Sampling-based inference and predictions in the hippocampus.

OTKA FK-125324 Final Report

Balázs B Ujfalussy
ujfalussy.balazs@koki.hu
Laboratory of Biological Computation, Institute of Experimental Medicine, Budapest, Hungary

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 was 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 is represented by encoding multiple alternative options sequentially.

With the Lab. of Neuronal Signalling, we set up virtual reality experiments and recorded the activity of hippocampal neurons using Ca imaging. We studied how hippocampal spatial and contextual representations develop in mice during learning a navigation task. We found, that although hippocampal spatial maps form early upon exploring an environment, emergence of task-relevant variables requires further experience.

We also analysed the contribution of dendritic nonlinearities to the spatially selective activity of hippocampal place cells. We found that the contribution of dendritic nonlinearities to the neural output is maximal when synaptic clusters, formed via local plasticity, are moderately large. Moreover, we found that synaptic inputs onto the same dendritic tree are processed by parallel functional architectures with distinct computational properties.

In the next sections I give a succinct description of the achievements of during the 5 years of the project. Further details can be found in our publications. As the workplan of the PD-125386 and the FK-125324 grants were largely overlapping, Sections 1 and 3.1 were also included in the final report of the PD grant.

1 Sampling-based representation of uncertainty during hippocampal theta sequences (PD and FK).

The aim of our proposal was to identify how neuronal activity during hippocampal theta oscillations contributes to efficient planning by representing the uncertainty associated with model based predictions. The first two steps to achieve this goal was to develop measures to discriminate competing forms of uncertainty representation (Aim 1, PD and FK) and to test them in previously recorded data (Aim 2, PD and FK). We completed this goal and presented the results on multiple conferences (Cosyne 2019, Areadne 2018, CCCN 2018, Bernstein Conference 2020) and is published as a Research Article in eLife (ref. 1). Since the work is published, I only provide a short summary here.

The hippocampus has been established as one of the brain areas critically involved in both spatial navigation and more abstract planning[2,3]. A crucial insight gained about the neural code underlying navigation is that neuron populations in the hippocampus represent the trajectory of the animal on multiple time scales: Not only the current position of the animal can be read out at the behavioral time scale[4], but also trajectories starting in the past and ending in the near future are repeatedly expressed on a shorter time scale at accelerated speed during individual cycles of the 6-10 Hz theta oscillation (Figure 1a, theta sequences[5]). These data provide strong support for a computational framework where planning relies on sequential activity patterns in the hippocampus delineating future locations based on the current beliefs of the animal[3,6]. Whether hippocampal computations also take into account the uncertainty associated with planning and thus the population activity represents the uncertainty of the encoded trajectories has not been studied previously.

Neuronal representations of uncertainty have been extensively studied in sensory systems[7–10]. Schemes for representing uncertainty fall into three broad categories (Figure 1b). In the first two categories (*product*[7] and Distributed Distributional Code[9], *DDC*, schemes) the firing rate of a population encodes a complete probability distribution over spatial locations *instantaneously* at any given time by representing the

parameters of the distribution[11]. In the third category, the *sampling scheme*, the population activity represents a single value sampled stochastically from the target distribution. In this case uncertainty is represented *sequentially* by the across-time variability of the neuronal activity[12].

To analyse the distinctive properties of the different coding schemes and to develop specific measures capable of discriminating them, we generated a synthetic datasets in which both the behavioural component and the neural component was precisely matched to experimental data. Then we tested these measures using both our synthetic datasets and place cell activity recorded in rats performing an open field navigation task[13].

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, when the encoding basis functions are localised, the width of the encoded distribution tends to decrease with the total number of spikes in the population[7]. In contrast with this prediction, the population firing rate did not decrease for predictions extending further into the future in the hippocampal data (Figure 1c-d).

In the DDC scheme, neuronal firing rate represents the overlap of the basis functions with the encoded distribution[9,14]. We used a maximum likelihood decoder to estimate the width of the distribution represented in the population activity. Although the decoded SD increased with theta phase in our synthetic DDC dataset (Figure 1e), we found no evidence for such systematic change during hippocampal theta sequences (Figure 1f).

In the sampling scheme the cycle-to-cycle variability of the neuronal activity is increased, compared to the mean scheme, by the stochasticity of the represented trajectory. We derived that this excess variability can be measured if we compare the cycle-to-cycle variability with the deviation between the encoded and the real trajectory of the animal (Figure 1g). The excess variability index (EV-index) reliably discriminates sampling based codes from mean codes across a wide range of parameters (Figure 1h-i). We found that the EV-index was consistently and significantly positive for all recording sessions (Figure 1j) indicating that the data is consistent with random sampling potential motion trajectories in successive theta cycles.

The magnitude of the EV-index was substantially larger when evaluated on real data than in our synthetic datasets, suggesting that the efficiency of the sampling algorithm is increased by reducing the correlation between trajectories sampled in neighbouring theta cycles[15] ("theta cycling"). We introduced the cycling index to measure the tendency of the encoded future positions to alternate around the real trajectory of the animal. We found that the cycling index was significantly positive for most of the recording sessions indicating frequent alternation and an efficient sampling process (Figure 1k-l).

We have shown that the hippocampal population code displays the signature of efficient sampling of possible motion trajectories but could not identify the hallmarks of alternative proposed schemes for coding uncertainty. Our analysis demonstrated that the neural code in the hippocampus shows hallmarks of probabilistic planning by representing information about uncertainty associated with the encoded positions.

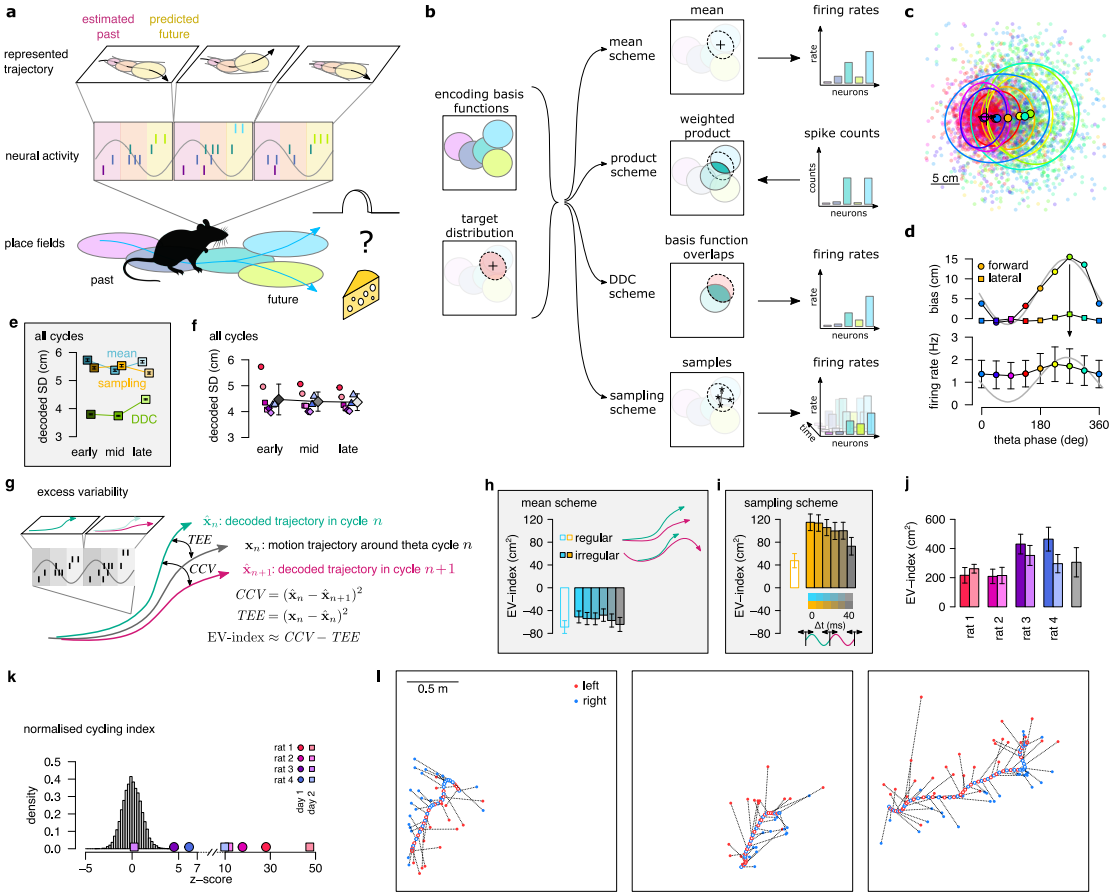

Figure 1. Theta sequences and uncertainty. **a** Hippocampal place cell activity representing possible trajectories in subsequent theta cycles during navigation. **b** Alternative schemes for representing a target probability distribution (left, bottom) by the activity of a population of neurons each associated with an encoding basis function (left, top). **c** Decoded position relative to the location and motion direction of the animal (black arrow) at eight different theta phases (colors as in d) in many theta cycles. **d** Bias of the decoded position (top) and population firing rate (bottom) as a function of theta phase. **e** Decoded SD (mean and SE across theta cycles) as a function of theta phase for the different schemes in the synthetic dataset. **f** Same as e for each experimental session (colors as in j) and mean and SD across animals. **g** Schematic definition of EV-index: the difference between cycle-to-cycle variabiliy (CCV) and trajectory encoding error (TEE). **h-i** The EV-index for the mean (h) and sampling (i) schemes with adding various amount of variability to the simulated data. Variability is added by simulating irregular theta-trajectories (inset in h) and applying various amount of jitter for segmenting the theta cycles (color code, inset in i). Error bars show SE. **j** EV-index calculated for all analysed sessions (color, mean and SE across theta cycles) and across session mean and SD (grey). **k** Cycling index (normalised) for all experimental sessions (colors) versus shuffle control (grey). **l** Real motion trajectories with the position of the animal (empty circles, last point is grey) and the corresponding position decoded from late theta phase (filled circles) in subsequent theta cycles. Theta cycles are colored by the orientation (left or right) of the decoded position wrt. the motion direction of the animal.

2 Learning related adaptation of behaviour and hippocampal neuronal representations to task structure.

The aim of the proposal was to test signatures of uncertainty representation in Ca imaging experiments recorded in collaboration with Judit Makara.

We originally proposed two different experiments. In the first one (Aim 3 of PD and FK), mice initially learn to distinguish two different corridors based on sensory cues. Uncertainty can be induced by applying conflicting sensory cues. In the second experiment (Aim 4 of FK) we were planning to optimize hippocampal imaging for population recordings and implement a sequential, non/spatial decision making task to test the representation of uncertainty during planning in non-spatial contexts. The focus of the research plan was towards experimental design and data analysis, and the actual experiments were performed in the laboratory of Signal Processing with additional support of different funding agencies.

During the past years we have been working towards achieving these goals, but progress with the experiments and data analysis was slower than we expected. In short, we 1) optimized imaging from large population of neurons simultaneously by using transgenic (Transgenic Thy1-GCaMP6s) mice (aim 4.1); 2) Mice could reliably learn to distinguish different virtual reality corridors (Aim 3); We investigated how hippocampal spatial and contextual representations develop during learning; 3) Even after repeated attempts by different experimenters and several sensory patterns, we could not teach the animals to distinguish corridors that could be morphed into each other, so we had to give up Aim 3.1-3.2. 4) We found that developing training protocol for complex behavioral experiments are extremely risky and time consuming, we focused on analysing data from the experiments that were safely working in the lab and did not trained mice in non-spatial decision making tasks.

In the following I will briefly describe our results in how hippocampal spatial and contextual representations develop during learning. These results have been presented in several conferences (FENS 2022, HHMI 2022, Hundoc 2022, MITT 2023), TDK conference of the Semmelweis University (2022 February, Attila Kelemen; 1st prize) and the University of Veterinary Medicine (2022 November). We are continuing these analyses and planning to include the results in a journal publication in the future. The analysis software developed by the support of the grant and is now routinely used by several students and postdocs in our laboratories. It is freely available in my github repository `https://github.com/bbujfalussy/ABmice`.

Establishing the relationship between neuronal and behavioural adaptation to the environment is a fundamental problem in neuroscience. The hippocampus contributes to flexible behaviour by representing task-relevant variables in an abstract form in monkeys performing simple decision making tasks. In rodents, the hippocampus supports spatial navigation by rapidly forming a sparse spatial code. How task-relevant spatial and non-spatial variables develop during learning and whether they are also represented in an abstract form is poorly understood.

It has been recently suggested that neuronal representations that share common elements across several task conditions are extremely useful for generalisation across those conditions because they enable cognitive flexibility and fast learning[16]. In the hippocampus, which is a key region in contextual learning and memory, variables relevant in a context dependent go/no-go task are represented in an abstract form that enables generalisation across conditions in monkeys[16].

Hippocampal representations have been most extensively characterised in rodents performing various spatial tasks[17]. Spatial tuning of hippocampal place cells emerge typically early and independent of the behavioural task[18] whereas task specific components of the hippocampal code emerges more slowly during prolonged experience[19–21].

Spatial codes are characterised by the rapid decorrelation of the representation of nearby positions[22]. Representation of distinct environments can be either independent ("global remapping") or, if the environments are similar, can involve neurons with similar spatial tuning curves that differ only in their firing rates ("rate remapping")[22,23]. Whether neurons can respond coherently during rate remapping and thus enable generalisation of the knowledge acquired at a specific location of one context to other locations even after remapping[24] is currently unknown.

Here we used 2P Ca2+ imaging to record the activity of CA1 pyramidal cells in Thy1-GCaMP6s mice in a contextual go/no-go task that required both goal-oriented navigation and using non-spatial environmental

features. Mice learned to slow down and specifically lick near the hidden reward zone in one of two virtual corridors that differed in their colour or pattern. Early during learning the spatial maps were substantially similar in the two corridors, but they diverged later: the code became more precise around the reward zone and the number of corridor selective neurons increased in both corridors.

Specifically, we studied hippocampal code in low and high performance sessions separately (Figure 2a-d). We found that the fraction of spatially tuned neurons increased from 20% (of active neurons) to 30% in the rewarded corridor in the high sessions (Figure 2e). In the low sessions, a substantial fraction of neurons had similar spatial tuning in the two corridors (Figure 2g) and the representation of the two corridors were highly correlated (Figure 2h). In the high session, corridor selective neurons emerged (Figure 2b,f) and the similarity of the representation of the two corridors decreased (Figure 2g-h).

Next we tested the trial-to-trial representation of the task relevant variables by simultaneously decoding both corridor identity and position. Decoding corridor identity was only possible in the high performance sessions (Figure 3a). Position decoding error was already below chance level in early sessions and decreased further near the reward zone of the rewarded but slightly increased in the unrewarded corridor (Figure 3b-c).

We tested the generalisation of the hippocampal spatial code across corridors and across positions (Figure 3d-e). The position code generalised better in low than in high sessions (Figure 3e), whereas corridor identity was encoded in an abstract form that generalised well across positions in high sessions (Figure 3d). We could also decode the corridor identity from lick and speed in high performance session (Figure 3f) and it also generalised well before the reward zone. The performance of the two decoders using behavioural and neural variables were correlated, but the predictions of the behavioural decoder lagged behind predictions from neuronal activity (Figure 3f-g).

Our results demonstrate that although hippocampal spatial maps may form early upon exploring an environment, emergence of abstract representation of task-relevant variables requires further experience. Encoding task-relevant information in an abstract form in the hippocampus could help the brain to quickly generalise newly acquired knowledge.

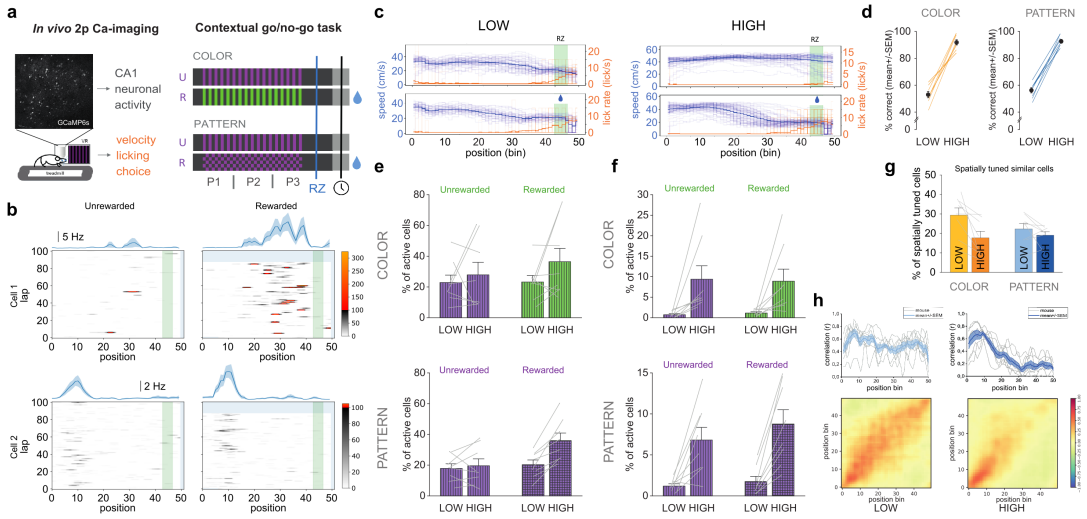

Figure 2. Tuning of hippocampal CA1 pyramidal cells during behavioral adaptation in a virtual contextual go/no-go task. **a** Experimental setup. Mice were head-fixed under the microscope, running on a linear treadmill in a virtual reality environment. Mice were trained to discriminate based on either COLOR or PATTERN between two corridors presented in a pseudorandom order. **b** Example activity (inferred spike rate in individual laps, bottom and average firing rate, top) of two neurons (top and bottom) in the two corridors (left and right). **c** Running speed (blue) and lick rate (orange) in low (left) and high (right) performance sessions in the unrewarded (top) and rewarded (bottom) corridors. Thick line shows average across laps. **d** Percent of correct choices for the analysed sessions. **e** Proportion of spatially tuned cells in the different corridors, tasks and performance. **f** Proportion of corridor selective neurons in the different corridors, tasks and performance. **g** Proportion of neurons with similar spatial selectivity across the two corridors in the different conditions. **h** Similarity of the spatial maps in the two corridors quantified by the cross-correlation (colormap) between the population vector at different positions (x and y coordinate) in the PATTERN task. Top: the diagonal of the cross-correlation matrix of individual animals (grey) and their mean and SD (blue). In panels e,f and g grey lines show individual mice, barplot shows mean and SE.

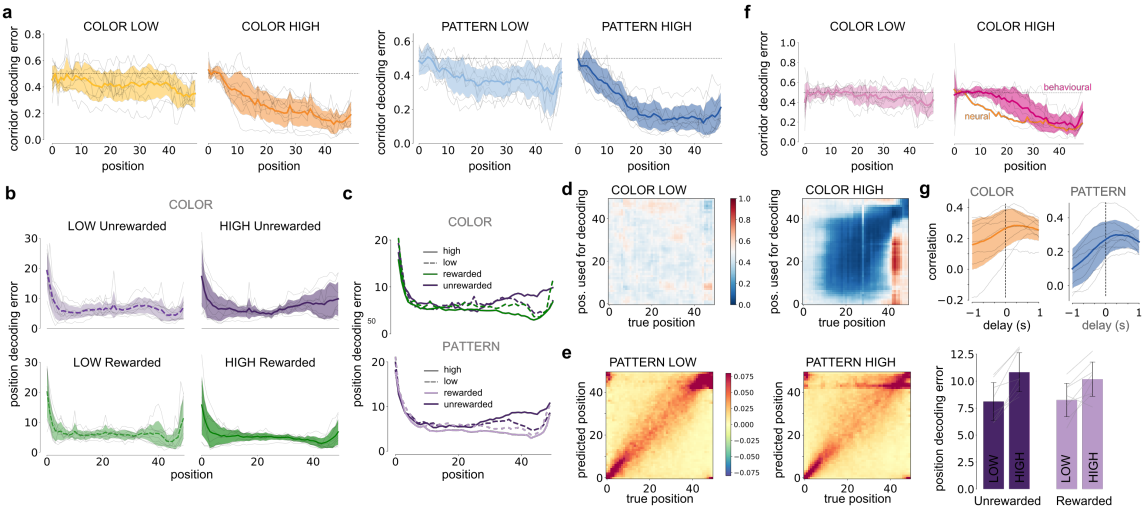

Figure 3. Population code adapts to task structure. **a** Probability of corridor decoding error (color: mean and SD; grey: individual sessions) as a function of position in low and high performance sessions in the color and in the pattern discrimination task. **b** Position decoding error (color: mean and SD; grey: individual sessions) plotted separately for the two corridors in the low and in the high performance sessions. **c** Mean position decoding error in the two tasks. **d** Example cross-position corridor decoding confusion matrix (one low and one high session) shows the probability of error as a function of the true position (x-axis) and the position from where neuronal data was used to train the decoder (y-axis). **e** Cross-corridor position decoding (e.g., decoder trained using laps from the rewarded and tested on the unrewarded corridor) confusion matrices averaged across low (left) and high sessions (middle) in the pattern discrimination task. Average position decoding error using for the two corridors in low and high sessions. **f** Probability of corridor decoding error using speed and licking data in low and high session in the color discrimination task. Mean of the corresponding neuronal decoder is replotted from a. **g** Frame-by-frame cross-correlation between the prediction of neuronal (a) and behavioral (f) decoder. positive delays indicate lead of the neuronal representation.

3 Disentangling hippocampal representations with unsupervised deep-learning

The Aim 5. of the project was to develop new tools to analyse non-spatial sequential data in the hippocampus. We achieved mixed and partial results here and the goal is not fully completed. The results were presented in several conferences (probabilistic AI summer school, Bernstein conference, Hundoc) and have been published in the Master Thesis of Márton Kis (ref. 25). After finishing his MSc in Computer Science, Marton has quit my research group for a job in the industry. We tried to continue this line of research with other students but they need more time to learn the necessary mathematical concepts and programming techniques to proceed with the project. Here I briefly describe our main conclusions.

We applied unsupervised dimensionality reduction techniques to discover the representations used by the hippocampus during cognitive tasks using beta Variational Autoencoders (βVAE). Learning disentangled representations is considered a central problem in representation learning and are indispensable tools to identify variables represented in the neuronal population activity in an unbiased way. The current state-of-the-art unsupervised disentanglement approaches are different flavours of the Variational Autoencoder (VAE) framework, a flexible generative model that aims to learn the joint distribution of observations and their latent generative factors. Several heuristics have been introduced to encourage disentangled representation of the posterior (e.g., Beta-VAE, Annealed-VAE, DIP-VAE) yet how their performance depends on the complexity of the mapping between the latent variables and the observation is not known.

We applied βVAEto a dataset monitoring the activity of hundreds of hippocampal neurons in rats completing an open field navigation task. We found that βVAEcan be used to efficiently learn a lower dimensional representation of neuronal data, but the factors learned by the model are not encoded in a disentangled, interpretable and linearly decodable way.

To investigate the relationship between learning disentangled representations and the nonlinearity of the mapping between the latent and the observed variables we generated several datasets where the latent variables were the 2-dimensional position of a Gaussian blob in a greyscale image and we systematically varied the size of the blob across datasets. This corresponds to simulating the activity of place cells regularly covering a 2-dimensional open arena. Confirming previous results we found, that VAEs can efficiently learn the smooth, disentangled representation of the 2D position (up to arbitrary rotation and scale) when the size of the blob is large (Figure 4a). However, disentangling failed when the size of the blob fell below a critical level. Interestingly, the learned representations displayed hallmarks of disentangling when studied in restricted regions of the state space, yet on the entire representation broke apart into several discontiguous regions (Figure 4b).

We found that this problem only arises when the mapping in highly nonlinear, i.e., learning the right mapping in one region of the state space does not facilitate finding the right mapping in distant areas. By using a supervised pre-training (i.e., training the encoder explicitly to map neuronal activities to position) we could force the network to learn the globally disentangled representation but the two solutions (supervised disentangled and unsupervised entangled) had nearly identical marginal likelihood indicating that this is not just a problem with the numerical optimisation techniques we used (i.e., local minima). We argue that this issue arises when learning the mapping mapping between the latent and the observed variables starts independently at different regions of the state space and current objectives do not encourage stitching these regions together to from a coherent representation. We tried solving this issue by keeping the capacity of the model low for an extended training period until the latent structure became globally disentangled. We speculate that this issue will be exaggerated in real-life tasks where the state spaces are often unbounded and the mappings are highly nonlinear, necessitating further investigation of this issue.

Additional work:

Because our recent results[26,27] as well as new data in the literature[18,28–31] indicated an important contribution of dendritic integration to location coding by hippocampal pyramidal neurons, we decided to further explore the synaptic and dendritic mechanisms responsible for the generation of the place code in the case of hippocampal pyramidal cells during theta sequences. We conducted two different research projects, where the PI (Balazs Ujfalussy) received part of his salary from the FK grant. The first one,

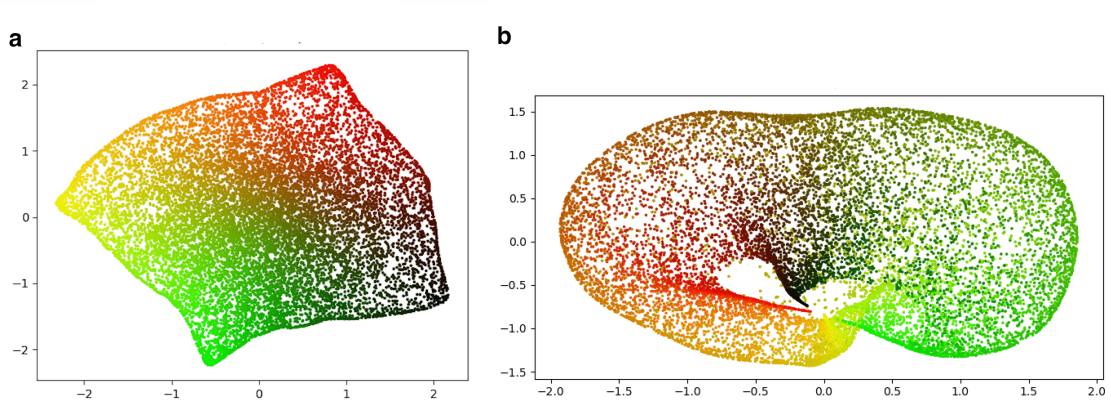

Figure 4. Unsupervised learning of latent variables. **a** Two dimensional latent embedding of 100 dimensional synthetic place cell data. Activity of 100 place cells were simulated at different spatial locations and the generated spike counts was used to train βVAEthat learned a two dimensional representation of the data. Each dot shows the value of the latent variables (x and y coordinate) for a given datapoint colored by the true location (not observed by the network during training, green and red pixel intensity corresponds to true x and y coordinates). In this case the algorithm found the intuitive mapping from neuronal activity to latent variables, and the latents can be interpreted as euclidean spatial coordinates. **b** Similar to panel a, for a different simulation, where the algorithm was unable to find the intuitive mapping from neuronal activity to latent variables, and the latent space remains distorted.

focusing on the impact of functional synapse clusters on neuronal response selectivity during *in vivo*-likeinput conditions is published in ref. 32 and is briefly summarized in Section 3.1. The second project, which was a collaboration with Kim Young Joon and Máté Lengyel (University of Cambridge), has been recently published in ref. 42 and is briefly presented in Section 3.2.

3.1 Impact of functional synapse clusters on neuronal response selectivity

Our goal in this project 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 (sVm) 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 by local dendritic mechanisms, action potential-based, global plasticity rules can not generate functional clustering. Thus, dendritic nonlinearities are indispensable for the appearance of functional synaptic clusters observed *in vivo*[30,33,34].

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

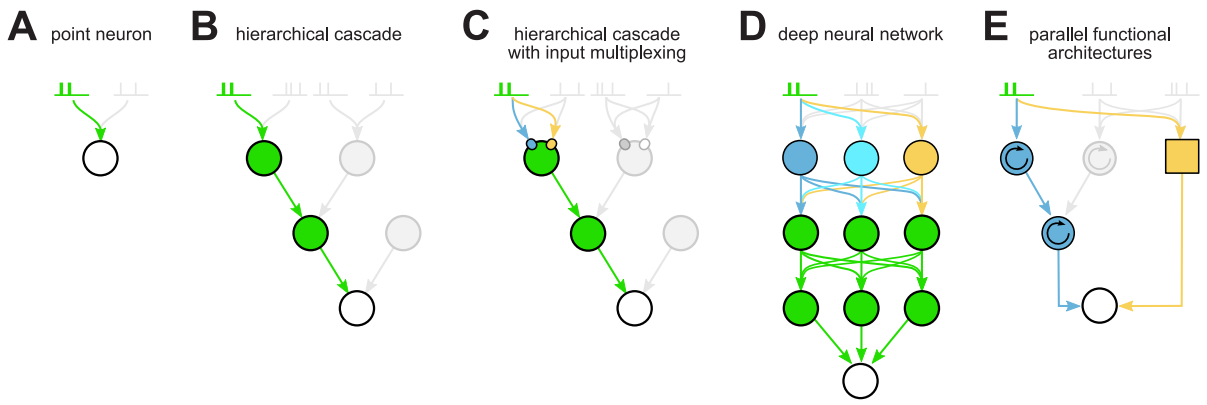

Figure 5. Comparison of computational models of the neuronal input-output transformation Each shape (circle or square) represents a functional "subunit", performing an elementary operation (e.g. summation followed by a sigmoid non-linearity) on the inputs (incoming arrows, from spikes arriving via afferent synapses, or from other subunits) to produce its output (outgoing arrow). The output of the cell is produced by the bottom subunit (open circle). Identical shapes indicate subunits performing similar operations. An example input spike train is highlighted in green, and the pathway processing this input is highlighted in colors. Splitting and mixing of colors indicates divergence and converge of this pathway: parts conveying distinct contributions of this input are shown in separate colors (blue and yellow), parts conveying mixed contributions are shown in green. **A** Point neuron model with no hierarchical processing[35–37]. **B** Hierarchical cascade: each spike train is only processed by a single pathway[27,38,39]. **C** Hierarchical cascade with input multiplexing[27]. Each spike train is processed by multiple rapidly converging pathways that are only kept separate at the level of inputs and use the same type of elementary operation. **D** Deep neural network[40]. Each spike train is processed by multiple rapidly converging pathways that diverge and converge in each step, always using the same type of elementary operation. **E** Parallel architectures (this study). The processing pathways of the same spike train only converge in the last step of the hierarchy, and these pathways belong to different functional architectures using different elementary operations (circles vs. square).

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) (ref. 32).

3.2 Parallel functional architectures within a single dendritic tree

The input-output transformation of individual neurons is a key building block of neural circuit dynamics. While previous abstract models of this transformation vary widely in their complexity, they all describe the underlying functional architecture as unitary, where each particular synaptic input makes a single contribution to the output of the neuron through subunits performing the same single type of elementary computation, either once or cascaded multiple times (Figure 5A-D). Here, we show that the input-output transformation of CA1 pyramidal cells is instead best captured by two distinct functional architectures operating in parallel (Figure 5E).

To study the functional architecture of neuronal input-output transformations under *in vivo*-like conditions, and in particular to establish formally whether its underlying architecture is unitary or parallel, we used a model-based approach that allowed us to rigorously characterize this transformation in a data driven way. For this, we first extended previously developed methods in which subunits only had static nonlinearities[27,38,39] by incorporating subunits capturing arbitrary nonlinear dynamics. Second, we fit all the parameters of the model, including those parameterizing the dynamical nonlinearities, as well as the architecture of the model in a data-driven way, with minimal prior assumptions. Using these methods, we characterized the input-output transformation of a hippocampal CA1 pyramidal cell during the "theta-state"[41].

In particular, our analysis revealed a striking dichotomy in the way dendritic Na^+ channels versus all other dendritic processes contribute to the somatic membrane potential. While the contributions of other processes, summarised in v_{noNa}, are responsible for most of the subthreshold variability in the somatic membrane potential, the contribution of Na^+ channels, represented by v_{Na}, has a major influence on the timing of somatic action potentials.

We found that predicting the contribution of dendritic Na^+channels to the output of the cell (v_{Na}) required a fundamentally distinct functional architecture from predicting the contributions of all other dendritic signals (v_{noNa}). In the v_{noNa} architecture, the entire dendritic tree functions as a *single* global subunit processing all synaptic inputs and expressing a *static* nonlinearity. This architecture is thus defined by the effective weights of synapses, which in turn mostly depend on their somatic distances. In contrast, in the v_{Na} architecture, the dendritic tree is best described by a cascade of *multiple* subunits, each of which corresponds to a small segment of the dendritic tree and processes only the corresponding subset of synaptic inputs. This architecture is mainly determined by how the synapses are assigned to different subunits, which in turn depends on the clustering of coactive synapses on the dendritic tree. The nonlinearity of subunits is *dynamic* and exhibits spiking-like behaviour that is activated by strong and synchronous local inputs.

Thus, synaptic inputs onto the same dendritic tree are processed by parallel functional architectures with distinct computational properties.

This work has been published in Cell Reports 42, 112386 (2023) (ref. 42).

References

[1] Ujfalussy, B.B. & Orbán, G. Sampling motion trajectories during hippocampal theta sequences. *Elife* **11** (2022).

[2] O'Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map* (Oxford University Press, 1978).

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

[4] Wilson, M.A. & McNaughton, B.L. Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058 (1993).

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

[6] Stachenfeld, K.L., Botvinick, M.M. & Gershman, S.J. The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).

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

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

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

[10] Walker, E.Y., Cotton, R.J., Ma, W.J. & Tolias, A.S. A neural basis of probabilistic computation in visual cortex. *Nat. Neurosci.* **23**, 122–129 (2020).

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

[12] 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).

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

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

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

[16] Bernardi, S. *et al.* The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell* **183**, 954–967.e21 (2020).

[17] Moser, E.I., Kropff, E. & Moser, M.B. Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.* **31**, 69–89 (2008).

[18] Sheffield, M.E.J., Adoff, M.D. & Dombeck, D.A. Increased prevalence of calcium transients across the dendritic arbor during place field formation. *Neuron* **96**, 490–504.e5 (2017).

[19] Lever, C., Wills, T., Cacucci, F., Burgess, N. & O'Keefe, J. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* **416**, 90–94 (2002).

[20] Priestley, J.B., Bowler, J.C., Rolotti, S.V., Fusi, S. & Losonczy, A. Signatures of rapid plasticity in hippocampal ca1 representations during novel experiences. *Neuron* **110**, 1978–1992.e6 (2022).

[21] Grienberger, C. & Magee, J.C. Entorhinal cortex directs learning-related changes in ca1 representations. *Nature* **611**, 554–562 (2022).

[22] Fyhn, M., Hafting, T., Treves, A., Moser, M.B. & Moser, E.I. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* **446**, 190–194 (2007).

[23] Leutgeb, S. *et al.* Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619–623 (2005).

[24] Jeffery, K.J., Gilbert, A., Burton, S. & Strudwick, A. Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. *Hippocampus* **13**, 175–89 (2003).

[25] Kis, M. Analysing hippocampal neural representations with probabilistic deep learning. Master's thesis, PPKE Faculty of Information Technology and Bionics, (2020).

[26] Ujfalussy, B.B., Makara, J.K., Branco, T. & Lengyel, M. Dendritic nonlinearities are tuned for efficient spike-based computations in cortical circuits. *Elife* **4** (2015).

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

[28] Sheffield, M.E.J. & Dombeck, D.A. Calcium transient prevalence across the dendritic arbour predicts place field properties. *Nature* **517**, 200–4 (2015).

[29] Bittner, K.C. *et al.* Conjunctive input processing drives feature selectivity in hippocampal ca1 neurons. *Nat. Neurosci.* (2015).

[30] Adoff, M.D. *et al.* The functional organization of excitatory synaptic input to place cells. *Nat Commun* **12**, 3558 (2021).

[31] Rolotti, S.V., Blockus, H., Sparks, F.T., Priestley, J.B. & Losonczy, A. Reorganization of ca1 dendritic dynamics by hippocampal sharp-wave ripples during learning. *Neuron* **110**, 977–991.e4 (2022).

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

[33] Wilson, D.E., Whitney, D.E., Scholl, B. & Fitzpatrick, D. Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. *Nat. Neurosci.* (2016).

[34] Scholl, B., Thomas, C.I., Ryan, M.A., Kamasawa, N. & Fitzpatrick, D. Cortical response selectivity derives from strength in numbers of synapses. *Nature* **590**, 111–114 (2021).

[35] McCulloch, W.S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Bio.* **52**, 99–115 (1943).

[36] Wilson, H.R. & Cowan, J.D. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys J* **12**, 1–24 (1972).

[37] Gerstner, W. & Kistler, W. *Spiking Neuron Models* (Cambridge University Press, 2002).

[38] Poirazi, P., Brannon, T. & Mel, B.W. Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. *Neuron* **37**, 977–87 (2003).

[39] Jadi, M., Behabadi, B., Poleg-Polsky, A., Schiller, J. & Mel, B. An augmented 2-layer model captures nonlinear analog spatial integration effect in pyramidal neuron dendrites. *Proceedings of the IEEE* **102**, 782–798 (2014).

[40] Beniaguev, D., Segev, I. & London, M. Single cortical neurons as deep artificial neural networks. *Neuron* **109**, 2727–2739.e3 (2021).

[41] Buzsáki, G. Theta oscillations in the hippocampus. *Neuron* **33**, 325–40 (2002).

[42] Kim, Y.J., Ujfalussy, B.B. & Lengyel, M. Parallel functional architectures within a single dendritic tree. *Cell Rep* **42**, 112386 (2023).