing Systems*, pages 1030–1037, 1992.

T. Solstad, C. N. Boccara, E. Kropff, M.-B. Moser, and E. I. Moser. Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909):1865–1868, 2008.

B. Sorscher, G. Mel, S. Ganguli, and S. A. Ocko. A unified theory for the origin of grid cells through the lens of pattern formation. *Advances in Neural Information Processing Systems*, 32, 2019.

B. Sorscher, G. C. Mel, S. A. Ocko, L. Giocomo, and S. Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *bioRxiv*, 2020.

K. L. Stachenfeld, M. Botvinick, and S. J. Gershman. Design principles of the hippocampal cognitive map. *Advances in neural information processing systems*, 27:2528–2536, 2014.

D. L. Yamins and J. J. DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356–365, 2016.

K. Yoon, M. A. Buice, C. Barry, R. Hayman, N. Burgess, and I. R. Fiete. Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature neuroscience*, 16(8):1077–1084, 2013.

K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *Journal of Neuroscience*, 16(6):2112–2126, 1996.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes]
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]