Sampling-based inference and predictions in the hippocampus.

OTKA PD-125386 Final Report

Balázs B Ujfalussy
Laboratory of Biological Computation and Laboratory of Neuronal Signalling, Institute of Experimental Medicine, Budapest, Hungary

Predicting the possible future consequences of actions is necessary for efficient planning. In volatile or partially observed environments animals can capitalise on representing the uncertainty associated with those predictions but whether and how uncertainty is represented during planning, and more specifically during navigation has remained elusive.

In this project, we investigated the neuronal representations in the hippocampus associated with planning during a goal-oriented spatial navigation task. We found no evidence for the instantaneous representation of the uncertainty in the data: the population activity typically encoded single locations at any given time point. However, the variability between subsequent theta cycles is larger than that expected from only encoding the most likely trajectory, and is consistent with random sampling from hypothetical future trajectories. Thus, in the hippocampus uncertainty during short term planning is represented by encoding multiple alternative options sequentially.

We also analysed the contribution of dendritic nonlinearities to the spatially selective activity of hippocampal place cells. We found that linear integration with a single global dendritic nonlinearity can accurately predict the response of neurons to naturalistic synaptic input patterns. Moreover, we demonstrated that the contribution of dendritic nonlinearities is maximal when synaptic clusters, formed via local plasticity, are moderately large.

In the next sections I give a succinct description of the achievements of during the 3 years of the project. Further details can be found in our publications. As large part of the work is still in preparation for publication I included an edited version of the current version of the manuscript as an Appendix to this report. As the workplan of the PD-125386 and the FK-125324 grants were largely overlapping I clearly indicated in the text which part of the work belongs uniquely to one of the grants (PD) or to both of the grants (PD and FK), as I did it in the progress reports.

Trajectory sampling model (Aims 1-2, PD and FK).

The aim of our proposal was to identify how neuronal activity during hippocampal theta oscillations contributes to efficient planning. The central hypothesis of our proposal was that the hippocampus implements a probabilistic generative model of the environment and uses this model to combine past and current sensory information to infer the current state of the world and predict future events or spatial trajectories. We further assumed that the hippocampus is able to consider multiple future options by sampling independent hypothetical trajectories. In the *sampling encoding scheme* uncertainty is represented sequentially by the multiple potential trajectories activated one after the other. Our first aim during the project was to elaborate this *trajectory sampling* theory and make experimentally testable predictions. The second aim of the project was to use data recorded during open field navigation to test the predictions of the theory.

In the proposal we reasoned that a signature of trajectory sampling is the increase of the average error between the true trajectory of the animal and the trajectory encoded by the population activity within a single theta cycle (trajectory encoding error; *TEE*): at early theta phase, when representing past the actual and the encoded positions are similar, whereas at late theta phase, when representing future, the actual and the encoded positions diverge. During our analysis we realised that the phase dependent increase of the *TEE* is a general signature of encoding predictions in a dynamic generative model: it is equally true when uncertainty is represented via sampling, when uncertainty is not represented at all or when uncertainty is represented using other, not sampling based schemes.

Therefore we analysed the statistics of trajectories sampled from dynamic generative models and developed a set of criteria to differentiate between competing encoding schemes. In particular, we decided to include four alternative encoding schemes, all widely used in machine learning and computational neuroscience for representing uncertainty (Figure 2): 1) The product scheme, where uncertainty is represented by a probability distribution encoded as the product of neuronal tuning functions. 2) The DDC scheme (also called as the mixture scheme) where uncertainty is represented by expectations of nonlinear tuning functions under the encoded probability distribution. 3) Sampling scheme, where uncertainty is represented by drawing samples from the encoded probability distribution. 4) Maximum a posteriori (MAP) scheme, where uncertainty is not represented and only the mean of the distribution is encoded. To analyse the distinctive properties of the different coding schemes and to develop specific test to discriminate them we generated a synthetic dataset that was carefully matched to the behavioral (Figure 8) and neuronal (Figure 10) data in these navigation experiments. We compared the properties of these 4 coding schemes and found that they are all consistent with several well-known features of the experimental data, including spatial selectivity (Figure 10), phase precession and theta sequences (Figure 4D-E).

We also identified statistical signatures of increased uncertainty associated with predictions in a dynamic model both at the single neuron level and at the population level. At the single neuron level, place fields are larger when estimated from the more uncertain, late theta phase spikes, while they are more compact when calculated from early theta phase, when they represent past position (Figure 3E). At the population level, the positions decoded from late theta phase spikes show a larger forward bias and greater spread than the positions decoded from early phase spikes (Figure 3F-G). Importantly, we found that these are fundamental features of predictions in a dynamical model and were similarly true for all 4 encoding schemes irrespective of the way uncertainty is represented (Figure 4G-H).

Next we identifed the hallmarks of each of the probabilistic coding schemes and set up a decision tree to discriminate these schemes from each other (Figure 5A). We used our synthetic dataset to demonstrate that our measures are sufficiently sensitive to discriminate between the alternative encoding schemes but also robust to variations in the implementational details of the schemes. Then we analysed a dataset consisting of 8 experimental sessions recorded from 4 rates exploring a $4\,\mathrm{m}^2$ large open arena to check the signatures of these encoding schemes (Pfeiffer and Foster, 2013[13]). We found no evidence supporting either the product scheme (Figure 5) or the DDC scheme (Figure 6) in this dataset or in another dataset recorded on a linear track (Figure 11, Grosmark et al., 2016[22]). These findings indicate that the hippocampal neurons encode single locations not distributions during theta sequences.

Finally, we compared the sampling and the MAP schemes. We introduced the cycle-to-cycle variability (CCV), which measures the variability between the trajectories encoded in subsequent theta cycles and the trajectory encoding error (TEE) which is the difference between the encoded and the real trajectories (Figure 7A). Using these measures, we defined the *sampling index* as the difference between CCV and TEE. We have shown that the sampling index is positive when trajectories are randomly sampled and negative when they represent the mean trajectory. We used extensive numerical simulations to confirm that this measure is robust against variations in the timing and speed of the sampled trajectories and jitter in the detection of sequence boundaries (Figure 7F-G). We calculated the sampling index in the 4 rats and 8 recording sessions and found that it was consistently positive indicating that the variability of the encoded trajectories during hippocampal theta sequences is large (Figure 7H-L). This large variability is consistent with random sampling of prospective locations and suggests that a dynamic generative model is implemented in the hippocampus.

These results have been presented on multiple conferences (Cosyne 2019, Areadne 2018, CCCN, 2018) and most recently as a contributed talk on the annual Bernstein Conference of the German Computational Neuroscience Society (2020) and on the Generative Adversarial Collaborations online workshop (2020). We are currently working on the manuscript that we are aiming to submit for journal publication in the following months.

Test the prediction of the sampling theory on the effect of ambiguity on neuronal and behavioural variability during navigation (Aim 3., PD and FK).

The goal of this specific aim was to test the predictions of the sampling theory in data recorded in the laboratory of Neuronal Signalling using in vivo Ca^{2+} imaging. The experiments were performed by various postdocs and students in the Laboratory of Neuronal Signalling, this PD grant supported the development of the recording and analysis software. During the first two years of the project we were mostly involved in developing the control software for the virtual reality environment. During the third year we focussed on establishing the data analysis pipeline. We developed a python-based analysis framework (https://github.com/bbujfalussy/ABmice) to investigate the spatial coding in the virtual corridors.

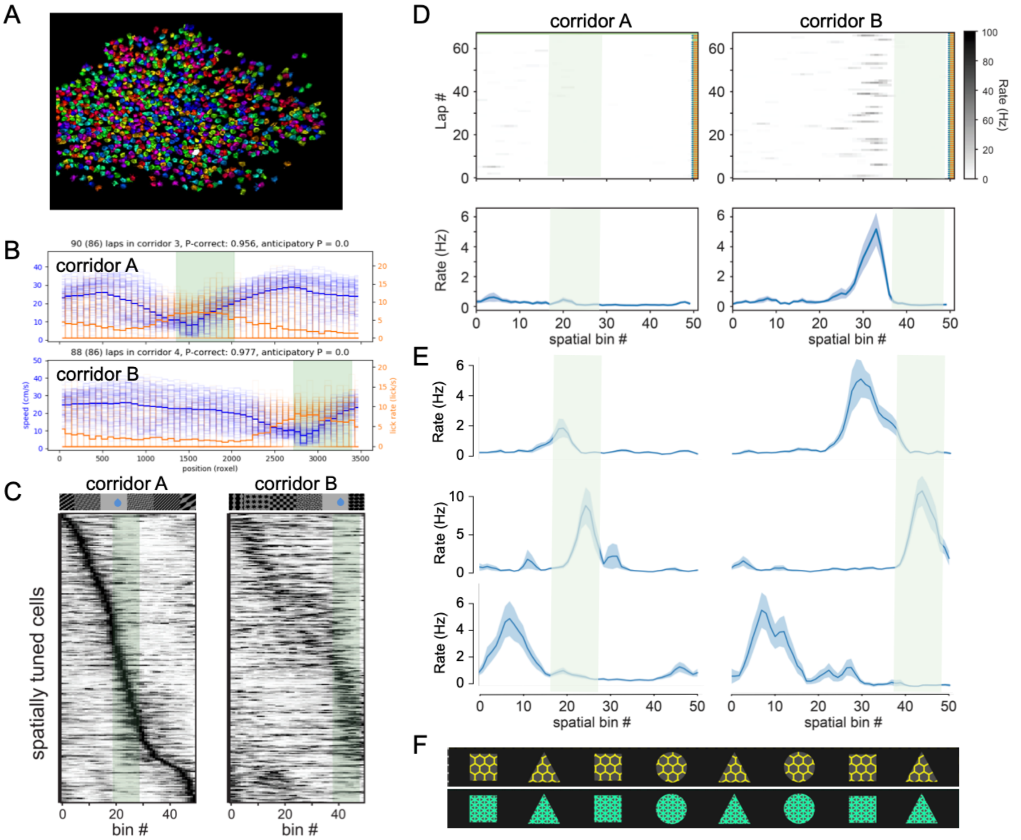

Figure 1. Imaging hippocampal neuronal activity during spatial navigation. **A** Behavior of a well trained mouse in the two corridors. The mouse typically runs at a high speed (30 cm/s), slows down and starts licking only before the reward zone (green). **B** Somata of the neurons recorded in a typical experiment. **C** Activity of spatially tuned cells in the two corridors sorted by their peak in corridor A. **D** Activity of a corridor selective place cell in each imaged laps (top) and mean and SE of the estimated firing rate of the cell as the function of position in the two corridors. **E** Other examples for rate remapping (top), reward-zone activated cell (middle) and non/selective place cell (bottom). **F** Corridors designed for the morphing analysis. The texture and the color distinguishes the two corridors while the shapes provide spatial cues.

`?⟨f:imaging⟩?`

We started the behavioral experiments with very different corridor patterns to test the spatial tuning properties of the hippocampal neurons in our experimental setup. We have successfully implemented the behavioural paradigm where mice learn to discriminate between two different virtual environments: Over the course of the training mice learn to slow down and start licking as they approached the reward location specific to the virtual corridors (Figure 1A).

We used C57BL/6J-Tg (Thy1-GCaMP6s) GP4.3Dkim/J mice f/+ expressing GCAMP-6 indicator in hippocampal CA1 neurons and we were able to monitor the activity of $\sim$ 1000 neurons during navigation in the virtual environment using an in vivo 2-photon microscope (Figure 1B). We calculated the tuning specificity, spatial information and spatial reliability of the recorded cells and performed a shuffling analysis to detect cells with significant spatial tuning. We found that in typical recording configurations 30-60% of the active cells (activity in at least 20% of the laps) had significant spatial tuning in a given corridor ($\sim$ 100 cells, Figure 1C) and corridors were sufficiently different to elicit distinct representations in the two environments (remapping, Figure 1C-E).

Based on these results we started to perform the morphing experiments where the similarity of the two corridors can be gradually increased in order to study the underlying spatial representations. We had to iterate over the patterns several times because the animals could not reliably discriminate the different corridors. The current version of the corridor pattern is shown in (Figure 1F). After we manage to successfully train the animals we will continue with the imaging and the analysis according to our original plans.

The analysis software developed by the support of the grant is freely available in my github repository `https://github.com/bbujfalussy/ABmice`.

Additional work:

Global, nonlinear dendritic computations in hippocampal polace cells (PD only). In a related project, we were interested in the synaptic and dendritic mechanisms responsible for the generation of the place code in the case of hippocampal place cells during theta sequences. We developped statistically principled methods using a hierarchical cascade of linear-nonlinear subunits (hLN) to model the dynamically evolving somatic response of neurons receiving complex, in vivo-like spatiotemporal synaptic input patterns including theta periodic inhibitory inputs and presynaptic CA3 place cell firing. We used the hLN to predict the somatic membrane potential of an in vivo-validated detailed biophysical model of a CA1 pyramidal cell. Linear input integration with a single global dendritic nonlinearity achieved above 90% prediction accuracy. Including multiple layers of local dendritic nonlinearities could improve predictions as much as 5%. Our data suggest that synaptic factors play the most important role in determining the tuning of hippocampal place cells. Our approach provides a data-driven characterization of a key component of cortical circuit computations: the input-output transformation of neurons during in vivo-like conditions.

This work appeared as an article in Neuron (Ujfalussy et al., 2018, Neuron 100, 579?592) [1]

Impact of functional synapse clusters on neuronal response selectivity (PD and FK). Following our Neuron paper, we further explored the synaptic and dendritic mechanisms responsible for the generation of the spatial tuning in hippocampal place cells during theta sequences and sharp waves.

Our goal was to understand the relative contribution of input-dependent and dendritic factors to the input-output transformation of CA1 pyramidal neurons under in vivo input conditions. We developed a computational technique to decompose the variability of the somatic voltage response of a model CA1 pyramidal neuron to in vivo-like inputs into three components: 1) dendritic factors i.e., differential spatial distribution of inputs active along the track including small-scale functional clustering or large-scale spatial heterogeneities (e.g., proximal vs. distal dendritic location). 2) variability in the inputs active along the track, including variations in the precise number of presynaptic neurons with place fields at a given location, their maximal firing rates or their synaptic strengths; 3) Trial-to-trial variability associated with stochastic biophysical processes (e.g., spiking and synaptic vesicle release). We found that large-scale dendritic spatial inhomogeneities in synaptic tuning properties did influence sVm, but small synaptic clusters appearing randomly with unstructured connectivity did not. With structured connectivity, 10-20 synapses per cluster was optimal for clustering-based tuning, but larger selectivity was achieved by 2-fold potentiation of the same synapses. We further showed that without nonlinear amplification of the effect of random clusters, action potential-based, global plasticity rules can not generate functional clustering.

Moreover, we analysed dendritic integration of pyramidal neurons during sharp wave activity and found that the impact of synaptic clustering on the neural response is small during SPWs and can even be negative when large clusters are combined with LTP and synaptic saturation becomes the dominant effect. Furthermore, using our computational model we showed that strong Na+ spikes can overcome synaptic saturation during SPWs and make the CA1 model neuron more responsive to large synaptic clusters.

Our findings indicate that 1) the selective responses of cortical neurons are primarily the consequence of the tuning of their synaptic inputs, 2) functional synaptic clustering matched to local dendritic properties can have additional role in refining those responses 3) plasticity of functional synapse clusters such as those observed *in vivo* requires local rather than global mechanisms, and 4) in turn, local plasticity by small synaptic clusters may lead to powerful tuning of somatic responses.

This work has been published in Nature Communications 11:1413 (2020) [2].

Changes in dendritic excitability by the LTP induction protocol cannot explain synaptic crosstalk (PD only). Finally, we also used the model to analyse the synaptic versus dendritic mechanisms of heterosynaptic crosstalk during synaptic plasticity experiments in CA1 neurons. In these experiments we found that input patterns evoking dendritic Na^+ or NMDA spikes can induce robust long term synaptic potentiation (LTP) not only on stimulated synapses but can also strengthen nearby, nonsynchronous synapses by local heterosynaptic plasticity crosstalk. Importantly, the LTP was not accompanied by changes in the somatic strength of dendritic Na spikes, a possible indicator of dendritic excitability.

We used or detailed biophysical CA1 pyramidal neuron model to explore whether local changes in the dendritic excitability can explain the changes in the response to non co-stimulated synapses. In our model we considered changes in either the passive dendritic properties (increasing the local membrane resistivity and decreasing axial resistivity) or voltage dependent ion channels (eliminating potassium channels). Contrary to the experimental observations, both of these manipulations increased the dV/dt of the somatic spikelets. On the other hand, increasing the AMPA conductance of the synapses by 40% (mimicking LTP) increased the amplitude of the EPSPs without changing the spikelets. These simulations made it unlikely that changes in dendritic excitability by so far described mechanisms could alone explain the increase in somatically measured amplitude of the test spine EPSPs, and suggest that crosstalk was most likely mediated by synaptic mechanisms.

This work has been published in The Journal of Neuroscience, 40(13):2593?2605, (2020) [3].

Appendix - Sampling-like representation of prospective locations during hippocampal theta sequences.

Model-based planning and predictions are necessary for flexible behavior in various cognitive tasks. Representing uncertainty of predictions in volatile environment is beneficial[4]. Representation of and computations with uncertainty has recently gained support in a wide range of cognitive phenomena[5]. Hippocampus is essential for temporal association memory[6–8] and model based planning[9]. Whether uncertainty associated with predictions is encoded in the hippocampus and how it is represented are not known.

During spatial navigation, the activity of neurons in the rodent hippocampus represents both the immediate location of the animal and its near future position[10,11], a property that has long been implicated in predictive models and planning[12]. At the population level this prospective coding is manifested by fast, sequential activity patterns emerging either during awake immobility and sharp waves[13] or during active exploration and theta oscillation[14]. Here we focus on *theta sequences* when place cells encoding past, present, and future locations become sequentially active within a single cycle of the 6-10 Hz theta oscillation outlining a possible trajectory of the animal[15] (Figure 2A). We interpret this activity pattern as the consequence of the animal repeatedly performing probabilistic inference about possible trajectories in a dynamical generative model. Critically, during a single theta sequence the uncertainty of the position representation is expected to increase systematically from the beginning, when neurons represent past location, towards the end of theta, when they encode future position (Figure 2A), thus providing a chance to identify how uncertainty is encoded in the population activity.

Several encoding schemes have been proposed for representing uncertainty in the activity of neural populations: In one class of schemes, an entire probability distribution can be represented by a snapshot of the neural activity either as a product[16] or as a mixture[17,18] of the underlying basis functions (loosely tuning curves) where the uncertainty is loosely proportional to the inverse of the gain or the diversity of the active neurons, respectively (Figure 2B). More formally, these proposals correspond to variational approximations of the underlying distributions using canonical parameters or mean parameters ([19]). Alternatively, the uncertainty can be encoded by the variability of the neural activity through time, when distributions are represented by random samples[20], a technique called Monte Carlo in statistics[21]. However, measures that could dissociate these alternative schemes based on experimental data are remarkably lacking.

In particular, dissociating these schemes requires; (i) to know what are the variables represented in the population activity; (ii) systematic changes in the uncertainty and (iii) recordings from large number of neurons across many repetitions to gain statistical power. Here we capitalize on the unique properties of hippocampal code during theta oscillation when the (i) population activity represents spatial location; (ii) uncertainty increases within a theta cycle and (iii) recordings of hundreds of neurons for 1000s of theta cycles are available[13,22] and apply a novel analysis framework to identify the neural representations of uncertainty during model based predictions.

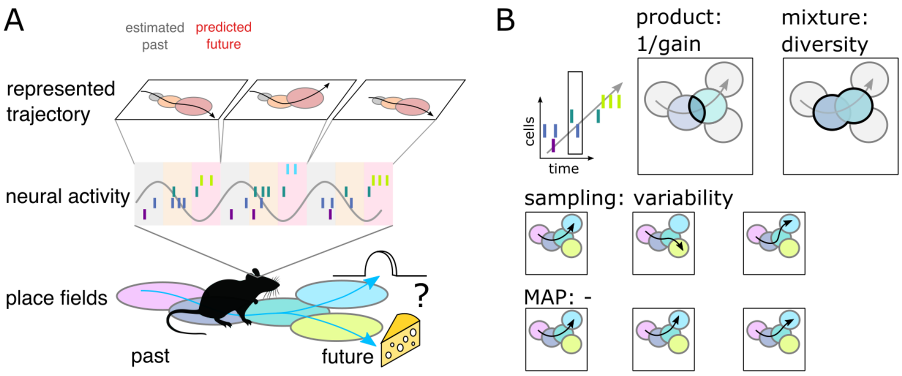

Figure 2. Theta sequences, uncertainty and variability. A Schematic showing that neural activity represent a possible trajectory in each theta cycle during navigation. **B** Possible schemes for representing uncertainty in neural population activity.

?⟨f:tseq⟩?

Results

In the present paper we first show that the trial-to-trial variability of theta sequences increases within a theta cycle both at the single cell and at the population level. Although these observations first seem to provide evidence for a sampling based representation, we demonstrate that they are direct consequences of predictions in a probabilistic system (the motion of the animal), and hence they are equally consistent with all four alternative coding schemes. In order to discriminate between between the different schemes we derive more specific measures and confirm, using simulated data matched statistically to real data, that our measures can reliably discriminate between the alternative frameworks. We then apply our analysis to experimental data and show that there are no signatures of representation of uncertainty in snapshots of the hippocampal population activity. Finally, we found that the variability between subsequent theta cycles is larger than that expected from encoding only the most likely trajectory, and is consistent with random sampling from hypothetical future trajectories. These results suggest a new perspective on the functional role of variability during hippocampal theta oscillation and provide a novel framework for distinguishing between conflicting proposals about how uncertainty is represented in the brain.

Neural variability increases within theta cycle. To investigate neural correlates of uncertainty during theta sequences, we analysed hippocampal spike trains recorded during theta oscillation in rats performing spatial navigation in a 2×2 m large open arena[13]. Individual neurons typically had location-related activity (a few example tuning curves are shown in Figure 3B, see also ref. 13), but their spike trains were highly variable (Figure 3A, ref. 23,24).

Next we studied the trajectories encoded during individual theta cycles by decoding the spike trains in overlapping 20 ms time bins (Figure 3C). Despite the large cell counts (68–242 putative excitatory cells in n=8 sessions from 4 rats), the decoding had a limited accuracy (Fisher lower bound on the decoding error in 20 ms bins: 16–30 cm vs. typical trajectory length ∼ 20 cm), yet in high spike count cycles we could approximately reconstruct the encoded trajectory (Figure 3C). We then compared the reconstructed trajectories to the actual trajectory of the animal. We observed that individual decoded trajectories displayed substantial deviation from the real trajectory of the animal: trajectories typically started near the actual location of the animal and then proceeded forward (Figure 3C)[15] often departing either to the left or to the right from the actual path taken by the animal[25].

To systematically analyse how this deviation depends on the phase of the trajectory we sorted spikes into three categories according to their phase in the theta cycle (early, mid and late) and compared the decoded position with the actual position shifted in time (Figure 3D). For early spikes, the minimum of the average decoding error was shifted by ∼-100 ms while for late spikes +500 ms shift minimised the decoding error indicating that on average, theta sequences started in the recent past and extended into the near future. Further, the minima of the decoding curves showed marked differences, with the position decoded from later spikes showing a greater average deviation from the true trajectory of the animal (Figure 3D).

At the single neuron level the increased variability at late theta phases was observed by the expansion of the place fields of the neurons were more compact when estimated from the more precise, early phase spikes than when calculated from late phase activity (Figure 3E). At the population level, the larger deviation between the true trajectory and the decoded future positions in the case of late spikes can be attributed to the increased variability in the encoded possible future locations: when we aligned the decoded locations relative to the current location and motion direction of the animal, we observed that the spread of the decoded locations increased in parallel with the forward progress of their mean within theta cycles (Figure 3F-G). This analysis confirmed that observed theta sequences typically started around the actual location of the rat and ended on average ∼20 cm ahead of the animal, similar to the value found in 1-dimensional mazes[14]. More importantly, this analysis demonstrated that cycle-to-cycle variabilityis larger at the end of the theta cycle when encoding future locations.

The parallel increase of the uncertainty and the cycle-to-cycle variabilityat the end of theta cycles could be taken as evidence favoring a sampling based representation. However, qualitatively, the increased variability is a direct consequence of variable 2-dimensional trajectories of the animal. To illustrate this, we plotted the location and direction aligned real trajectories and found that their spread similarly increases when we move away from the present either towards the past or the future (Figure 3H-I). Furthermore, in the next sections we use a synthetic dataset to demonstrate that the increase in the

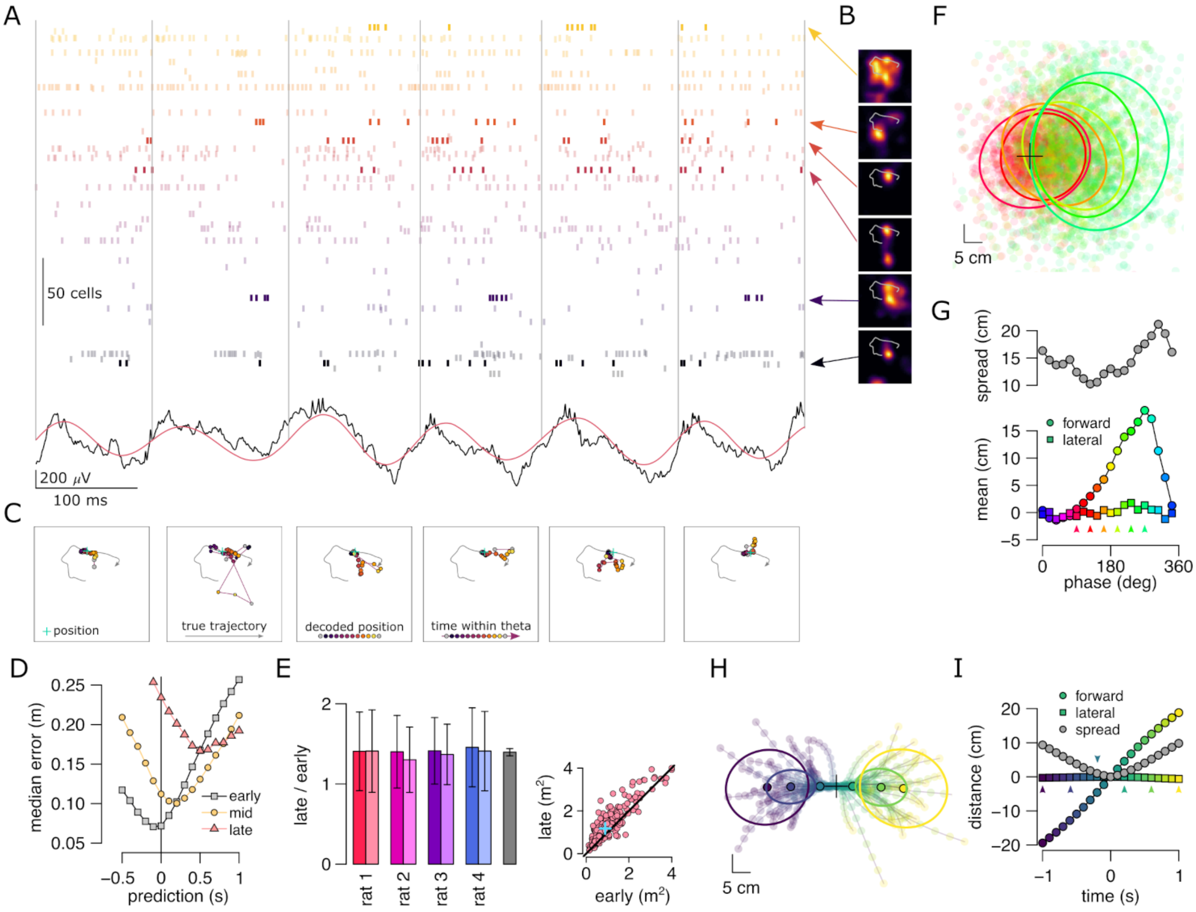

Figure 3. Theta sequences, uncertainty and variability. A Example spiking activity of 250 cells (top) and row (black) and theta filtered (red) local field potential (bottom) for 6 consecutive theta cycles (vertical lines). **B** Place field of 6 selected cells in a 2x2 m large open arena. Grey line indicates a run episode analysed in B-D. **C** Position decoded in overlapping 20 ms time bins during the 6 theta cycles shown in A. Time within theta cycle is color coded, grey circles indicate bins on the border of two cycles. Grey line shows the true trajectory during the run episode, green cross indicates the location of the animal in each cycle. **D** Decoding error for early, mid and late phase spikes (highest 5% spike count cycles) as a function of the temporal shift of the animal's position. For the analysis in panels D-F each theta cycle was divided into 3 parts with equal spike counts. **E** Relative place field size in late versus early theta spikes for the 8 sessions (error bars: SD over cells). Grey bar: average and SD over the 8 sessions. Inset: Place field size (ratemap area above 10% of the max firing rate) estimated from late vs. early theta spikes. Only putative excitatory cells are included. We shifted the reference positions with Δt that minimised decoding error to estimate ratemaps (see panel D). Median is indicated with blue cross. **F** Location and direction aligned decoded positions, and 0.5 confidence interval (CI) ellipses for 6 different theta phases. Black cross: current position. **G** Bias (bottom, mean of the decoded positions) and spread (top, $(\det \Sigma)^{1/4}$) of decoded positions as a function of theta phase ($n_{sp} \geq 13$) for R1D2. **I** Location and direction aligned movement trajectories and 0.5 confidence interval ellipses for $-1 \geq \Delta t \geq 1$. **J** Bias and spread of movement trajectories as a function of time in an example session (R1D2).

?⟨f:var⟩?

variability during theta cycles is consistent with all the 4 representational schemes and use these datasets to develop specific tests to discriminate the different coding schemes.

Synthetic data: testbed for discriminating the encoding schemes.

To analyse the distinctive properties of the different coding schemes and to develop specific test to discriminate them we generated a synthetic dataset that was carefully matched to the behavioral and neuronal data in these navigation experiments.

Our synthetic dataset had three consecutive levels. First, we simulated the motion of a rat in a two dimensional open arena by allowing smooth changes in the speed and motion direction and matching the kinematics of the real and the simulated animal (Figure 8). Second, similar to the real situation, the simulated the animal did not have access to its true position x_t in time t, but had to infer it from the sensory inputs it had observed in the past $\mathbf{y}_{1:t}$ (Methods). To perform this inference and predict its future position, the animal used a model of its own movements in the environment. In this dynamic generative model the result of the inference is a distribution over possible trajectories starting in the past and ending in the future $P(\vec{\mathbf{x}}_t|\mathbf{y}_{1:t})$, where $\vec{\mathbf{x}}_t = x_{(t-t_p):(t+t_f)}$.

Third, in our simulations the hippocampal population activity encoded the resulting trajectory at an accelerated speed in each theta cycle such that the trajectory started in the past at the beginning of the theta cycle and arrived to the predicted future states (locations) by the end of the theta cycle (Figure 4A-C). We simulated the activity of 200 hippocampal pyramidal cells (*place cells*) whose firing rates depended on the encoded spatial location using either of the 4 different coding schemes (Methods): In product and DDC codes a snapshot of the population activity at any given time encoded the posterior distribution about the position at time t' based on the sensory information up to time t: $P(\mathbf{x}_{t'}|\mathbf{y}_{1:t})$. In the MAP code, the population encoded the most likely trajectory and in sampling a single trajectory was sampled randomly in each theta cycle from the posterior distribution of possible trajectories (Figure 9). The posterior distribution was updated in each of the 100 ms long theta cycles. Importantly, all of the 4 encoding models yielded single neuron and population activity dynamics consistent with the majority of known feature of hippocampal population activity including spatial tuning, phase precession (Figure 4D) and theta sequences (Figure 4E, see also Figure 10).

Next, we repeated the our previous analysis regarding the changes in neural variability within a theta cycle on the synthetic dataset. First we decoded the population activity in early, mid and late theta phases and compared the estimated position with the actual trajectory of the simulated animal. Although simulated trajectories started further in the past and terminated earlier than in the real data (Figure 3D), the deviation between the decoded position and the actual trajectory increased throughout the theta cycle irrespective of the coding scheme (Figure 4F; see also Figure 10). Moreover, place fields were significantly larger when estimated from late than early phase spikes (Figure 4F, $\mathrm{P} < 10^{-16}$ for each schemes). Finally, when aligned to the current position and motion direction of the simulated animal, the spread of the decoded locations increased with the advancement of their mean within theta cycles in all 4 coding schemes (Figure 4G).

Our analyses thus confirmed that the increased variability of the hippocampal neural activity at late versus early theta phase is consistent with encoding trajectories from a dynamic model of the environment irrespective of the representation of the uncertainty. This synthetic datasets provides a functional interpretation of the hippocampal population activity during theta oscillations and offer a unique framework for testing representation of uncertainty.

Testing the product scheme: gain

In the following sections we will identify the hallmarks of each of the probabilistic coding schemes set up a decision tree to discriminate these schemes from each other (Figure 5A).

To discriminate the product scheme from other representations we capitalise on the specific relationship between response intensity of neurons and uncertainty of the represented variable. In a product representation, a probability distribution over a feature, such as the position is encoded by the wighted product of the neuronal tuning functions where the weights are the number of spikes fired by the neurons. When tuning functions are localised (e.g., Gaussian[16]), the spread of the encoded distribution tends to

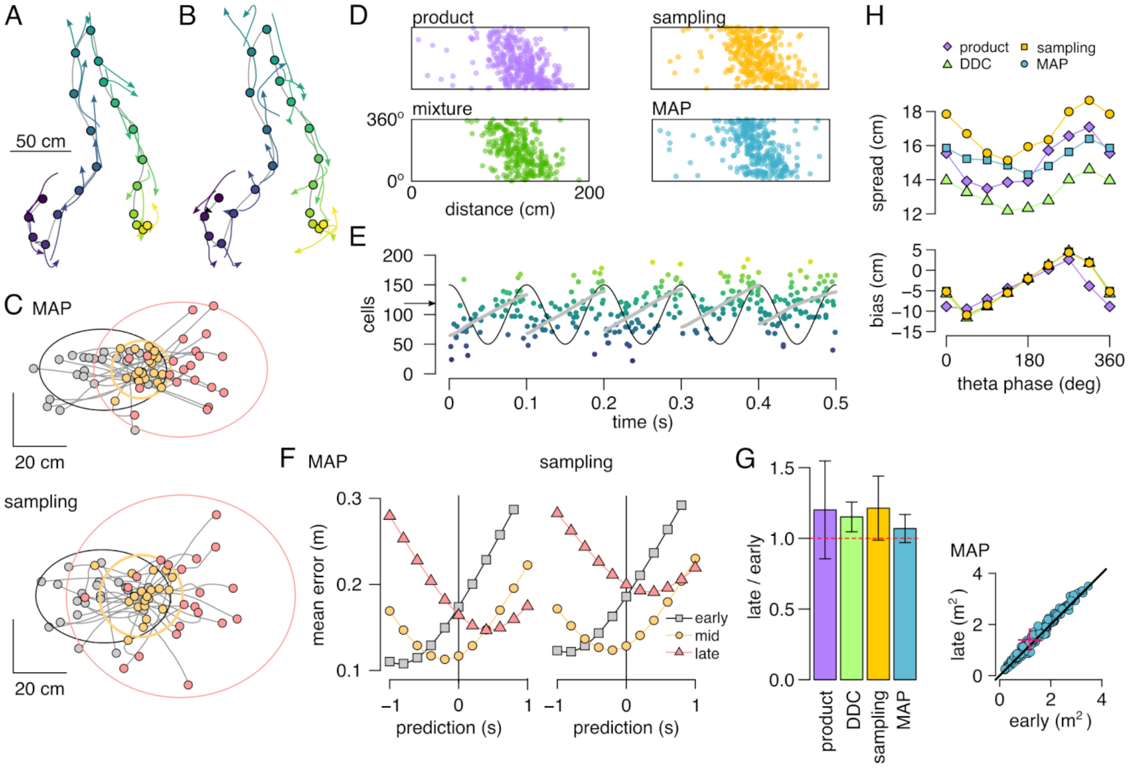

Figure 4. Theta sequences in simulated data. A Trajectory of the simulated rat (grey, 20 s) together with its own inferred and predicted most likely (MAP) trajectory segments (colored arrows) in 21 locations separated by 1 s (filled circles). These trajectories (and many others corresponding to locations between the highlighted) were represented in a theta cycles using 1 of the 4 possible schemes. **B** Same as panel A, but trajectories sampled from the posterior distribution. **C** Example represented trajectories aligned to the position and direction of the simulated animal in the MAP (top) and in the sampling (bottom) scheme. Ellipses indicate 50% CI of all theta cycles. Color code indicates start (grey), mid (yellow) and end (pink) of the trajectories. **D** Theta phase as a function of position in the 4 schemes. **E** Simulated hippocampal population activity in MAP code in a linear track (200 x 10 cm) during an idealised 10 Hz theta oscillation. Grey lines show the represented trajectories in each theta cycle. **F** Decoding error from early, mid and late phase spikes (highest 5% spike counts) as a function of the temporal shift of the simulated animal's position in a MAP (left) and in sampling code (right). **G** Relative place field size in late versus early theta spikes for the 4 different models (error bars: SD over 200 cells). Inset: place field size estimated from late vs. early theta spikes in the MAP code. Median is indicated with red cross. **H** Decoding bias (bottom) and spread (top) as a function of theta phase for the 4 different encoding models ($n_{sp} > 0\%$ - product model is biased otherwise).

⟨f:synth⟩

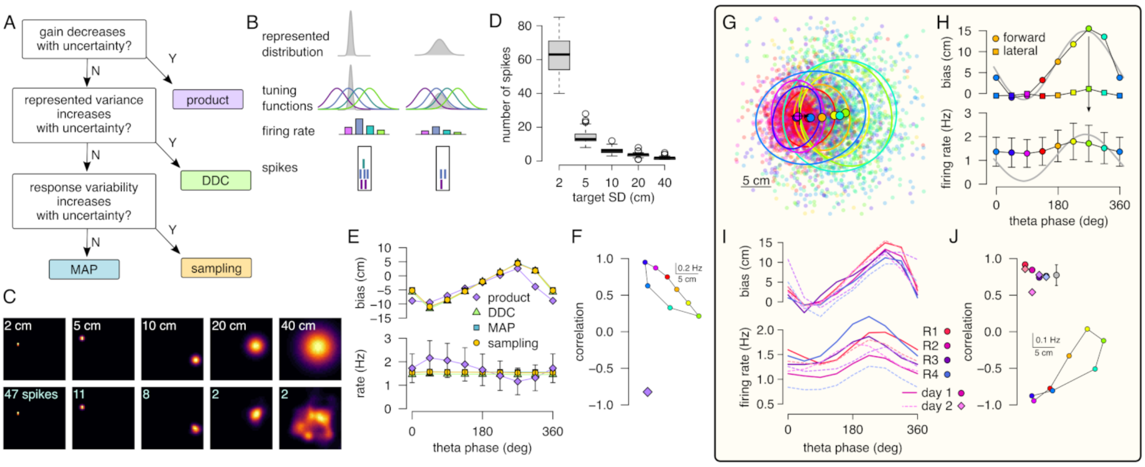

Figure 5. Product scheme: population gain decreases with uncertainty. **A** The decision-tree for identifying the representational scheme. **B** Schematic of encoding a narrow (top) and a wide (bottom) distribution using the product scheme with 4 neurons. In the simplest case, the SD is represented by the gain of the population. **C** Representing distributions of increasing uncertainty in the product form using realistic, localised basis functions (similar to Figure 3B). Top: target isotropic Gaussian distribution (the number shows the SD). Bottom: Approximate distribution with the total number of spikes used in the approximation. **D** The number of spikes required for approximating a distribution decreases as a function of the target SD. **E** Firing rate (bottom) and decoding forward bias (top) as a function of theta phase for the 4 schemes in our simulated dataset. Only the product scheme predicts a systematic change in the firing rate. **F** Correlation between firing rate and forward bias for the product scheme. Inset: Firing rate as a function of forward bias in the product scheme. Color code is the same as in H. **G** Decoded position relative to the location and motion direction of the rat (black arrow) in different theta phases using high spike count cycles in an example animal ($n_{sp} > 95\%$). Filled circles indicate mean, ellipses indicate 50% CI of the data. (Similar to Figure 4G, but from 120° and not 20 ms binning.) **H** Bias of the relative decoded position (top) and the firing rate (bottom) as a function of theta phase for session R1D2. Grey line in top and bottom show cosine fit to the forward decoding bias ($n_{sp} > 95\%$). **I** Decoding bias (top) and firing rate (bottom) for all animals and sessions (line type). **J** Correlation between firing rate and forward bias for all recorded sessions. Grey symbol: mean and SD across the 8 sessions. Inset: Firing rate as a function of forward bias R1D2. Bias in H and I used theta cycles with $n_{sp} > 95\%$. Firing rate plots used show average over all cycles.

`?⟨f:prod⟩?`

decrease with the spike counts, thus the variance is related to the gain population response (Figure 5B). To demonstrate this, we approximated isotropic 2-dimensional Gaussian distributions with real tuning functions derived from the data (Figure 3B) in the form of

$$\mathcal{N}(x|\boldsymbol{\mu}, \sigma^2\mathbf{I}) \propto \exp\Big(\sum_i s_i \phi_i(x)\Big) \qquad (1)\ \{?\}$$

where $e^{\phi_i(x)}$ is the tuning function of neuron i and s_i is the number of spikes it fired (see Methods for details). We found that the approximations were reasonably accurate, especially for distributions with lower variance (Figure 5C). Importantly, the average number of spikes used for encoding decreased systematically with the SD of the encoded distribution throughout the range relevant for hippocampal theta sequences (Figure 5D; see the range of spread values in Figure 3G,I). Therefore we used the systematic variation of the population gain with uncertainty as a hallmark of the product scheme.

Next we used this criteria to discriminate the different coding schemes in the synthetic dataset. We decoded the spikes using 8 overlapping 120° windows and identified the theta phase with the maximal forward bias as the end of the theta cycle encoding future predictions (Fig. Figure 5E top). Then we calculated the average number of spikes in a given 120° window as a function of the theta phase. We found that only the product scheme predicted a systematic modulation of the firing rate within the theta cycle (Figure 5E, bottom). In this case the peak of the firing rate coincided with the theta phase encoding backward positions, when the uncertainty is minimal. The three other coding schemes do not constrain the firing rate of the population to represent probabilistic quantities, and thus the firing rate was independent of the theta phase or the encoded uncertainty.

After testing on synthetic data, we returned to the experimental dataset[13] and applied the same analysis to neuronal activity recorded from freely navigating rats. We first decoded the spikes during theta oscillation falling in 8 overlapping 120° window and aligned the decoded locations relative to the animals position and motion direction. With this coarser time resolution, we confirmed that the encoded location varied systematically within the theta cycle from the beginning towards the end of the theta cycle both when considering all theta cycles (**??**) or when focusing on theta cycles with the highest 5% spike count (Figure 5F-G cf. Figure 3F-G). The spread of the encoded distribution typically had a maximum near the end of the theta sequences encoding the forward predictions (Figure 5G). We also confirmed previous observations[26] that the population firing rate also had a maximum at the end of the theta sequences when representing the most uncertain future positions (Figure 5G, bottom). This result was consistent across all the recording sessions (Figure 5H) and was also confirmed in an independent dataset where rats were running on a linear track (Figure 11A-D, Grosmark et al., 2016[22]).

This observation is in sharp contract with the prediction of the product encoding scheme where the maximum of the firing rate should be near the beginning of the theta sequences (Figure 5E). The other encoding schemes are neutral about the theta modulation of the firing rate, and therefore they are all consistent with the observed effect.

Testing the mixture scheme: diversity

Next, we set out to discriminate the mixture encoding scheme from MAP and sampling. In codes using the mixtures schemes neuronal activity corresponds to the mean parameters of the encoded distribution[19]. Specifically, in this scheme the firing rate of neurons represents the expectation of their tuning functions under the encoded distribution[17,18]. Naturally, in this scheme the diversity of the co-active neurons increases with the variance of the encoded distribution (Figure 6A). When the encoded variance is small, only neurons with overlapping tuning functions will fire synchronously, while a large diversity of neurons become co-active when encoding a distribution of high variance (Figure 6A).

To establish the feasibility of the analysis, we first explored this coding scheme in a simplified synthetic dataset without spatial exploration and trajectories. We encoded isotropic Gaussian distributions of increasing SD, $\sigma = \{2, 5, 10, 15, 20\}$ [cm] in the spiking activity of a neural population using experimentally derived tuning curves (similar to Figure 3B) and, instead of correlating σ with a chosen measure of the neuronal diversity, we built a maximum likelihood decoder of the parameters of the encoded distribution. We found that using a sufficiently long time window ($\Delta t = 100\,\mathrm{ms}$) we could reliably estimate the encoded mean and SD (Figure 6B, see also Figure 12). However, the decoder of the SD became substantially biased when we decreased the time window to the relevant range ($\Delta t \sim 20\,\mathrm{ms}$). We found that this bias can be significantly reduced if we correct for the differences between the wider empirical tuning curves and the narrower tuning functions by sharpening the tuning curves (Figure 12). The remaining bias made trial by trial decoding of the represented distributions unfeasible in short time scales. However, when averaged across many 1000 cycles, the decoder was still sufficiently sensitive to the encoded SD to identify systematic changes (Figure 6B, right). Therefore, we used the SD decoded from the neural activity as a hallmark of DDC encoding (Figure 6A).

Next, we turned to our synthetic dataset encoding trajectories during simulated movement. In each of the three remaining coding schemes (mixture, sampling and MAP) we decoded the population activity in three separate windows of the theta cycle (early, mid and late) assuming that it represented the distribution of the possible trajectories (Figure 6C inset) in the mixture form. Similar to our experience in the simplified dataset (Figure 6B) we found a significant increase in the average decoded SD in the mixture scheme from early to late theta phases, whereas the decoded SD was independent of the theta phase for MAP and sampling codes (Figure 6C-D). We found that although compensating for the over-dispersion of the estimated tuning curves reduced the bias of the SD decoding, this was not necessary to discriminate between the coding schemes (Figure 12). Thus, the decoded SD is a reliable measure to discriminate the mixture scheme from sampling or MAP encoding.

After testing on synthetic data, we repeated the same analysis on the 2D real dataset. We divided each theta cycle into three sections with equal spike counts (early, mid and late) and decoded the population activity assuming that it represented the a distribution of trajectories according to the mixture scheme in each theta cycles. We decoded the population activity to obtain the estimate of the mean and the SD of the encoded trajectories in early, mid and lates phases of the theta. We found that the decoded SDs had a nearly identical distribution in the three theta phases for all recording sessions (Figure 6E). The mean of the decoded SD did not change significantly or consistently across the recording session neither when we analysed all theta cycles (Figure 6F) nor when we focused on the half of the cycles with higher than median spike count (Figure 6G) or when we did not compensate for the over-dispersion of the estimated

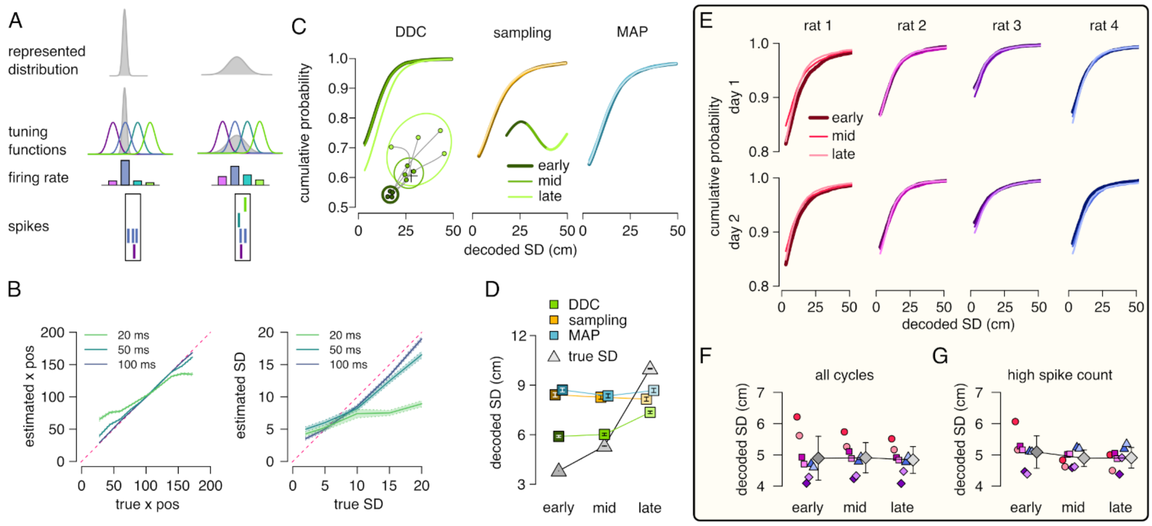

Figure 6. Mixture scheme: diversity increases with uncertainty. **A** Schematic of encoding a narrow (top) and a wide (bottom) distribution with spike-based DDC using 4 neurons. Intuitively, the SD is represented by the diversity of the co-active neurons. **B** Decoding in a spike-based DDC scheme. Estimation of the mean (left) and the SD (right) of the encoded distribution becomes biased when the observation window is too short (B). **C** Cumulative probability of the decoded SD of the represented distribution at different theta phases for the DDC, sampling and MAP schemes in the simulated dataset. **D** Mean and SE of decoded SD as a function of theta phase for the different schemes in the simulated dataset. Grey triangles show the tru SD. Only the DDC code predicts a slight, but significant increase in the decoded SD at late theta phases. **E** Cumulative probability of the decoded SD from spikes in early, mid and late phase of the theta cycles (all cycles) for the analysed sessions. **F-G** Mean and SD across animals of the decoded SD from early, mid and late theta spikes using all theta cycles (F) or with $n_{sp} > 50\%$ (G).

`?⟨f:mix⟩?`

tuning curves (Figure 12). We obtained similar results in a different dataset with rats running on a linear track (Figure 11E-I, Grosmark et al., 2016[22]). We conclude that we could not find signatures of a mixture scheme in the hippocampal population activity during theta oscillation. These findings indicate that the hippocampal neurons encode single locations not distributions during theta sequences.

Hallmark of sampling: variability

Sampling-based codes are characterised by large and structured trial-to-trial neuronal variability[27]. However, we saw that the divergence of possible future trajectories led to a systematic increase in the variability within theta cycles that was independent of the underlying encoding scheme (Figure 3-4). In sampling the cycle-to-cycle variabilityis further increased by the randomness of the selection of the represented trajectory in each theta cycle. In order to discriminate sampling from the MAP scheme we quantified the cycle-to-cycle variability(CCV) by comparing it to the difference between the encoded and the real trajectories (trajectory encoding error, TEE; Figure 7A). Specifically, we defined the *sampling index* as the difference between CCV and TEE. We have shown (Methods), that the sampling index is positive when trajectories are randomly sampled and negative when they represent the mean trajectory.

We then we calculated the sampling index in our synthetic dataset for population activity representing either randomly sampled trajectories or mean trajectories (MAP). In particular, we decoded the population activity in three separate windows of the theta cycle (early, mid and late) using a standard static Bayesian decoder[28] and computed the difference between the decoded locations across neghbouring theta cycles (CCV) and the difference between the decoded position and the true location of the animal (TEE). As expected both TEE and CCV increased from early to late theta phase both for MAP (Figure 7B-C) and sampling codes (Figure 7D-E). We found that it is the magnitude of the increase that is the most informative of the identity of the code: the cycle-to-cycle variabilityincreases more within theta than trajectory encoding error in the case of sampling (Figure 7E) whereas increase in the trajectory encoding error is larger in the MAP encoding (Figure 7C). The sampling index measures this difference and was negative for MAP (Figure 7F) and positive for sampling codes (Figure 7G). To test the robustness of the sampling index against various confounding factors influencing the cycle-to-cycle variabilitywe varied

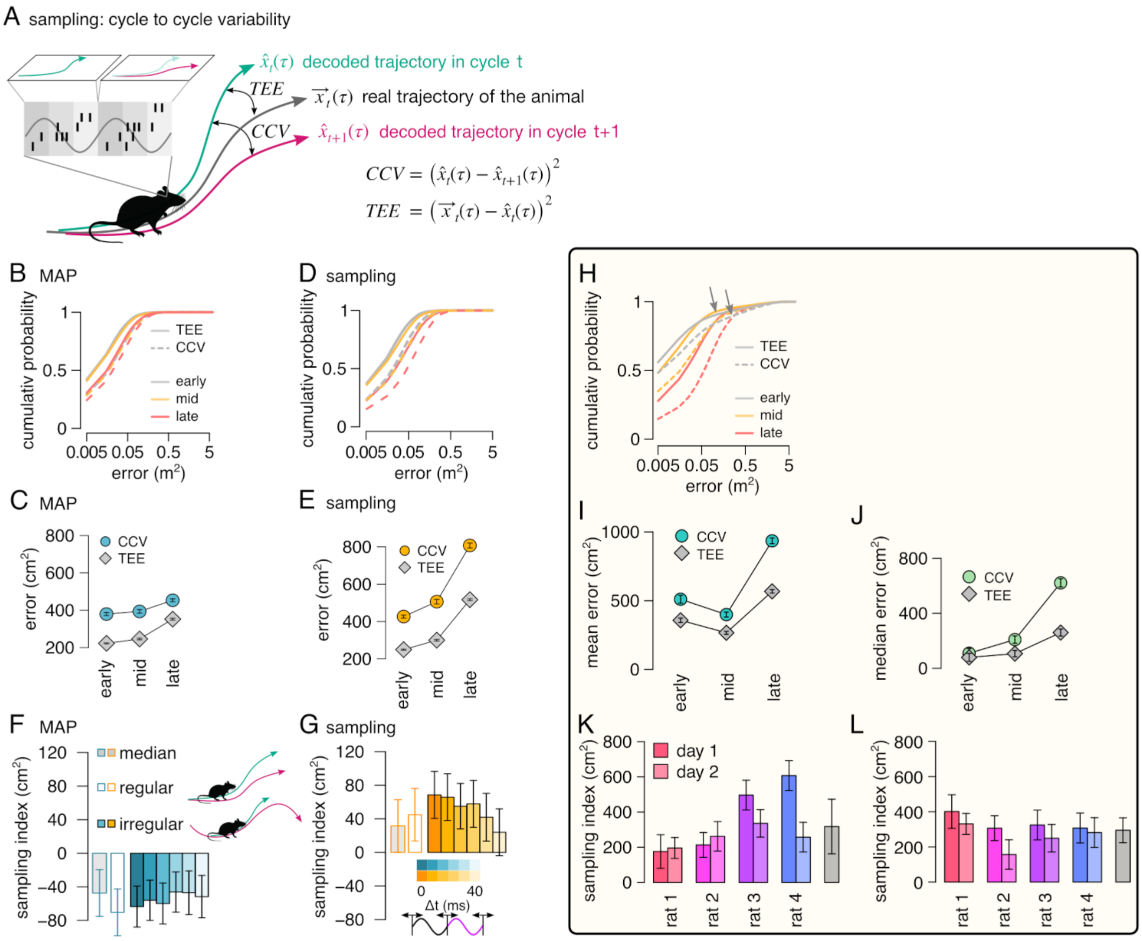

Figure 7. Sampling scheme: cycle-to-cycle variabilityincreases with uncertainty. A To discriminate sampling from MAP encoding we need to calculate the sampling index which measures the magnitude of cycle-to-cycle variability (CCV) relative to the trajectory encoding error (TEE). **B-C** Cumulative distribution of CCV (dashed) and TEE (solid) for early, mid and late theta phase (colors) in the MAP (B) and sampling (C) schemes using simulated data. Note the logarithmic x axis. **D-E** Mean CCV and TEE calculated from early, mid and late phase spikes in the MAP (D) and sampling (E) schemes. **F-G** The sampling index for the MAP (F) and sampling (G) schemes calculated from the median (grey) or mean (colors) across all theta cycles with simulating regular or irregular theta-trajectories (left inset) and applying various amount of jitter for segmenting the theta cycles (right inset). **H** Cumulative distribution of CCV (dashed) and TEE (solid) for early, mid and late theta phase (colors) for session R1D2. Note the logarithmic x axis. **I-J** Mean (I) and median (J) of CCV and TEE calculated from early, mid and late phase spikes for session R1D2. **K-L** Sampling index calculated for all analysed sessions (color) and across session mean and SD (grey) using the mean of all theta cycles (K) or the median of all cycles (L).

⟨f:sam⟩

the speed and the length of the encoded trajectories (irregular vs. regular trajectories, Figure 7F inset, Methods) and the jitter of the boundaries of the theta cycles (jitter 0-40 ms, Figure 7G inset, Methods). We found that the sampling index was robust against these variations, reliably discriminating sampling based codes from MAP codes across a wide range of parameters.

Finally we repeated the same analysis on the dataset recorded from rats the 2D spatial arena. We calculated the cycle-to-cycle variability and the trajectory encoding error and found that although typically both the error and the variability increased during theta, their distributions had high positive skewness (Figure 7H). The outliers preferentially occurred at early theta phase (grey arrows in Figure 7H), when they increased the mean of the distribution above its mid theta value (Figure 7I). The median was more robust to the effect of the outliers (Figure 7J). We observed similar effect on synthetic data when the jitter on the sequence boundaries was large (not shown). Therefore we calculated the sampling index using both the mean (all cycles: Figure 7K or median (Figure 7L) across theta cycles or including only high spike count cycles ($\mathrm{n_{sp}} > 50\%$). We found that the sampling index was consistently positive for all recording sessions. This data indicates that the large cycle-to-cycle variabilityof the encoded trajectories during hippocampal theta sequences is consistent with random sampling from the posterior distribution of possible trajectories in a dynamic generative model.

Methods

Theory

To study the neural signatures of the probabilistic coding schemes during hippocampal theta sequences we developed a coherent theoretical framework which assumes that the hippocampus implements a dynamic generative model of the environment. The animal uses this model to estimate its current spatial location and predict possible consequences of its future actions. Since several possible positions might be consistent with the recent sensory inputs and multiple options are available to choose from, these alternative possibilities could be represented in the neuronal activity during theta oscillation. Within this framework, we interpreted theta sequences as representing inferences and predictions in the dynamic generative model.

We describe this generative process at three hierarchical levels. First, we generated random but smooth goal trajectories in the 2-dimensional box. These trajectories represented the intended position for the animal at each time step. Second, similar to the real situation, the animal had to use a model of its environment and his own motion to infer its trajectory from the noisy or ambiguous sensory cues and calculate the next motor command to maintain the intended trajectory. Third, the inferred trajectory is represented by the activity of the simulated hippocampal neurons using either of the 4 different representational schemes. Importantly, we used this algorithm only to sample trajectories from the right distribution in the simulated data and we do not assume that calculating hypothetical trajectories would involve similar algorithmic steps when implemented in the hippocampal network. The flexibility of this hierarchical framework enabled us to fit qualitatively the experimental data both at the behavioral (Figure 8) and the neural (Figure 10) level. In the following paragraphs we will describe these levels in detail.

Smooth goal trajectory. To generate flexible movement trajectories allowing efficient inference we adopted the a linear Gaussian state space model[29] (Figure 9A). The flexibility of the motion was achieved by generating complex goal trajectories that the simulated animal intended to follow. To maintain tractability, the simulated animal had access to the goal trajectories and both sensory and motor noise had Gaussian distribution.

The goal trajectory was generated by starting from the center of the box and independently changing the magnitude and direction of the velocity. Specifically, at time step n of duration $\Delta t = 0.1\,\mathrm{s}$ (corresponding to the length of a theta cycle) we first changed the direction of motion by adding a random 2-dimensional vector of length $\mu_{\Delta v} = 3.5\,\mathrm{cm/s}$ to the velocity $\bar{v}_n$. Next, we changed the speed $\bar{\nu} = |\bar{v}|$ according to an Ornstein-Uhlenbeck process:

$$\bar{\nu}_n = \bar{\nu}_{n-1} + (\mu_\nu - \bar{\nu}_{n-1})\frac{\Delta t}{\tau_\nu} + Q_\nu \epsilon \sqrt{\Delta t} \tag{2}$$

?e:goalv?

Table 1. Summary of the symbols used in the model.

symbol	meaning
n	index of time step or theta cycle
x_n	position at theta cycle n
y_n	sensory input
u_n	motor command
$\bar{x}_n$	goal position
$y_{1:n}$	past sensory input until theta cycle n
μ_n	mean of the filtering posterior
Σ_n	covariance of the filtering posterior
φ	theta phase
$\vec{\mathbf{x}}_n \equiv x_{(n-n_p):(n+n_f)}$	trajectory of the animal around theta cycle n
$\vec{\boldsymbol{\mu}}_n(\varphi)$	posterior mean trajectory at theta cycle n
$\vec{\mathbf{\Sigma}}_n(\varphi)$	posterior variance of trajectory at theta cycle n
$\tilde{\mathbf{x}}_n$	trajectory sampled from $P(\vec{\mathbf{x}}_n \| y_{1:n}, u_{1:n})$
$\phi_i(\vec{x})$	tuning function of cell i - firing rate as a function of the *encoded* position
$\psi_i(x)$	tuning curve of cell i - firing rate as a function of the *real* position
λ_i	firing rate of cell i
$\vec{\mathbf{s}}_n(\varphi)$	spikes recorded in theta cycle n encoding trajectory $\vec{\mathbf{x}}_n$
$\hat{\mathbf{x}}_n(\varphi)$	trajectory decoded from the observed spikes assuming direct encoding (Eq. 16)
$\hat{\boldsymbol{\mu}}_n(\varphi)$	estimated trajectory mean assuming DDC encoding (Eq. 17)
$\hat{\mathbf{\Sigma}}_n(\varphi)$	estimated trajectory variance assuming DDC encoding (Eq. 17)

?⟨t:symbols⟩?

with parameters $\mu_\nu = 16\,\text{cm/s}$, $\tau_\nu = 2\,\text{s}$, $Q_\nu = \sqrt{2/\tau_\nu}\,\sigma_\nu$, $\sigma_\nu = 10\,\text{cm/s}$ and $\epsilon \sim \mathcal{N}(0, 1)$ when $2\,\text{cm/s} \leq \nu \leq 80\,\text{cm/s}$ and $\epsilon = 0$ otherwise. The goal trajectory was generated by integrating the velocity signal:

$$\bar{x}_n = \bar{x}_{n-1} + \bar{v}_n\,\Delta t. \tag{3}$$

`e:goalx`

When the goal trajectory reached the boundary of the box, the trajectory was reflected from the walls by negating the corresponding component of the velocity. The goal trajectory influenced the motion of the simulated animal by providing an external control signal (motor command) as we describe it in the next paragraph.

Inference and movement in the generative model. We assumed that the simulated animal aims at following the goal trajectory, but does not have access to its own position, but has to infer it from his noisy sensory inputs. To follow the goal trajectory, the simulated animal calculated motor commands to minimise the deviation between its intended position and its estimated location. Specifically, we adapted the standard linear Gaussian state space model:

$$x_n = x_{n-1} + u_n + \varepsilon_u \qquad \varepsilon_u \sim \mathcal{N}(0, Q) \tag{4}$$

?e:evolution?

$$y_n = x_{n-1} + \varepsilon_y \qquad \varepsilon_y \sim \mathcal{N}(0, R) \tag{5}$$

{?}

where x_n is the true location of the animal, y_n is the sensory input and u_n is the external control (motor command), and Q and R are diagonal noise covariance matrices with $Q_{ii} = 2.25\,\text{cm}^2$ and $R_{ii} = 225\,\text{cm}^2$.

The estimated position was represented by the Gaussian filtering posterior:

$$P(x_n | y_{1:n}, u_{1:n}) = \mathcal{N}(\mu_n, \Sigma_n^2) \tag{6}$$

?e:filtPost?

where the parameters μ_n and Σ_n denote the mean and the covariance. The parameters were updated in each timestep using the standard Kalman filter algorithm[29] by combining the previous posterior with the current sensory input and the last motor command:

$$\begin{aligned} \mu_n &= \mu_{n-1} + u_n + K_n(y_n - (\mu_{n-1} + u_n)) \\ \Sigma_n &= (I - K_n)(\Sigma_{n-1} + Q) \end{aligned} \tag{7}$$

?e:filt?

where $K_n = (\Sigma_{n-1} + Q)(\Sigma_{n-1} + Q + R)^{-1}$ is the Kalman gain matrix.

The motor command was a low-pass filtered version of the deviation between the intended position ($\bar{x}_t$) and the estimated position of the animal (posterior mean, μ_{n-1}):

$$u_n = (1-\alpha)u_{n-1} + \alpha(\bar{x}_t - \mu_{n-1}) \tag{8}$$?e:drive?

with $\alpha = 0.25$.

To make predictions about future position we defined the distribution of trajectories consistent with all past observations and motor commands: $P(\vec{\mathbf{x}}_n | y_{1:n}, u_{1:n})$, where $\vec{\mathbf{x}}_n \equiv x_{(n-n_p):(n+n_f)}$ is a trajectory starting n_p steps behind in the past and ending n_f steps ahead in the future, and the lower index emphasizes that we are conditioning on the inputs and motor commands up to time step n (Figure 9B-C, Table 1). We call the distribution $P(\vec{\mathbf{x}}_n | y_{1:n}, u_{1:n})$ the *trajectory posterior*. Importantly, the trajectory posterior has a rich conditional dependence structure: To sample trajectories from the posterior distribution, we started each trajectory from the filtering posterior and proceeded first backwards, sampling from the conditional smoothing posterior, and then forward, sampling from the generative model.

To sample the past part of the trajectories ($m < n$) we capitalised on the following relationship:

$$\begin{aligned} P(x_{m-1} | x_m, y_{1:n}, u_{1:n}) &\propto P(x_{m-1} | y_{1:m-1}, u_{1:m-1})\, P(x_m | x_{m-1}, u_m)) \qquad (9) \\ &= \mathcal{N}\big(\Lambda_{m-1}(Q^{-1}(x_m - u_m) + \Sigma_{m-1}^{-1}\mu_{m-1}, \Lambda_{m-1}\big) \qquad (10) \end{aligned}$$

?e:smooth? {?}

where the first term of Equation 9 is the filtering posterior (Equation 6) and the second term is defined by the evolution process (Equation 4) and $\Lambda = (\Sigma_{m-1}^{-1} + Q^{-1})^{-1}$. We started each trajectory by sampling its first point independently from the filtering posterior (Equation 6) and applied Equation 9 recursively to elongate the trajectory backward in time and generate historical trajectories sampled from the smoothing posterior.

To generate samples in the forward direction ($m \geq n$) we implemented an ancestral sampling approach. First, a hypothetical goal trajectory was generated as in Equation 2-3 starting from the last goal location $\bar{x}_n$. Next, we calculated the hypothetical future motor command, u_{m+1} based on the difference between the next goal location $\bar{x}_{m+1}$ and current prediction for m, x_m as in Equation 8. Finally, we sampled the next predicted position from the distribution

$$P(x_{m+1} | x_m, u_{m+1}) = \mathcal{N}(x_m + u_{m+1}, Q). \tag{11}$$ {?}

To elongate the trajectory further into the future we repeated this procedure multiple times.

We introduce $\vec{\boldsymbol{\mu}}_n(\varphi) = \mathbb{E}[P(\vec{\mathbf{x}}_n | y_{1:n}, u_{1:n})]$ to denote the average over the possible trajectories that is the function of the theta phase φ. This notation emphasizes our central assumption, that past, present and future parts of the represented trajectory unfold sequentially within each theta cycle. Similarly, we also defined the variance of the trajectories, $\vec{\boldsymbol{\Sigma}}_n(\varphi)$. We used an approximate, diagonal covariance matrix and ignored the covariances between different trajectories and theta phase.

The motion profile of the simulated animal, including the distribution and the auto-correlation of the speed, acceleration and heading was similar to that of the real animals (Figure 8). Figure 9B-E illustrates the inference process in the model by showing a short segment of the true trajectory of the simulated animal centered on its location at time step n as well as trajectories starting in the past and extending into the future sampled from the posterior distribution. As expected, the variance of these hypothetical trajectories $\vec{\boldsymbol{\Sigma}}_n(\varphi)$ increased consistently from the past towards the future (from the beginning to the end of the theta cycle; illustrated by the increasing diameter of the ellipses in Figure 9C), while their mean $\vec{\boldsymbol{\mu}}_n(\varphi)$ tracked accurately the true trajectory of the animal (Figure 9C). Mean trajectories and trajectories sampled from the trajectory posterior at subsequent theta cycles are compared in Figure 9D-E (see also Figure 4A-B).

?⟨s:encoding⟩?

Encoding the posterior by the firing of place cells We assumed, that in each theta cycle the sequential activity of hippocampal place cells represents the trajectory posterior at an accelerated speed such that the encoded trajectory starts in the past at the beginning of the theta cycle and arrives to the predicted future states (locations) by the end of the theta cycle[15,23]. Each place cell i was associated with a tuning function $\phi_i(\vec{x})$ mapping the *encoded position* $\vec{x}$ to the firing rate of cell i. We also define the empirical tuning curve $\psi(x)$ as the average firing rate as a function of the *real position*. Importantly, in our model the firing rate of the place cells was controlled by the encoded trajectories and not the true location of the animal. The way the trajectories were encoded by the firing of the place cells was different in the four encoding schemes:

1. In the *mean* encoding scheme the instantaneous firing rate of the place cells is controlled by the mean of the trajectory posterior
$$\vec{\lambda}_i(\varphi) = \phi_i(\vec{\boldsymbol{\mu}}(\varphi)) \tag{12} \{?\}$$
That is, the cell's firing rate changes within the theta cycle according to the value of its tuning function at the different points of the mean trajectory.

2. In the *sampling* scheme, the encoded trajectory is sampled from the trajectory posterior, $\tilde{\mathbf{x}}(\varphi) \sim P(\vec{\mathbf{x}}_n|y_{1:n}, u_{1:n})$ and the instantaneous firing rate is the function of the sampled trajectory:
$$\vec{\lambda}_i(\varphi) = \phi_i(\tilde{\mathbf{x}}(\varphi)) \tag{13} \{?\}$$

3. In the *product* scheme the firing rate is controlled by both the trajectory's posterior mean and variance:
$$\vec{\lambda}_i(\varphi) = \phi_i(\vec{\boldsymbol{\mu}}(\varphi))\,\frac{\sigma_0}{\vec{\sigma}(\varphi)} \tag{14} \{?\}$$
where $\sigma = (\det \Sigma)^{1/4}$ and $\sigma_0 = 5.5\,\mathrm{cm}$. This is similar to the MAP encoding model, except that the population firing rate is scaled by the inverse of the posterior uncertainty. Note, that in standard PPC codes[16] the population gain is scaled by the inverse of the variance (instead of SD) which would predict much larger changes in the firing rate which is inconsistent with the experimental data.

4. In the *DDC* scheme the instantaneous firing rate of cell i is the expectation of the basis function i under the encoded time point of the trajectory posterior:
$$\vec{\lambda}_i(\varphi) = \int \phi_i(\vec{x})\, P(\vec{x}(\varphi)|y_{1:n}, u_{1:n})\, d\vec{x} \tag{15} ?\text{e:ddc}?$$

In each encoding models, spikes were generated from the instantantous firing rate $\vec{\lambda}_i(\varphi)$ as an inhomogeneous Poisson process:
$$\vec{s}_i(\varphi) \sim \mathrm{Poisson}(\vec{\lambda}_i(\varphi))$$

The population activity in all of the four encoding schemes is consistent with many defining characteristics of hippocampal data during exploratory behavior e.g., neurons show place cell activity, phase precession and theta sequences (Figure 4D-E, Figure 10).

Decoding

To discriminate the different encoding schemes from each other we decoded the hippocampal neuronal activity in two different ways: first, assuming that a single position is encoded by the population activity (consistent with the mean and sampling schemes) and second, assuming that a distribution is encoded via the DDC scheme. We used static decoders to ensure that the variance of the decoder is independent of the theta phase as opposed to dynamic decoders, where the variance can be larger around the ends of the trajectories.

Single point decoding. We performed static Bayesian decoding independently in each temporal bin at different phases of the theta cycle. The estimated position at theta phase φ is then:
$$\begin{aligned}\hat{x}(\varphi) &= \int x\, P(x|\mathbf{s}(\varphi))\, dx \\ &= \int x\, \frac{\prod_i P(s_i(\varphi)|x) P(x)}{P(\mathbf{s})}\, dx \end{aligned} \tag{16} \boxed{\text{e:dec\_direct}}$$
where prior $P(x)$ was estimated by the 5 cm binned position of the animal and we used a Poisson likelihood with spatial tuning curves $\psi_i(x)$, estimated from the 5 cm binned spike counts. Both the prior and the tuning curves were smoothed with a 10 cm Gaussian kernel. We binned spikes either into windows of fixed duration (20 ms, Figure 3A-C) fixed theta phase (120°, Figure 5) or into three bins with equal number of spikes within a theta cycle (Figure 6-7 and everywhere else).

DDC decoding. The DDC decoder assumes that at each time point in the theta cycle an isotropic Gaussian distribution over the locations is encoded by the population activity via DDC encoding Equation 15. To ensure that the theta phase dependence of the firing rates do not introduce a bias in the decoded variance, we divided each theta cycle to three windows (early, middle and late) with equal number of spikes. As linear decoding of DDC codes from spikes is very inaccurate, we performed independent maximum likelihood decoding of the parameters of the encoded distribution. The estimated mean and the variance in bin φ is then:

$$\{\hat{\boldsymbol{\mu}}(\varphi),\, \hat{\sigma}(\varphi)^2\} = \underset{\boldsymbol{\mu},\, \sigma^2}{\arg\max} \prod_i P(\mathbf{s}(\varphi)|\lambda_i(\boldsymbol{\mu}, \sigma^2)) \tag{17}$$ `e:dec_ddc`

where

$$\lambda_i(\mu, \sigma^2) = \int \phi_i(x)\, \mathcal{N}(x; \mu, \sigma^2)\, dx \tag{18}$$ {?}

where $\phi_i(x)$ is ideally the tuning function of neuron i used during the encoding process (Equation 15). However, in practice we do not have access to the true tuning functions but only to the empirical tuning curves, $\psi(x)$. We found that the tuning curves are typically more dispersed than the tuning functions when the encoded and the true location is not the same or when the encoded distributions have nonzero variance (Figure 12C, F). The difference between the size of the tuning functions used for encoding and decoding can introduce a substantial bias in decoding the variance of the distribution (Figure 12A-C). To reduce this bias, we devised a non-parametric deconvolution algorithm that could efficiently sharpen the tuning curves to approximate the original tuning functions for a large variety of tuning curves shapes (Figure 12D,F). Using these sharpened tuning curves substantially reduced the bias of decoding the variance in synthetic data (Figure 12E).

Sampling index

We introduced the sampling index to discriminate the mean scheme fro the sampling scheme. The sampling index is the difference between the cycle-to-cycle variabilityand the trajectory encoding errorand is positive for sampling and negative for the mean scheme. In this section we provide definitions of these quantities, derive their expected value for the sampling and the mean scheme and show how the sampling index can be estimated from data.

Cycle-to-cycle variability (*CCV*). We defined cycle-to-cycle variability as the difference between the trajectories decoded from the neural activity in two subsequent theta cycles (Figure 9E):

$$CCV_n(\varphi) = \left(\hat{\mathbf{x}}_n(\varphi) - \hat{\mathbf{x}}_{n-1}(\varphi)\right)^2 \tag{19}$$?e:ccv?

where φ indicates theta phase. The expected value of χ in the case of MAP encoding is

$$\mathbb{E}_{\mathrm{MAP}}[CCV(\varphi)] = 2\epsilon^2 + \mathbb{E}[\zeta^2(\varphi)] \tag{20}$$?e:traj_ch_MAP?

where $\zeta^2 = (\vec{\boldsymbol{\mu}}_n - \vec{\boldsymbol{\mu}}_{n-1})^2$ is the change of the mean encoded trajectory between subsequent theta cycles and ϵ^2 is the variance of the decoder. Note, that on average ζ increases with φ as new observations have larger effect on uncertain future predictions than on the estimated past positions. The variance of the decoder reflects the finite number of observed neurons in the population and their stochastic spiking. As we divide each theta cycle to equal spike count bins, the decoder variance is independent of the theta phase φ.

In the sampling scheme, when the population activity encodes independent samples drawn from the trajectory posterior, the cycle-to-cycle variabilityalso depends on the posterior variance $\vec{\boldsymbol{\Sigma}}_n(\varphi)$:

$$\mathbb{E}_{\mathrm{sam}}[CCV(\varphi)] = 2\epsilon^2 + 2\mathbb{E}[\vec{\boldsymbol{\sigma}}^2(\varphi)] + \mathbb{E}[\zeta^2(\varphi)] \tag{21}$$?e:traj_ch_sam?

where $\vec{\boldsymbol{\sigma}}_n^2(\varphi) = \sqrt{\det\left(\vec{\boldsymbol{\Sigma}}_n(\varphi)\right)}$. Thus the hallmark of the sampling scheme is $CCV(\varphi) > 2\epsilon^2 + \zeta(\phi)$. However, we can only measure CCV and estimate ϵ, but don't have access to ζ nor $\vec{\boldsymbol{\sigma}}^2$ in the experimental data to discriminate sampling from the mean scheme. Therefore we defined another measure, the trajectory encoding errorto estimate the posterior variance $\vec{\boldsymbol{\sigma}}^2$.

Trajectory encoding error (*TEE*). We defined trajectory encoding erroras the difference between the true 2-dimensional trajectory of the rat $\mathbf{x}_n$, and the trajectory decoded from the neural activity $\hat{\mathbf{x}}_n$ (Figure 9E):

$$\varepsilon(\varphi) = \big(\vec{\mathbf{x}}_n - \hat{\mathbf{x}}_n(\varphi)\big)^2 \tag{22}$$

`?e:enc_err?`

When comparing decoded and physical trajectories, we assumed a fixed correspondence between theta phase φ and temporal shift along the trajectory. Specifically, for each animal we calculated the decoding error for early, mid and late phase spikes as a function of the temporal shift of the animal's position (Figure 3D) and compared the decoded position with the true location at a fixed time shift that minimised the average error at that phase.

The expected value of *TEE* in the case of MAP encoding is

$$\begin{aligned} \mathbb{E}_{\mathrm{MAP}}[\mathit{TEE}(\varphi)] &= \mathbb{E}\Big[\big(\vec{\mathbf{x}}_n - \vec{\boldsymbol{\mu}}_n(\varphi) + \vec{\boldsymbol{\mu}}_n(\varphi) - \hat{\mathbf{x}}_n(\varphi)\big)^2\Big] \\ &= \mathbb{E}[\vec{\boldsymbol{\sigma}}^2(\varphi)] + \epsilon^2 \end{aligned} \tag{23}$$

`{?}`

where we used the fact that the encoded trajectory is the mean $\vec{\boldsymbol{\mu}}_n(\varphi)$, and if the model of the animal is consistent, then the expected difference between the posterior mean and the true location equals the variance of the posterior, $\mathbb{E}[(\mathbf{x}_n - \vec{\boldsymbol{\mu}}_n)^2] = \vec{\boldsymbol{\sigma}}^2$.

In the case of sampling the encoded trajectory is $\tilde{\mathbf{x}}_n$ and the *TEE* is increased by the difference between the mean and the sampled trajectory:

$$\begin{aligned} \mathbb{E}_{\mathrm{sam}}[\mathit{TEE}(\varphi)] &= \mathbb{E}\Big[\big(\vec{\mathbf{x}}_n - \vec{\boldsymbol{\mu}}_n(\varphi) + \vec{\boldsymbol{\mu}}_n(\varphi) - \hat{\mathbf{x}}_n(\varphi))\big)^2\Big] \\ &= \mathbb{E}[\vec{\boldsymbol{\sigma}}^2(\varphi)] + \mathbb{E}\Big[\big(\vec{\boldsymbol{\mu}}_n(\varphi) - \tilde{\mathbf{x}}_n(\varphi) + \tilde{\mathbf{x}}_n(\varphi) - \hat{\mathbf{x}}_n(\varphi)\big)^2\Big] \\ &= 2\mathbb{E}[\vec{\boldsymbol{\sigma}}^2(\varphi)] + \epsilon^2 \end{aligned} \tag{24}$$

`{?}`

Sampling index. We define the *sampling index* as

$$\xi(\varphi) = \mathbb{E}[\overline{\mathit{CCV}}(\varphi)] - \mathbb{E}[\overline{\mathit{TEE}}(\varphi)] \tag{25}$$

`?e:SI?`

where $\overline{\mathit{CCV}} = \mathit{CCV} - 2\epsilon^2$ and $\overline{\mathit{TEE}} = \mathit{TEE} - \epsilon^2$. It can be shown, that $\mathbb{E}[\zeta^2(\varphi)] = \alpha\,\mathbb{E}[\vec{\boldsymbol{\sigma}}^2(\varphi)]$ with $0 < \alpha < 1$. Therefore the sampling index is positive for sampling and negative for MAP:

$$\xi_{\mathrm{MAP}}(\varphi) = (\alpha - 1)\,\mathbb{E}[\vec{\boldsymbol{\sigma}}^2(\varphi)] < 0 \tag{26}$$

`{?}`

$$\xi_{\mathrm{sam}}(\varphi) = \alpha\,\mathbb{E}[\vec{\boldsymbol{\sigma}}^2(\varphi)] > 0 \tag{27}$$

`{?}`

As $\vec{\boldsymbol{\sigma}}^2$ is expected to increases within a theta cycle, the sampling index is most distinctive at late theta phases. To estimate the decoding error ϵ^2 we started with the observation that at early theta phases when $\vec{\boldsymbol{\sigma}}^2$ is small and both *TEE* and *CCV* are dominated by ϵ. Thus we approximated

$$\mathbb{E}[\overline{\mathit{CCV}}(\varphi)] \approx \mathbb{E}[\mathit{CCV}(\varphi)] - \mathbb{E}[\mathit{CCV}(\text{early})] \tag{28}$$

`{?}`

$$\mathbb{E}[\overline{\mathit{TEE}}(\varphi)] \approx \mathbb{E}[\mathit{TEE}(\varphi)] - \mathbb{E}[\mathit{TEE}(\text{early})] \tag{29}$$

`{?}`

Our analysis on simulated data also confirmed that estimating the decoding error this way is more robust than using the Cramér-Rao bound, as $\mathbb{E}[\mathit{CCV}(0)]$ also contains non-trivial second order terms, e.g., correlation between the decoding error in two subsequent theta cycles imposed by arena borders. These second order terms have a differential contribution to *CCV* and *TEE*, but are approximately constant within theta cycles. We also found that the actual decoding error is slightly larger than the error estimated from the Fisher information and underestimating ϵ biases the Sampling index towards more positive values.

Therefore throughout the paper we evaluated the sampling index as the difference between the *CCV* and *TEE*evaluated on late vs. early spikes:

$$\xi = \mathbb{E}[\overline{\mathit{CCV}}(\text{late})] - \mathbb{E}[\overline{\mathit{TEE}}(\text{late})] \tag{30}$$

`?e:SI2?`

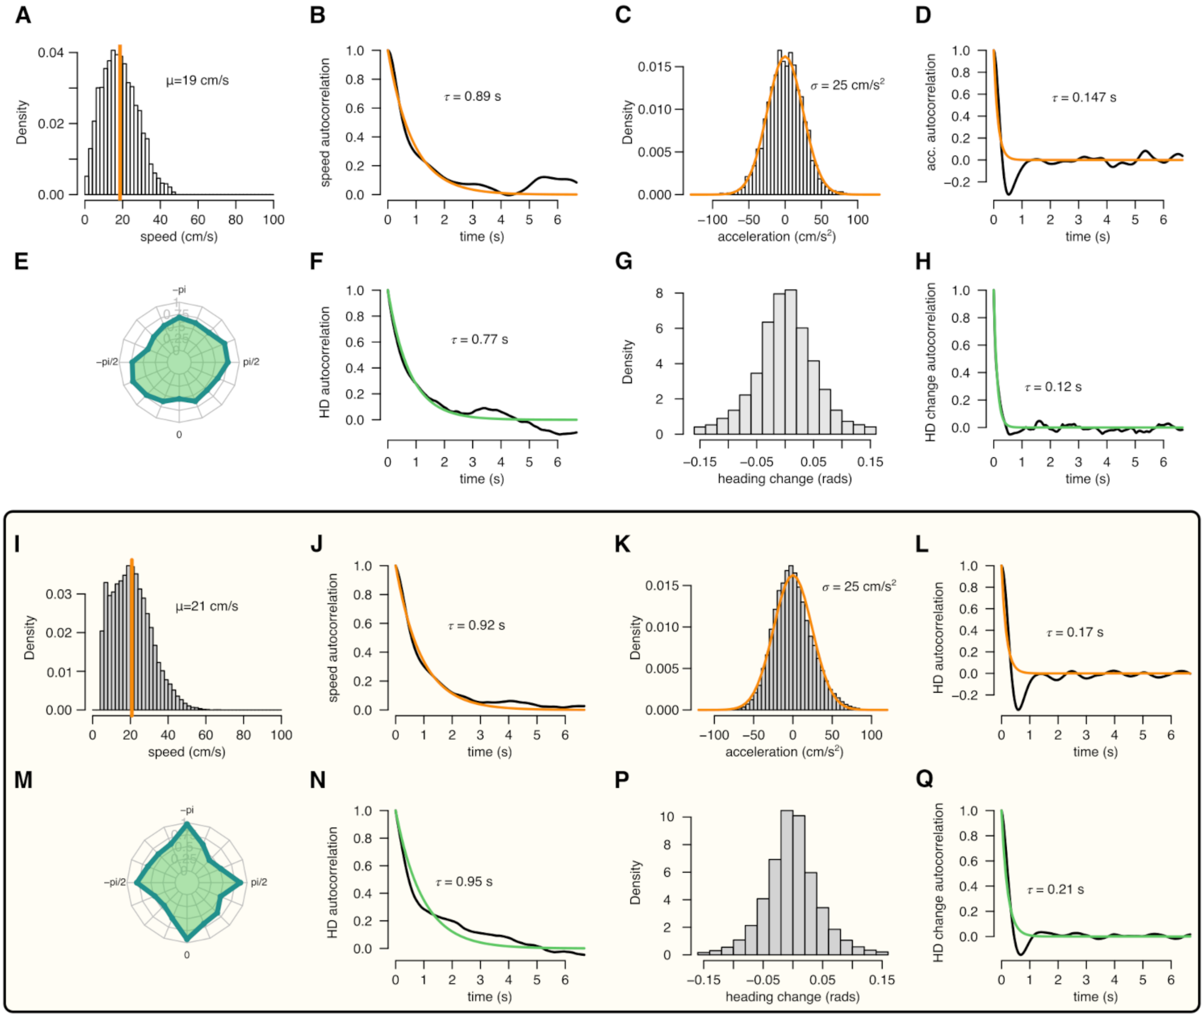

Figure 8. Comparison of the motion profile of the simulated animal and one of the analysed experimental sessions. A-H Motion profile of the simulated animal. **A** Histogram of the running speed. **B** Autocorrelation of the running speed. **C** Histogram of the acceleration. **D** Autocorrelation of the acceleration. **E** Polar plot if the running direction. **F** Autocorrelation of the head direction. **G** Histogram of the head direction change. **H** Autocorrelation of the head direction change. **I-Q** Same as A-H for experimental session rat 1 day 2. Note, that for the experimental data we only analysed continuous running periods, where $v > 5$ cm/s.

?⟨sf:motion⟩?

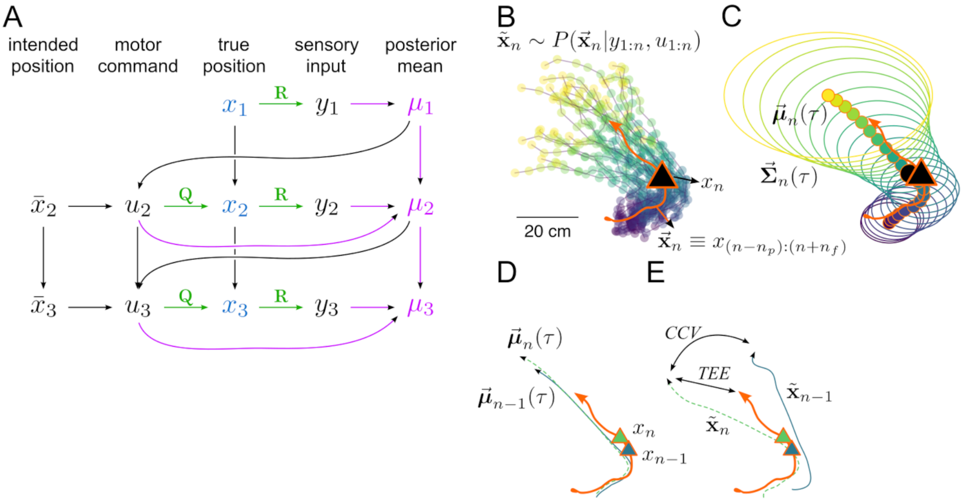

Figure 9. Inference and movement in the generative model. A Graphical model of the processes underlying the generation of the simulated animal's trajectory. At the unknown true location, x_t, the animal receives a noisy sensory input, y_t, and combines it with the previous motor command, u_t, to update its current estimate of its location $P(x_t|y_{1:t}, u_{1:t}) = \mathcal{N}(\mu_t, \Sigma_t)$. Based on its current position estimate and the intended position for the next time step, $\bar{x}_{t+1}$, it calculates the next motor command, u_{t+1}, that is used to generate the new position x_{t+1}. Green arrows represent sensory and motor noise. Purple arrows represent inference of $P(x_t|y_{1:t}, u_{1:t})$ represented here by the posterior mean μ_t. **B** The true trajectory of the simulated animal (orange) with it inferred past (blue) and predicted future (yellow) positions (colors) represented as 25 trajectories sampled at position x_t (triangle). **C** Same as B except that distribution of trajectories are represented by covariance ellipses. **D** True trajectory (orange) and posterior mean trajectories at two subsequent timesteps (theta cycles, blue and green). **E** Similar to D but with trajectories sampled independently from the corresponding posterior distribution. Trajectory encoding error(TEE) denotes the difference between the true and the sampled trajectory whereas cycle-to-cycle variability(CCV) is the difference between trajectories sampled in subsequent theta cycles.

`?⟨sf:repr⟩?`

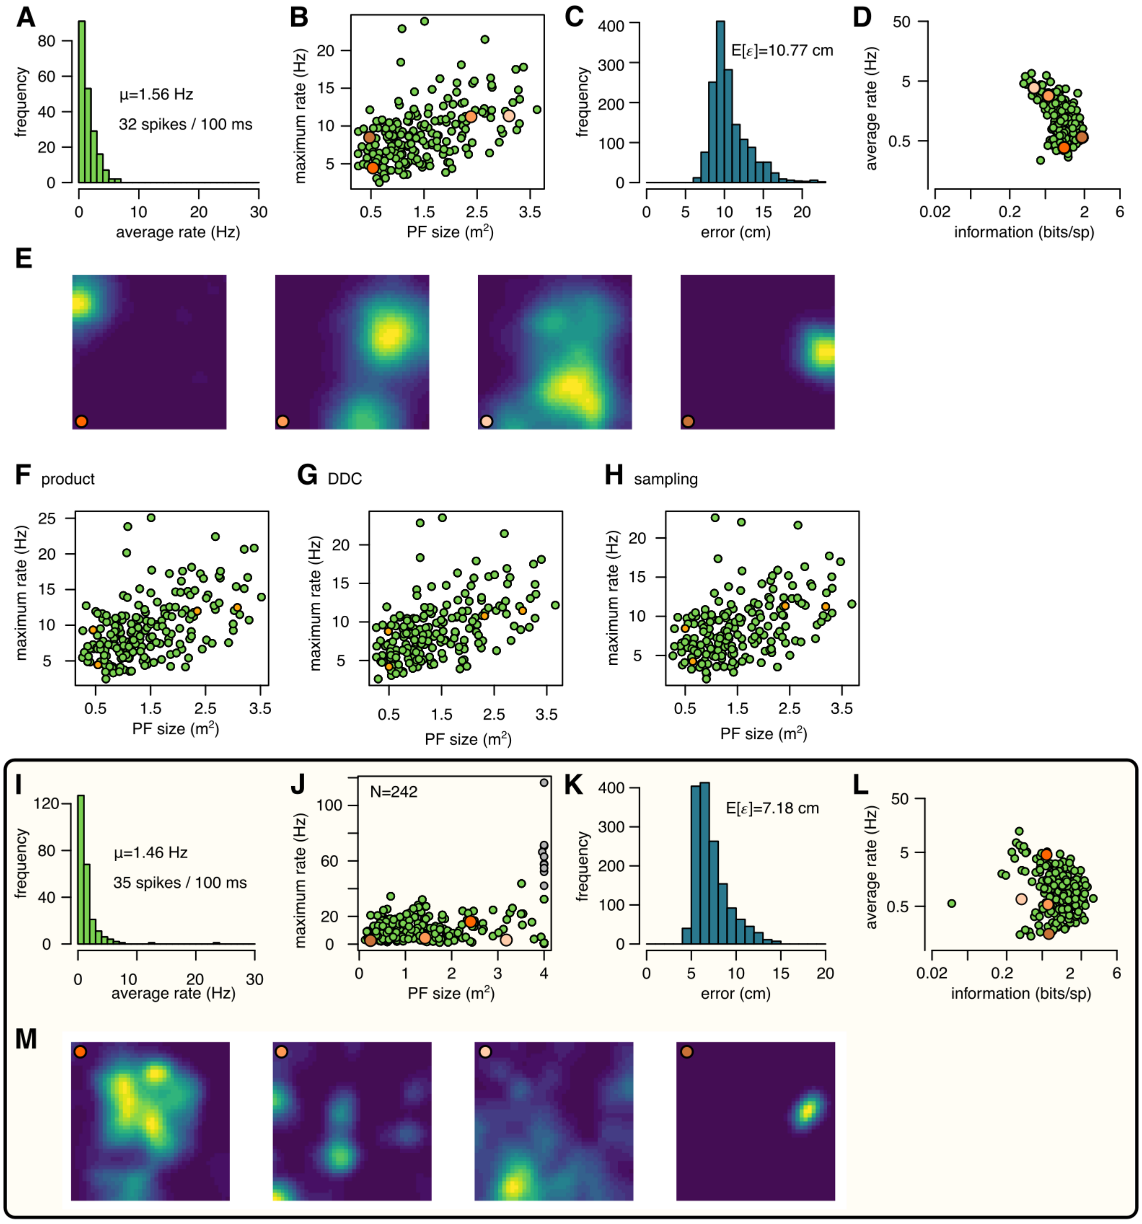

Figure 10. Place cell firing in the synthetic and in the experimental data. A-E Place cell activity in the simulated data using MAP encoding. **A** Histogram of the average firing rate of the 200 simulated neurons. Inset shows the average firing rate across cells and the average number of spikes in a 100 ms theta cycle. **B** Maximum firing rate as a function of place field size (area over the 10%of the maximum rate). Orange indicates cells with place fields shown in E. **C** Histogram of the estimated Fisher lower bound for decoding error. **D** Average firing rate as a function of spatial information. **E** Normalised ratemaps of the 4 cells selected (orange in panels B and D). **F-H** Similar to panel B for for the product (F), DDC (G) and the sampling (H) encoding scheme. **I-M** Same as A-E for experimental session rat 1 day 2. Note that panel J includes putative inhibitory cells (grey) whereas other panels show only putative excitatory neurons. See also Supplementary Figure 3 of Pfeiffer and Foster (2013)[13].

?⟨sf:match⟩?

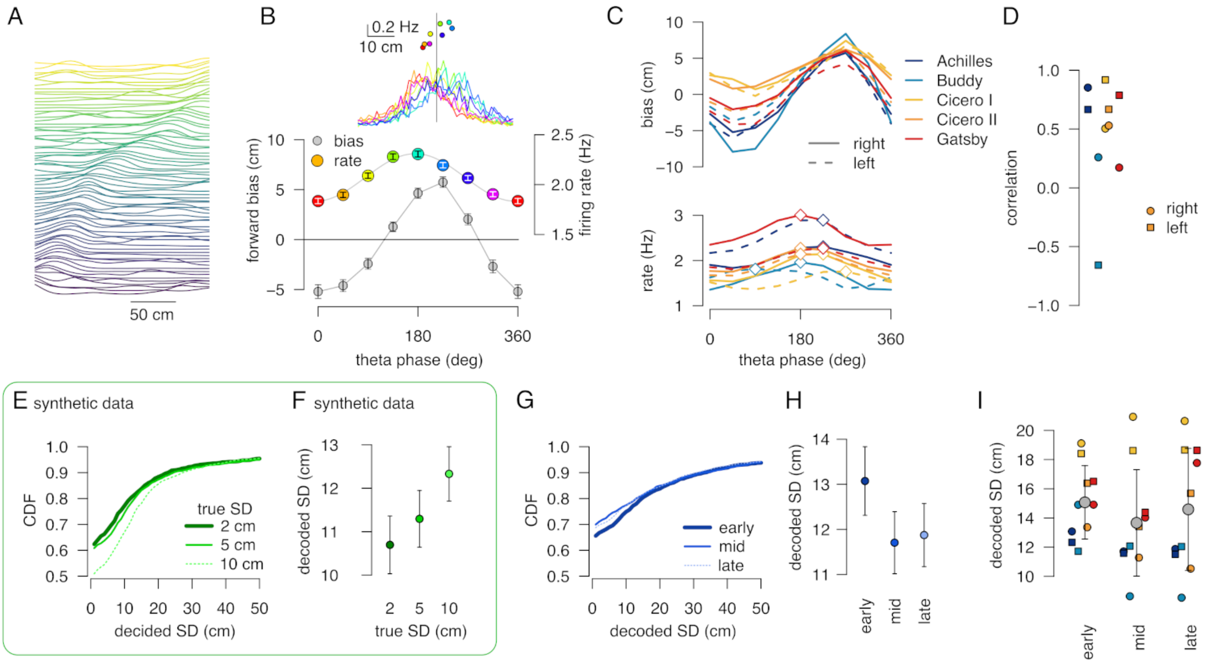

Figure 11. Linear track data and analysis. A Normalised firing rate of all putative excitatory neurons recorded in a single session (Achilles, up) order by the location of the peak activity on rightward runs. **B** Median forward bias and average firing rate as a function of the theta phase calculated in 120° windows with 45° shift. Inset: the distribution of the decoded location relative to the position of the animal (vertical line) as a function of theta phase (color) and the firing rate as a function of the median decoded location (forward bias, bottom). Error bars show SEM. **C** Median forward bias (top) and average firing rate (bottom) as a function of the theta phase for all recording sessions. Peak forward bias is shifted to 270°, diamonds indicate the peak of the firing rate. Color indicate animal name and line style indicates running direction. **D** Correlation between median forward bias and firing rate is typically positive except for one outlier. We note, that the number of recorded neurons was the lowest for the rat Buddy (N=15 cells with place field on right and N=16 on left runs.) **E** Cumulative distribution of the estimated SD for different distributions encoded with a 1 dimensional DDC scheme. Gaussian distributions of different SD were encoded using the tuning functions derived from the place fields shown in A with the same number of theta cycles (N=1846) and bin width ($\Delta t = 36\,\text{ms}$) as in the experimental data (G-H). **F** Mean decoded SD at different values of the true SD. At this short time window the decoder is still biased (see Figure 6), but the decoded SD correlates with the true SD. Error bars show SEM. **G** Cumulative distribution of the decoded SD for early, mid and late theta phases for the same session as A-B. We used sharpened tuning functions to reduce the decoding bias. **H** Mean decoded SD at different theta phases. Error bars show SEM. **I** Mean decoded SD at different theta phases for all recorded sessions. Color code is the same as in panel C-D. Only one of the animals (Gatsby) showed an increase in the decoded SD within the theta cycle. Grey symbols show mean and SD across animals.

`?⟨sf:linear⟩?`

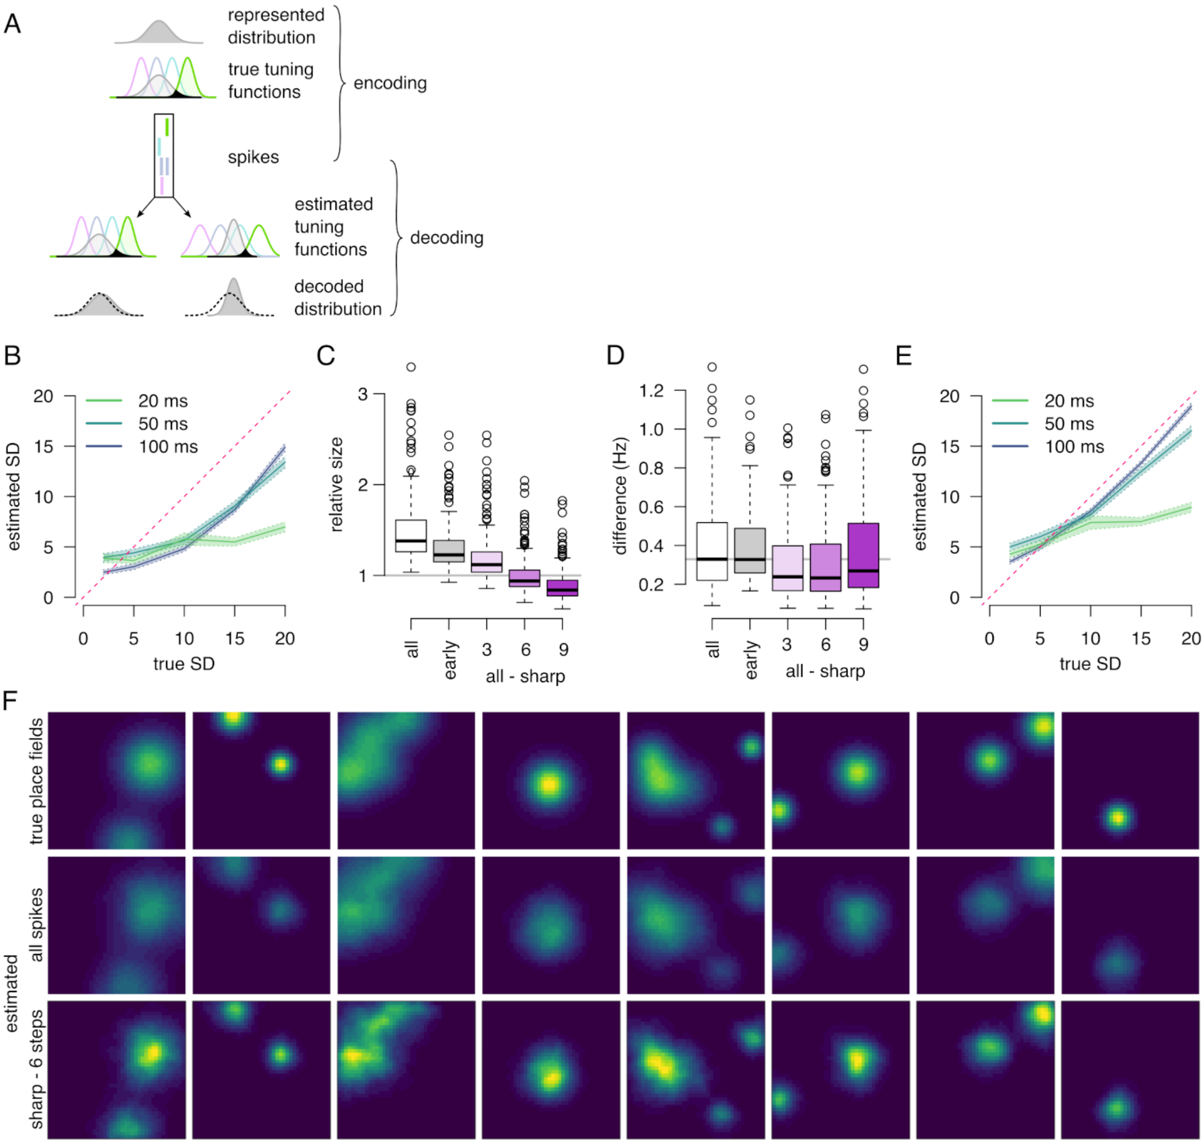

Figure 12. Reducing the bias of decoding the SD in the DDC scheme. A Illustration of the problem of decoding bias. In the DDC scheme the firing rate of the neurons is proportional to the overlap between the encoded distribution and the tuning curve of the neurons (black shading). It is possible to approximately reconstruct the encoded distribution from the spikes observed in the population activity and the tuning functions (bottom left). However, when the estimated tuning functions are different than the original, the reconstruction will be biased: Wider tuning functions showing similar overlap with a narrower distribution thus lead to systematic underestimation of the variance of the encoded distribution (bottom right). **B** Systematic bias in decoding the SD of a 2-dimensional Gaussian distribution in a spike-based DDC scheme with different temporal windows. Estimated tuning functions become wider than the true tuning functions due to three factors: encoding other than the current position, encoding distributions of large variance and smoothing during the estimation. **C** Ratio of the estimated and the true tuning function sizes (place field sizes) for different estimation strategies: all spikes, only spikes from early theta phase and after 3-9 sharpening steps. The size of the estimated tuning curve is closest to the original after 6 sharpening steps. **D** Average absolute difference between the true tuning function and the estimated across 200 cells is decreased after algorithmic sharpening the tuning curves. The difference is minimal after 3-6 sharpening steps. **E** Same as B after sharpening the estimated tuning functions used for decoding. (Figure 6B, right repeated here for comparision.) **F** Examples for the original 2-dimensional tuning functions in the synthetic data (top), estimated tuning functions from all observed spikes (middle) and after 6 sharpening steps (bottom).

?⟨sf:mix⟩?

References

`ujfalussy18` [1] Ujfalussy, B.B., Makara, J.K., Lengyel, M. & Branco, T. Global and multiplexed dendritic computations under in vivo-like conditions. *Neuron (accepted)* (2018).

`ujfalussy20` [2] Ujfalussy, B.B. & Makara, J.K. Impact of functional synapse clusters on neuronal response selectivity. *Nat Commun* **11**, 1413 (2020).

`mago20` [3] Magó, Á., Weber, J.P., Ujfalussy, B.B. & Makara, J.K. Synaptic plasticity depends on the fine-scale input pattern in thin dendrites of ca1 pyramidal neurons. *J. Neurosci.* **40**, 2593–2605 (2020).

`glimcher03` [4] Glimcher, P. *Decisions, Uncertainty, and the Brain* (MIT Press, 2003).

`doya06` [5] (eds. Doya, K., Ishii, S., Pouget, A. & Rao, R.P.N.). *Bayesian Brain: Probabilistic Approaches to Neural Coding* (MIT press, 2006).

`wallenstein98` [6] Wallenstein, G.V., Eichenbaum, H. & Hasselmo, M.E. The hippocampus as an associator of discontiguous events. *Trends in neurosciences* **21**, 317–323 (1998). PMID: 9720595.

`suh11` [7] Suh, J., Rivest, A.J., Nakashiba, T., Tominaga, T. & Tonegawa, S. Entorhinal Cortex Layer III Input to the Hippocampus Is Crucial for Temporal Association Memory. *Science* **334**, 1415–1420 (2011). PMID: 22052975.

`kitamura14` [8] Kitamura, T. *et al.* Island Cells Control Temporal Association Memory. *Science* **343**, 896–901 (2014).

`miller17` [9] Miller, K.J., Botvinick, M.M. & Brody, C.D. Dorsal hippocampus contributes to model-based planning. *Nat. Neurosci.* **20**, 1269–1276 (2017).

`frank00` [10] Frank, L.M., Brown, E.N. & Wilson, M. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* **27**, 169–178 (2000).

`wood00` [11] Wood, E.R., Dudchenko, P.A., Robitsek, R.J. & Eichenbaum, H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–33 (2000).

`lisman09` [12] Lisman, J. & Redish, A.D. Prediction, sequences and the hippocampus. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **364**, 1193–201 (2009).

`pfeiffer13` [13] Pfeiffer, B.E. & Foster, D.J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).

`gupta12` [14] Gupta, A.S., van der Meer, M.A.A., Touretzky, D.S. & Redish, A.D. Segmentation of spatial experience by hippocampal θ sequences. *Nature neuroscience* **15**, 1032–1039 (2012).

`foster07` [15] Foster, D.J. & Wilson, M.A. Hippocampal theta sequences. *Hippocampus* **17**, 1093–1099 (2007). PMID: 17663452.

`ma06` [16] Ma, W., J.M., B., Latham, P. & Pouget, A. Bayesian inference with probabilistic population codes. *Nat. Neurosci.* **9**, 1432–1438 (2006).

`zemel98` [17] Zemel, R., Dayan, P. & Pouget, A. Probabilistic interpretation of population codes. *Neural Comput.* **10**, 403–430 (1998).

`vertes18` [18] Vértes, E. & Sahani, M. Flexible and accurate inference and learning for deep generative models. *Advances in Neural Information Processing Systems*, 4166–4175 (2018).

`wainwright08` [19] Wainwright, M.J. & Jordan, M.I. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning* **1**, 1–305 (2008).

`fiser10` [20] Fiser, J., Berkes, P., Orban, G. & Lengyel, M. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.* **14**, 119–30 (2010).

`prml` [21] Bishop, C. *Pattern Recognition and Machine Learning* (Springer, 2006).

`grosmark16b` [22] Grosmark, A., Long, J. & Buzsáki, G. Recordings from hippocampal area CA1, PRE, during and POST novel spatial learning. *CRCNS.org* **http://dx.doi.org/10.6080/K0862DC5** (2016).

`skaggs96` [23] Skaggs, W.E., McNaughton, B.L., Wilson, M.A. & Barnes, C.A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–72 (1996).

`fenton98` [24] Fenton, A.A. & Muller, R.U. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proc Natl Acad Sci U S A* **95**, 3182–7 (1998).

`kay20` [25] Kay, K. *et al.* Constant sub-second cycling between representations of possible futures in the hippocampus. *Cell* **180**, 552–567.e25 (2020).

`feng15` [26] Feng, T., Silva, D. & Foster, D.J. Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *J. Neurosci.* **35**, 4890–902 (2015).

`orban16` [27] Orbán, G., Berkes, P., Fiser, J. & Lengyel, M. Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron* **92**, 530–543 (2016).

`zhang98` [28] Zhang, K., Ginzburg, I., McNaughton, B.L. & Sejnowski, T.J. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* **79**, 1017–44 (1998).

`murphy12` [29] Murphy, K.P. *Machine Learning: A Probabilistic Perspective* (MIT Press, Cambridge, MA, 2012).