
Variability of Grid-Cell Activity

Johannes Nagele

Dissertation
der Graduate School for Systemic Neuroscience
der Ludwig-Maximilians-Universität
München

vorgelegt von
Johannes Nagele
aus München

München, den 05.06.2019

Erstgutachter: Prof. Dr. Andreas Herz

Zweitgutachter: Dr. Martin Stemmler

Drittgutachter: Dr. Kevin Allen

Tag der mündlichen Prüfung: 27.09.2019

Contents

Summary	ix
1 Introduction	3
1.1 Spatial Representations in the Brain	3
1.2 Grid Cells in the Entorhinal Cortex	7
1.3 Variability of neural spiking	11
1.3.1 Sources of neural variability	11
1.3.2 Interpretations of neural variability	12
1.3.3 Analysis of spiking variability	13
2 Untethered firing fields and intermittent silences	17
2.1 Abstract	17
2.2 Introduction	19
2.3 Results	20
2.4 Discussion	36
2.5 Materials and Methods	41
2.6 Author contributions	45
2.7 Acknowledgements	45
2.8 Funding information	46
2.9 Declaration of interest	46
2.10 Figures	47
2.11 Tables	55
2.12 Supplementary Figures	56
2.13 Supplementary information	67
2.14 References	72
3 Spike afterpotentials shape the burst activity of grid cells	83
3.1 Abstract	83
3.2 Introduction	84
3.3 Materials and Methods	85
3.4 Results	87
3.4.1 Grid cells differ in the voltage deflections following an action potential	87
3.4.2 Grid cells differ in their spike-train characteristics	88

3.4.3	Post-AP dynamics explain the spike-train characteristics of bursty neurons	89
3.4.4	Spike-train characteristics of bursty cells are conserved across experimental conditions	90
3.4.5	Spatial response properties are shared across all three cell group . .	91
3.4.6	Bursty grid cells: One continuum or two clusters?	92
3.5	Discussion	93
3.6	References	95
3.7	Figures	98
4	Discussion	109
4.1	Complementary work	113
4.2	Time scales and spatial scales of analysis - the impact of binning	117
4.3	Outlook	118
4.4	Conclusion	123
A	Additional information on Fenton's variability measure σ_z^2	125
A.1	The relation of variability σ_z^2 to chi-Square testing	125
A.2	The upper bound on variability and its relation to the zero inflated Poisson parameter	125
B	The Hidden Markov model and its relationship to the ZIP model	129
C	Tuning displacement in 2D recordings	133
D	Jitter transformation and the Beta-Poisson model	139
E	Miscellaneous	143
E.1	Pooled return map of grid cells	143
E.2	mecPhysio Viewer - A graphical user interface for hands-on cell classification	144
E.3	Bin size effect	145
E.4	Grid field detection in 2D: An algorithm inspired by physics	146
	Acknowledgements	159

List of Figures

1.1	Hippocampus in the rat's brain	4
1.2	Hippocampus in the human's brain	5
1.3	Morphology and spatial firing of Grid Cells in the MEC	6
1.4	Simultaneously recorded ensemble of grid cells	8
2.1	Variation in the spike counts on single crossings through firing fields	47
2.2	Example recordings of a grid-cell in the open field and on the linear track .	48
2.3	Zero inflation in 1D, 2D and on different timescales	49
2.4	Trial alignment example from virtual reality	50
2.5	Joint alignment of trials of five simultaneously recorded MEC neurons . . .	51
2.6	Spike trains within a firing field cover shorter distances than expected from firing-rate maps	52
2.7	Coherence of the state transitions	53
2.8	Noise correlations governed by displacement of tuning curves	54
2.9	z-score distributions across data sets, conditions and species	58
2.10	Zero inflation in different speed ranges	59
2.11	Correlation of zero-inflation (ZI) on the linear track versus the open field .	60
2.12	Zero-inflation across parameters and time scales, conditions and species . .	61
2.13	Effects of trial realignment	62
2.14	Drop in zero inflation due to trial alignment	63
2.15	Testing for coherent shifts of grid fields across simultaneously recorded cells	64
2.16	Consistent tuning displacement across contexts	65
2.17	The shifted input LNP model	66
3.1	Spike afterpotentials of MEC cells from mice moving in virtual corridors .	100
3.2	Spike-time autocorrelations of MEC cells from mice moving in VR	101
3.3	Group-level analysis of MEC cells from mice moving in virtual corridors . .	102
3.4	Spike-time autocorrelations of MEC cells from mice moving in open arenas	103
3.5	Comparison of spike-train characteristics and spatial coding across data set	104
3.6	Cluster structure of bursty grid cell	105
3.7	Stability of afterpotential parameters	106
3.8	Robustness of cluster analysis	107

4.1	The number of clusters and the typical distances	122
C.1	Jitter model for tuning curves: The Beta-Poisson distribution	134
C.2	Parameters signalling jittered firing fields in 1D and 2D	136
C.3	Linear track simulation based on 2D parameters	137
E.1	Pooled return map of grid cells	143
E.2	mecPhysio Viewer	144
E.3	The bin size effect	145
E.4	Self-organized grid clustering	146

List of Tables

2.1	Conjunctive tuning of grid cells and zero inflation	55
2.2	Common field displacements for pairs of spatially modulated neurons . . .	56
2.3	Shifts and distances between firing fields	57
C.1	Correlations of different models	135

Summary

Action potentials of grid cells in the entorhinal cortex of navigating rodents occur every two seconds on average. If one considers the precise temporal sequence of these events, however, it can be seen that they rarely occur in isolation. In fact, the intervals between successive action potentials can be on the order of a few milliseconds. Mapped to the trajectory of the animal, a clear clustering of the action potentials in space can be observed as well. The places where the density of such events is particularly high are called firing fields and are arranged in a hexagonal grid.

Regardless of the cell characteristics, the number of spikes observed on different crossings of a field varies strongly. The time between subsequent field crossings is on the order of seconds. We found out that one cause of spike-count variability is that the exact position of the firing fields is not stable over time. In addition, the shifts of the fields were correlated across simultaneously recorded cells. This kind of non-stationarity in the grid-cell network allows conclusions to be drawn about the functioning of this system. Furthermore, dynamic field locations imply that common methods for data analysis of grid-cell recordings can be problematic.

Furthermore, we found out that a subset of grid cells, which have particularly high firing rates when crossing a field, can be associated with a peculiarity in the shape of their action potentials: The spikes of some cells are followed by a short afterdepolarization (DAP). At the same time, we discovered cells with even smaller and extremely stereotypical intervals between their spikes. This group of neurons, however, exhibited less pronounced DAPs. Cells with and without DAP did not differ in their spatial firing behavior. Our results imply that different burst behaviors are not directly related to different types of spatial coding. In addition, we suggest that bursting of grid cells could be altered via the mechanisms of DAP formation.

In summary, this work shows how details of neuronal activity on two different time scales provide fundamental insights into the processes of spatial navigation.

Overview


The results of my PhD research is presented in the form of two manuscripts. The articles are included as single chapters, preceded by a general introduction and review of some previous findings. In what follows, I will provide a brief overview of the organization of the thesis. In the introduction, I recapitulate the neuronal basis of spatial representations in mammals and discuss neural variability. In particular, I review experimental and theoretical results on excess variability in place- and grid-cell spiking. Then I state my research questions.

In the first manuscript, reliability of positional grid and place cell activity in two dimensional environments are compared and evidence for a specific source of trial-to-trial variability in grid cell spiking is presented (Chapter 2). To this end, we demonstrate that an uncertainty of the animal's internally represented location is sufficient to explain the bulk of variance of grid-cell spiking within firing fields. As a consequence of that view on grid-cell spiking, where the main source of noise comes from shifted tuning curves, we challenge the widely accepted hypothesis that grid cells themselves are representing the cognitive map of space. Instead, considering our results, we support the idea of grid cells as being part of a distance estimation complex, encoding relative distances rather than the actual position. In the appendix, I derive a doubly stochastic spiking model based on our observations and show that it outperforms previously suggested models in explaining spike count distributions.

The study considered looked at neural spiking on the behavioral timescale but left out another source of variance, particularly the effect of spiking motifs on small timescales of a couple of milliseconds, i.e., burst activity. In the second manuscript (chapter 3), we therefore look at virtual reality in-vivo recordings of the membrane potential of grid cells in mice and compare the variety in spiking to tetrode recordings in freely running mice in the open field. We identify clusters of cells which differ substantially in their electrophysiological parameters and spiking patterns. A mechanistic explanation in terms of depolarizing afterpotentials (DAPs) is presented. DAPs are short positive voltage deflections that follow action potentials. These results are discussed jointly in the last chapter and some future research questions are stated. Finally, in the appendix I provide some details on individual aspects and additional projects that I was involved during my doctoral studies.

*Under carefully controlled
experimental circumstances,
an animal will behave
as it damned well pleases.*

THE HARVARD LAW OF ANIMAL BEHAVIOR



Chapter 1

Introduction

At children's birthday parties in Germany, a game called "Topfschlagen" is often played. Your eyes are blindfolded, your friends turn you around several times to disorient you, and now, armed with a wooden cooking spoon, you have to find a pot turned up-side down (and the sweets covered beneath it). With a little luck, banging the spoon against objects in your vicinity, you will eventually find the pot, which you then joyfully make ring out with a percussion fanfare. To the delight of the observers one typically manages to completely lose one's way- often with several unexpected bumps into pieces of furniture, or even getting stuck under a table. Upon removing the blindfold, one often finds oneself in a completely unexpected location. The surprise comes about because we always maintain a notion in our heads of where we are. How is it that we always have a notion in our heads of where we are even when it is completely off like many times during of such a game? Or should we rather speak of a hallucinated position of the self in such cases, because the imagined place has little to do with the actual place?

In everyday life, we are not blindfolded and can use all our senses to locate ourselves. In that situation, is the imagined position objective, or is it still just an imaginary position?

The answer to this question fills philosophy books and has much to do with proprioception and the question of the self. We won't address such questions in this work. Rather, we will investigate the question of whether and how subjective location is reflected in the physiology and behavior of neurons known to be involved in spatial information processing.

1.1 Spatial Representations in the Brain

Over the last five decades a large body of anatomical, physiological, and theoretical research accumulated evidence for the idea that the hippocampus and adjacent brain areas in the mammalian temporal lobe are crucial for memory and navigation related tasks [1]. Together with the dentate gyrus (DG) and the subiculum (Sub) the hippocampus proper forms an anatomical complex, the hippocampal formation. While some authors consider the entorhinal cortex to be a part of the hippocampal formation, in this paper we stick to a limited definition in order to distinguish the two areas more clearly. The hippocampus

proper itself is subdivided into three main areas: The cornu ammonis (CA1, CA2 and CA3). The hippocampal formation of mammals as illustrated in Fig. 1.1 has a long history in evolution [2].

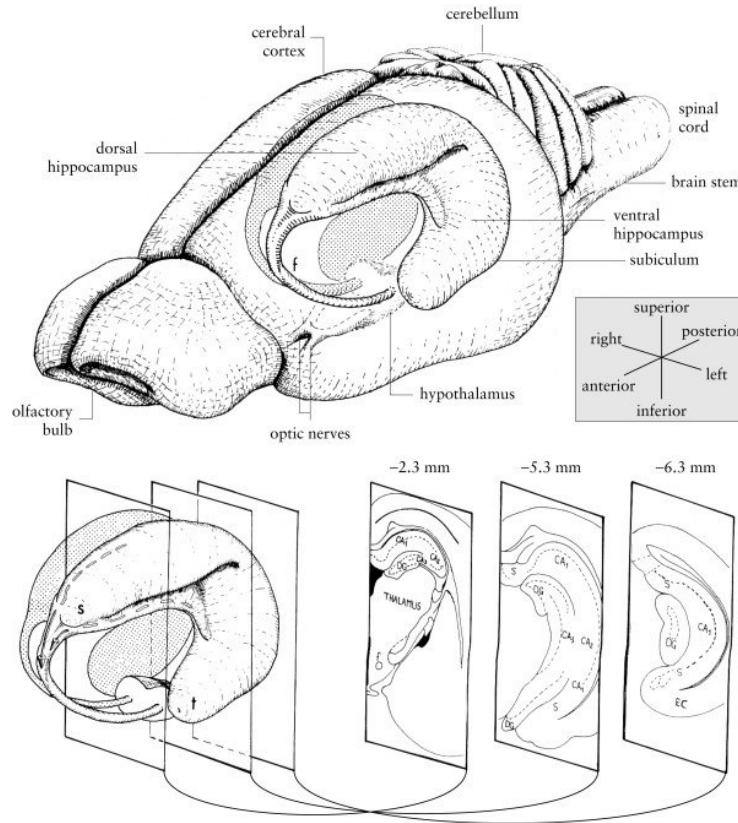


Figure 1.1: Anatomical structure of the rat hippocampus. Drawings of the rat brain showing the three-dimensional organization of the hippocampus and neighboring structures. Three coronal sections through the left hippocampus are shown at the bottom right of the figure, with their approximate anteroposterior coordinate relative to bregma. CA1, CA2, CA3: cornu ammonis fields 1–3; DG: dentate gyrus; EC: entorhinal cortex; f: fornix; s: septal pole of the hippocampus; S: subiculum; t: temporal pole of the hippocampus. Taken from Cheung et al. 2005 [3]. Available via license: CC BY 2.0, <https://creativecommons.org/licenses/by/2.0/>

In the 1970s the surprising discovery of a certain group of principal neurons, the “*place cells*”, in areas CA1 and CA3 of the rat by O’Keefe and Dostrovsky [4] led to fundamental insights into how the mammalian brain represents spatial memories: The neural activity of these cells depends on the location of the animal independent of the task that the animal is performing. Basically, each of these cells is mainly inactive except when the animal is located in one or several widespread, specific and spatially confined regions called *place fields* or *firing fields*, which can be found in any part of the environment. On average these cells elicit an action potential only about two times per second over a whole recording session [5] but within fields the spike rate is higher [6]. From the activity of many of these neurons the position of the animal can be decoded in any moment. The precision

depends on the number of neurons decoded simultaneously and is about 8cm if a population with 33 neurons is analyzed [7]. The smallest place fields cover an area of 10cm². The activity of place cells in open field environments is independent of the direction in which the animal traverses the field, therefore exhibit *omnidirectionality*. However, this can change when obstacles are placed into the environment or the animal is put on a linear track [8]. Place cells can be driven by both, visual and other sensory cues as well as by locomotion [9]. Over the time it turned out that the functional variety of place cells indeed is very complex: Besides the simple place cells that only correlate with location there exist object and experience related place cells as well as “misplace cells” which signal novelty or the unexpected lack of an object [10]. The discovery of “social place cells” [11, 12, 13] was an exciting twist to the story. The existence of place cells provided support for the concept of “cognitive maps” as suggested by Edward C. Tolman already in 1948 [14]. Further evidence for that idea came with the discovery of similar coding principles for head direction [15].




Figure 1.2: Left: Modified drawing of the neural circuitry of the rodent hippocampus as drawn by Santiago Ramón y Cajal (1911). DG: Dentate Gyrus. Sub: Subiculum. EC: Entorhinal Cortex. CA1-CA3: Hippocampus proper (CA2 unlabelled). Right: 3D model of the human Hippocampus. The hippocampal region is highlighted in red. Image generated by Life Science Database (LSDB, <http://lifesciencedb.jp>). Available via license: CC-BY-SA-2.1-jp, <https://creativecommons.org/licenses/by-sa/2.1/jp/deed.en>

The focus of this work is on another group of principal cells which have the interesting property that they have multiple firing fields arranged in a striking hexagonal lattice. These cells are called “*grid cells*” and have been found in rats, mice, bats, monkeys and human [16, 17, 18, 19, 20, 21]. These neurons are located in the pre- and parasubiculum and in layers 2, 3 and 5 of a small region adjacent to the hippocampus, the medial part of the entorhinal cortex (EC) (see Fig. 1.2) [22]. The superficial layers of the EC (i.e. Layers 2 and 3) provide inputs to the hippocampus whereas the deep layers receive feedback from the hippocampus as sketched in the inset of Fig. 1.2. The lateral (LEC) and the medial entorhinal cortex (MEC) process substantially different informations: In LEC mainly object related processing is found whereas the MEC neurons tend to reflect spatial

and proprioceptive aspects [23]. In layer 2 of the MEC grid firing was associated to two principal cell types, stellate and pyramidal cells [24, 25], see Fig. 1.3.

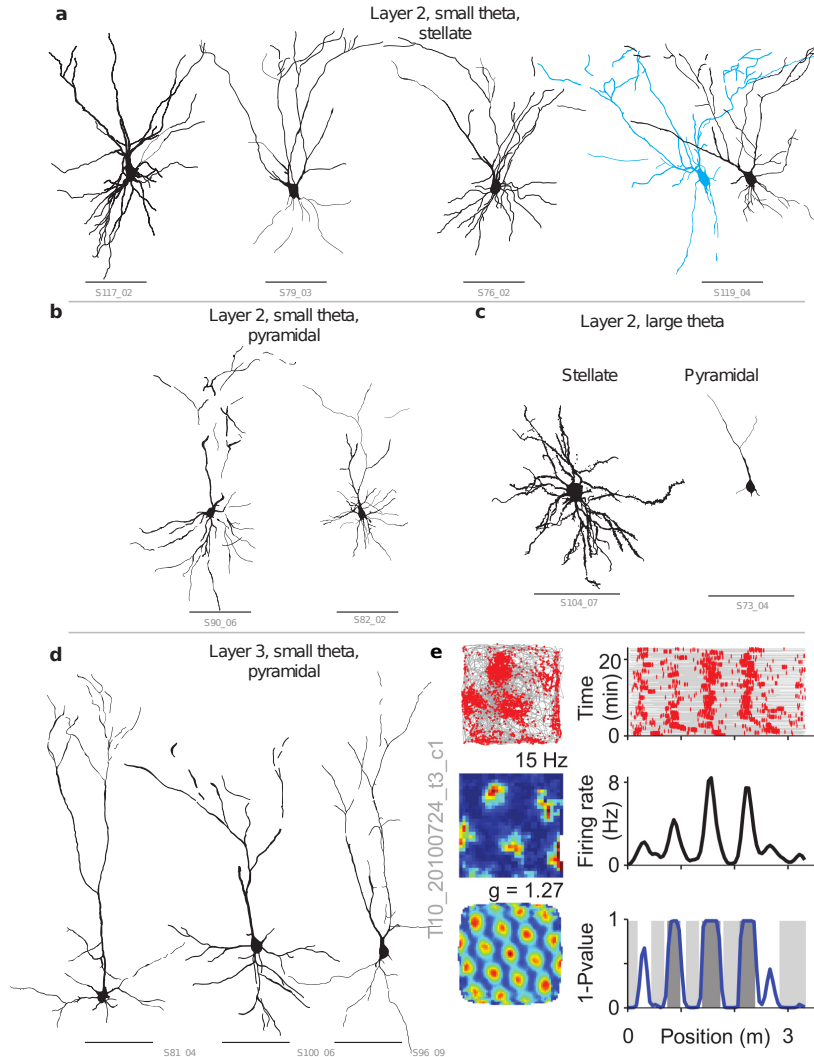


Figure 1.3: Morphology of grid cells in the MEC and grid-cell spiking on the virtual linear track. **a**, **b**, **c** and **d**: Morphology of stellate and pyramidal grid cells (except blue). Details of the imaging and identification can be found in [26]. Shown are examples of neurons suspect to a large theta fluctuations in the membrane potential (**c**) and examples with small theta fluctuations (**a**, **b** and **d**) in different layers of the MEC. Codes in grey are the IDs of the cells for all panels. Scale bars are $100\mu\text{m}$. **e**: Upper left: Grid-cell spiking in a $50\text{cm} \times 50\text{cm}$ box (trajectory (grey) and spikes (red)). Upper right: Same but showing multiple runs on a virtual linear track. Middle row: Tuning curve estimates in 2D (left: Positional firing rate colorcoded, 15Hz peak rate) and 1D (right: Trial averaged and smoothed tuning curve). Bottom left: Spatial autocorrelogram of the 2D tuning estimate with grid score (g). Bottom right: Detected firing fields (dark grey), regions between fields (light grey) and field probability (blue) based on a shuffling procedure. Reprinted by permission from Springer Nature: Nature, “Membrane potential dynamics of grid cells”, Domnisoru et. al, Copyright 2013 [26].

1.2 Grid Cells in the Entorhinal Cortex

In the absence of visual and other external cues spatial memories can still be triggered by proprioception. The network that is thought to perform the transformation of proprioceptive informations to a mental map on which spatial memories can be pinned and on which trajectories can be planned is the grid-cell network in the EC. It was suggested to be the coordinate system or the inner GPS of the brain (as discussed in [27]) such that it provides a neural metric for space [28]. That system might provide the neural substrate for dead reckoning via path integration [29].

As mentioned earlier, grid cells are characterized by the fact that the firing fields of each grid cell form a periodic grid. In two dimensional arenas a hexagonal arrangement of the firing fields can be observed when marking the positions at which a grid cell fired an action potential. These patterns were often referred to as honeycomb patterns [30], hexagonal [17], triangular [31] or rhombic lattices [32]. Characteristic of this particular structure is that each field is surrounded by six neighbouring fields and that all fields have the same distance to their neighbours (the grid period or spacing), at least in theory. This pattern occurs relatively frequently in nature, since objects with comparable repulsive properties arrange themselves hexagonally under pressure (e.g. during certain self-organization processes), whereby the densest sphere packing is achieved. In addition, it should be noted that even with uniformly distributed random points in 2D, the expected value of the number of nearest neighbors is six [33] and consequently the mean angle between two edges of a tessellation of such points as in the hexagonal lattice is $\pi/3$ [34].

Because grid cells unfold the specific structure of their fields in two-dimensional arenas, their properties are also usually studied in such environments. Typically, the animal is put into a box of about half a square meter, but experimenters often vary the geometry of the environment. To ensure that the animal covers as much space in the arena as possible typically food pellets are thrown in from time to time which encourages the animal to forage. Recording duration is in the order of tens of minutes, at the very least. To acquire recordings as the animal performs standardized movements, one lets the animals run on a linear track. In such 1D corridors, the animal can run back and forth, but has limited freedom to move side-wards. In virtual reality (VR) setups, animals are highly constrained: the physical location of the rodent is fixed as it travels on a rotating ball or treadmill; its movements are translated into the optic flow of scenes projected onto screens in the animal's field of view. By changing the gain factor between movement and optic flow, VR setups provide the possibility to disentangle locomotion cues and visual cues. On the other hand, tactile feedback and other modalities are missing in the VR [35]. Furthermore, head fixed preparations for intra-cellular recording of grid cells are possible in the VR [26]. The cells are clustered in modules [36] and the scale of the grids increases from dorsal (30 to 40cm) to ventral ($> 100cm$) [37]. The scale of a grid, or grid period, is typically measured by detecting the dominant peaks in the spatial autocorrelogram of a grid pattern. In Fig. 1.4 the firing fields of simultaneously recorded grid cells that fall into different modules are shown and sorted by their grid period. The maximal number of modules that was observed during one recording was four, but due to sampling issues and finite scales of the




Figure 1.4: Firing rate maps of 25 simultaneously recorded grid cells from the Stensola et al. data [36]. The firing rates are computed as number of spikes in a bin (5cm width) divided by the dwell time in that bin and subsequent Gaussian kernel smoothing ($\sigma = 5\text{cm}$). Each panel shows the firing rate map of one grid cell computed from a 20min recording of a male long evans rat together with the peak firing rate corresponding the the brightest bin. The cells are arranged according to their grid period (defined as the location of the first peak in the radial profile of a grid cell's autocorrelogram). All cells shown were labelled as grid cells by Stensola et al. [36].

experimental environments there are limitations in detecting very small and very large grid spacing. Estimates for the total number of modules range from four to hundreds [38]. The ratio between adjacent grid scales is about 1.4 [36] and is close to values being predicted by models for optimal coding of the animal's location using grid cells [39]. Besides the lattice spacing, within modules the cells' firing rate maps share the same orientation of the grids [36]. The distribution of grid phases is compatible with a uniform coverage [40]. In 1D environments, for example, when the animal is running on a linear track as described above, the spatial firing patterns of grid cells are consistent with slices through 2D lattices [40, 41].

Grid cells are defined by their functional properties. To this end, the *grid score* [42] has been used as an easy-to-use marker. It is computed as follows: From the spatial firing rates of a grid cell a two-dimensional spatial autocorrelation, the *autocorrelogram*, is computed (see Fig. 1.3, panel e, bottom left). After removing the central peak and all but the six closest fields to the center the correlogram is rotated by 60° and 120° (set A) and by 30° , 90° and 150° (set B). For each rotation a correlation coefficient with its original version is computed. Due to the 60° symmetry of a perfect hexagon the coefficients are expected to be large for rotations in set A and small for rotations in set B. Finally, the grid score is defined as the minimum difference between correlation coefficients of set A rotations and set B rotations. Values larger than 0 indicate a hexagonal symmetry. This measure can be biased due to limited amount of data. To account for this shuffling procedures are used [43]. There are also extensions to local measures of hexagonal symmetry [44]. Independently of the symmetry of the firing fields arrangement their spatial information content (SI) can be estimated to detect irregular and distorted grid patterns as well [45]. Briefly, the SI is an estimate of the mutual information of spiking and the animal's position in the arena.

The hexagonal arrangement of a grid cell's firing fields triggered numerous computational models of grid formation: Initially it was suggested that grid cells result from the superposition of plane waves [46], oscillatory interference [47, 48] or from continuous attractor dynamics [49, 50]. These models predict place-cell activity to result from a superposition of many grid cells. However, place cells are observed even when grid cells activity is reduced or inactivated [51, 52, 53, 54, 55]. This motivated a growing body of research on models of grid-cell formation based on place cell inputs [56, 57, 58, 59] as well as on self-organized formation of grid-like units in (semi-) supervised neural networks performing spatial tasks [60, 61]. It was demonstrated that a grid-cell network can be used for positional coding [39] and goal-oriented navigation [62] as well as for estimating distances [63, 28]. In the last years more general roles of grid cells beyond the classical idea of a cognitive map for spatial representations were suggested [64, 65]. Support for these ideas came from experiments indicating grid like activity during navigation in visual, auditory, odor and even abstract spaces in the entorhinal cortex and parts of the prefrontal cortex [66, 67, 68, 69, 70, 71]. We will discuss some implications of these ideas in the final discussion section of this dissertation.

Most of the computational models for grid formation mentioned above also address some of the less obvious dynamical properties of grid cells. For example, not all grid cells respond only to the animal's location. Many grid cells show conjunctive tuning to

additional stimulus dimensions: Grid cells were found that are modulated by head direction [42], by speed [72], contextual cues [73], goal locations [74, 75], or combinations of multiple cues [76]. Other factors alter the activity of grid cells as well: It was shown that grid cells realign their pattern consistently when place cells remap in novel situations [77]. In such situations also an expansion of the grid lattice was observed [78]. Grid cells remain inactive when the animal is not running.

Further factors impact grid-cell spiking: Short *bursts* of action potentials followed by longer inter spike intervals are observed, see Fig. E.1 in the appendix. Bursts are high firing-rate events that are characterized by a series of two or more spikes within a short time window, are observed across brain regions and come in various fashions. Besides that the precise characteristics of such events are very diverse [79], generally bursting neurons are grouped into intrinsically or input-driven bursting. Bursts have been shown to have many different purposes like synchronization [80], band-pass filtering allowing for multiplexing [80, 81] and robustness of information transmission [82]. In the subiculum intrinsically bursting neurons are thought to be utilized to relay signals to other brain areas via strengthening the signal of individual promising stimuli into long lasting burst sequences, such that new stimuli get emphasized and postsynaptic processing circuits are activated [83]. Typically such packages of spikes are fired by grid cells predominantly at specific phases of an oscillatory modulation of the membrane potential at a frequency around 7Hz to 10Hz (theta oscillation). The frequency and amplitude of that oscillation vary over time and are correlated with speed [84]. The phase at which the first spike of such a burst is elicited precesses with respect to that oscillation [85, 86]. Some cells are stronger modulated by theta oscillations than others [26], some elicit bursts frequently and others only sparsely [87]. Furthermore grid cells were shown to replay previously experienced trajectories on a fast timescale [88, 89]. Recent studies show that firing rates differ across fields of the same grid cell different to what is expected from a regular lattice where each node is just a copy of another [73, 90, 91, 92].

Whenever we talk about “grid-cell spiking” we are referring to the firing of action potentials (APs). However, APs themselves come in different shapes: The APs of some grid cells in the MEC are followed by a depolarizing afterpotential (DAP), a small depolarization hillock in the membrane potential maximum around 4ms after the AP [93]. The amplitude of DAPs is in the order of 10mV but varies from cell to cell and is less pronounced in pyramidal cells than in stellate cells [94, 95]. The origin and the functional role of a DAP is not clear but the depolarization might reflect a window of opportunity for a subsequent spike right after and therefore being involved in producing bursts.

All these findings and many more experimental observations increase the required complexity for models of grid-cell spiking. Progress has been made in incorporating some of the described dynamics into the computational models, but none is able to capture all of these dynamical properties of grid cells yet. Decoding models of grid cells typically rely on idealized grid-cell activity patterns to be perfect hexagons in space, having equal firing rates in individual firing fields and ignore temporal noise-correlation structure within the system.

The variety of effects on the activity of grid cells shows how difficult it is to model

grid-cell spiking already at a purely phenomenological level. In other words, even simply mimicking a grid cell's response in a simulation is problematic. Thus, hypothesis testing based on simulated data of spatially modulated cell firing in the entorhinal cortex needs to be done with care. In practice, typically shuffling procedures are performed, for example for cell classification as mentioned earlier or to perform statistical tests. The simplicity of such approaches makes them favorable but can result in high false-positive rates in the detection of grid cells [43].

The above observations indicate that grid cells do not always fire the same spike pattern when the animal crosses one of its firing fields (even if the field is traversed identically). This in turn leads to a large variability in the spatial firing of grid cells compared to what is expected from simple statistical models. This variability is not only an artifact and a mirror of our ignorance, but it is also very useful to measure our success in further developing spiking grid-cell models. As generally in the sciences, if we want to understand a phenomenon, we try to explain its peculiarities and thereby reduce the errors in our predictions. In this dissertation, we want to follow this path and contribute to understanding the function of grid cells and spatial navigation in mammals.

1.3 Variability of neural spiking

Under the assumption that a brain performs computations, it is natural to expect that neurons always behave identically when presented the same conditions and inputs. In this way, calculations could be performed reliably. As a result, the nervous system would be assumed to always make the same decisions leading to identical behavior under identical circumstances. However, the neural activities of behaving animals show a large trial-to-trial variability, even in laboratory experiments. This variability raises fundamental questions about the nature of neural encoding of information and decision-making. The causes of variability have been the subject of ongoing debate in brain research since the first half of the last century, when physiologists found that the response of nerve cells can be highly variable across repeated and well controlled stimulation. In the last two decades it became evident that even without any stimulation or behavior, spontaneous activity of cortical neurons can be subject to large fluctuations. For example, work by Britten and colleagues demonstrated that behavioral choices and neural activities correlate even when the stimulus was held constant [96]. In that situation the neural firing is highly variable with respect to one's expectation given the stimulus.

In the following sections we review potential origins of such variability, discuss some general aspects of its interpretation and introduce statistical measures for its investigation before we continue with the first manuscript.

1.3.1 Sources of neural variability

One potential origin of variable neural discharge patterns are thermal fluctuations at the level of ion channels or unknown internal states of the animal which can not be controlled

in behavioral experiments. This can be problematic under a reductionist approach in which changes in neural response are tried to be explained by external stimuli. The lack of information about unknown factors that influence the processes in the brain lets the changes of a dependent variable (like the neural activity) appear to be random noise. In their review paper Renart and Machens [97] formulate this issue as follows: “*Since neural activity is influenced by many variables, identifying these hidden variables and locating their sources is one of the hard problems of systems neuro science*”. Determinants of neural firing can be more or less controllable and are produced external to or internally within the observed system (the organism). Typical internal variables are factors like attention, arousal, fatigue and motivation, which can be controlled to some extent [97]. Other factors, like chaotic dynamics within the networks, cannot be controlled without direct intervention in the network itself. The typical external variables are designed by the experimenter as part of the stimulus design. Even so, the degree to which a variable is external or internal is not always well defined. For example, thermodynamic fluctuations and prior experiences are neither exclusively external nor internal. It is clear that many of these factors themselves might not be mutually independent; chaotic behavior might be triggered by state changes or changes in the temperature.

1.3.2 Interpretations of neural variability

When the brain is studied in vitro or under anesthesia, many unknown influences are eliminated. Yet this situation is highly artificial. Indeed, one would like to understand how neural computations are performed under natural conditions. In such situations, attempts to decode neural activity can fail dramatically due to unreliable and variable responses. Even when a decoder is trained to achieve optimal accuracy on the training data and some particular test set(s), the following Gedankenexperiment illustrates the difficulty in identifying hidden variables: Suppose that we trained a decoder to predict an animal’s actions in the near future based on the neural firing patterns while the animal was behaving. Then we use this decoder to predict the actions from spontaneous activity in the absence of behavior and stimuli. The decoder will still predict outcomes from its known set of actions. Two very different results are plausible: First, the decoder predicts with a *low degree of certainty* a sequence of actions that appears random. This would be the case for spontaneous random activity which was not captured during training. Second, the decoder predicts a sequence of actions with *high certainty* close to the original test set. This latter scenario might ensue if the animal generated stimuli internally or hallucinated a sequence of actions from the training data during the recording. For both of the cases the *performance* of the decoder was low because the animal did not act and therefore the predictions did not match to the animal’s actual behavior. But unlike in the first scenario, the interpretation of the neural variability was not random noise in the second scenario: From the high certainty in the decoding it can be concluded that there was a hidden variable which determined the chronological sequence of the coded actions.

So far we were mostly concerned about neural variability from the perspective of the ob-

server, basically highlighting the lack of knowledge and difficulties in controlling variables. We, as the observers, controlled the animal's actions and made sure there is no external stimulus in this thought experiment. We designed the decoder, defined the training set and evaluated the variable responses.

But what does variability actually mean for the organism itself? From the thought experiment in the previous paragraph it becomes clear that certainty and performance of a decoder are not identical. Although the latter statement seems trivial, it is worth to emphasize that highly certain predictions which are far off from the true values do not necessarily imply that the decoder is erroneous. Biological organs are thought to be adapted via evolution [98, 99]. In the brain, which has the capability to learn and adjust, there is a high chance for that some decoding network is structured in alignment precisely to what its inputs are encoding in healthy subjects. Typically scientists only have vague ideas about what is encoded in the brain area or network of interest. What we identify to be random noise based on such hypothesis might carry important information for a postsynaptic network. Furthermore, also random activity as such can have a functional meaning: On the one hand, noise can be beneficial for coding [100], learning and plasticity [101] and is a necessary condition for the idea of “Neural Darwinism” [102]. On the other hand high variability is related to psychiatric disorders, where it is reported to be increased excessively [103]. However, the high degree of randomness of neural activity reported previously got questioned in the last years [97, 104, 105].

It has been long known that measured randomness is often due to the necessary ignorance with which we approach new insights:

Randomness is only a measure of our “ignorance of the different causes involved in the production of events”.

LAPLACE, 1825

Applied to neuroscience this quote highlights the importance of both accurate experimental design (control of stimuli) as well as careful interpretations of neural variability. A third component is to define the “event” itself such that it matches to what a study is intended to draw conclusions on. For example, neurons can exhibit random fluctuations on small timescales but act reliable on larger scales and vice versa or encode different modalities on different timescales. Therefore, when defining the number of spikes in timewindows of a certain width to be the quantity of interest (event), only statements about variability on that time scale can be drawn. Related to that issue is the general problem of quantifying variability itself as described in the next sections.

1.3.3 Analysis of spiking variability

In neuroscience the Fano factor (FF) and the coefficient of variation (CV) have been used to describe variability in spike trains. The FF is defined as

$$FF = \frac{\sigma_W^2}{\mu_W} \quad (1.1)$$

where σ_W^2 is the variance and μ_W the mean of spike counts in a time window W across trials. Whereas the FF is used to measure trial-by-trial variability the CV is used to quantify the variability in the inter spike interval (ISI) distribution:

$$CV = \frac{\sigma_{ISI}}{\mu_{ISI}} \quad (1.2)$$

with σ_{ISI} being the standard deviation of the ISI distribution and μ_{ISI} its mean. In the case of a homogeneous Poisson process both measures are equal to one in the limit of large N (large number of trials for the FF and large number of events or spikes for the CV). In the case of the CV being equal to zero there is no noise in the time series and all ISIs are equal. Similarly all spike counts are equal across trials in the case of FF being equal to zero. The term *underdispersion* describes situations in which the values are smaller than one whereas excess variability, that is values significantly larger than one is termed *overdispersion*.

CV and FF are specifically useful tools for analyzing inter spike intervals or spike counts collected from homogeneous renewal processes. In the presence of serial interval correlations the situation changes: Negative serial interval correlations (i.e. due to spike frequency adaptation, see Benda and Herz 2003 [106]) affect the FF in the following way: $\lim_{T \rightarrow \infty} FF = CV_\infty^2 (1 + 2\Sigma)$ with $\Sigma = \sum_{i=1}^{\infty} \Sigma_i$ (see McFadden 1962 [107]). Here Σ_i is the i th-order linear correlation coefficient (the expected correlation between those intervals (ISI_k, ISI_{k+i}) that are separated by $i - 1$ intermediate ISIs). Both measures are sensitive to trial-to-trial variations in the firing rate. More robust to this sort of fluctuations is the trial averaged $\overline{CV} = 1/N \sum_i^{trials} CV_i$ as described by Nawrot 2010 [108].

In case of rate-modulated spike trains (inhom. firing rates) where the modulation is consistent across trials the FF is advantageous to the standard CV. That is because in such situations the CV captures the variations in the firing rate. One solution for that is a demodulation of the spike times using a model of the firing rate over time [109]. Other non-parametric solutions reduce the effect of slow rate variations via pooling ISIs locally and measuring the CV on small time scales from these ensembles [110].

Parametric strategies for estimating neural variability do not rely on FF, CV or corrections of those. Such approaches include Bayesian approaches of estimating rates and regularity simultaneously from data using an inhomogeneous-Poisson assumption with constraints on the variations in firing rate (like smoothness / serial correlations, as for example in Fenton et al. 2010 [111] and Rad and Paninski 2011 [112]) or by estimating parameters for a gamma process [113]¹

In essence, model free variants typically suffer from strongly biased or uncertain estimates

¹The regularization of a the rate model or the smoothness that can be tolerated to achieve Poisson spiking itself is a measure of variability: Without letting nearby data influence the model at one given point (which is the role of binning, smoothing and regularization) but treating each data point independently when modelling expectations, the most informative estimate is the data itself. On the other hand, considering any data to influence the model at any point equalizes the estimate. That way we are left with a homogeneous model even when rate modulations are presumed. Hence, without any grouping or smoothing, we arrive at an extreme case of *overfitting* and as the model is exactly the data there is no noise left, whereas when smoothing extremely we end up in scenario of *underfitting* and might overestimate the variability dramatically.

whereas the estimation of a firing rate model or tuning curve is at the heart of the parametric approaches which has to be done with care as well.

In the first manuscript in this dissertation (Chapter 2) we use a strategy to measure the variability of grid cells that belongs to the first group, similar to what has been applied to place-cell recordings previously [111, 114]. Then we study the nature of that variations more in detail. Furthermore, we provide a model of grid cell trial-to-trial variability. In fact, in that chapter a neuron's behavior is compared in the very same stimulus region but at very different moments in time. However, in the second manuscript (chapter 3) we turn towards the other side of the spectrum: In a model free study we analyse the detailed temporal patterns of grid-cell spiking in the range of milliseconds, independent of the stimulus at that moment, and find three different physiological classes of grid cells that differ in their temporal spiking characteristics.

The general issue that comes with trying to attribute noise to firing-rate variations or spike-timing variations is discussed in Amarasingham et al. 2015 [115]. In the discussion of this dissertation we will refer to these issues and provide an outlook on how the two perspectives on neural variability can lead to different insights into neural computation.

A good approach to measure neural variability is one that generalizes across scales:

*From milliseconds to seconds, from one to another context,
and from single action potentials to the symphony of interconnected networks.*

That is as difficult as understanding the brain.

Untethered firing fields and intermittent silences: Why grid-cell discharge is so variable

Johannes Nagele^{a,b,c}, Andreas V.M. Herz^{b,c}, Martin B. Stemmler^{b,c}

^a*Graduate School of Systemic Neurosciences, Munich*

^b*Bernstein Center for Computational Neuroscience Munich*

^c*Ludwig-Maximilians-Universität München*

Abstract

Grid cells in medial entorhinal cortex are notoriously variable in their responses, despite the striking hexagonal arrangement of their spatial firing fields. Indeed, when the animal moves through a firing field, grid cells often fire much more vigorously than predicted or do not fire at all. The source of this trial-to-trial variability is not completely understood. By analyzing grid-cell spike trains from mice running in open arenas and on linear tracks, we characterize the phenomenon of “missed” firing fields using the statistical theory of zero inflation. We find that one major cause of grid-cell variability lies in the spatial representation itself: firing fields are not as strongly anchored to spatial location as the averaged grid suggests. In addition, grid fields from different cells drift together from trial to trial, regardless of whether the environment is real or virtual, or whether the animal moves in light or darkness. Spatial realignment across trials sharpens the grid representation, yielding firing fields that are more pronounced and significantly narrower. These findings indicate that ensembles of grid cells encode relative position more reliably than absolute position.

Keywords:

Grid cell | Entorhinal Cortex | Spatial navigation | Trial-to-trial variability |
Overdispersion | Excess variability | Firing-field jitter and drift | Tuning
curve | Zero inflation | Noise correlations

Introduction

Path recall and path finding are crucial skills, yet the neurons that represent space in the mammalian brain are surprisingly variable in their discharge during spatial navigation, as first shown by Fenton and Muller (1998) for place cells. While grid cells in the medial entorhinal cortex (MEC), for instance, tend to fire at specific locations in space that map out a hexagonal grid (Fig.1A), close inspection of individual trajectories through a particular firing field (Fig.1B) reveals that a grid cell will often fire *no* spikes when the animal passes through the field; in other instances, the cell will fire many more spikes than expected from Poisson statistics.

A grid cell’s hexagonal spatial firing-rate map by itself does not capture these highly variable spike count statistics. Hippocampal place cells, which have only a few isolated firing fields, also exhibit strongly fluctuating firing (Fenton and Muller, 1998). A number of explanations for place-cell spike-count variability have been proposed: neuronal sensitivity to task, action, or sensory variables that are unrelated to spatial location, changes in selective spatial attention (Fenton et al., 2010), flickering between multiple maps of space (Jezek et al., 2011), or “knowledge-guided fluctuations” (Jackson and Redish, 2007; Prerau et al., 2014; Kelemen and Fenton, 2016). Another potential source of variability in neurons with spatial selectivity is error accumulation during path integration (Hardcastle et al., 2015).

To study the origin of grid-cell variability, we analyze data from experiments in rats and mice moving in two-dimensional arenas and in mice on real and virtual linear tracks. We show that two distinct modes are necessary to

explain the trial-to-trial variability of grid-cell activity. In one mode, cells fire stochastically when the animal is in a firing field; in the other mode, cells remain absolutely silent. In addition, grid cells' firing fields, over time, drift in space, as seen across repeated experimental trials on the linear track. These findings provide a mechanistic explanation for the excess variability in grid-cell activity.

Results

Grid cells are silent much more often than predicted by the firing-rate map

For spatially modulated neurons, such as place cells and grid cells, firing-rate maps allow one to calculate the expected number of spikes along any trajectory through the animal's environment. The true number of spikes fired, however, will vary from run to run (Figs. 1A and 1B), even if the animal takes the same path (Fig. 2). To better understand this trial-to-trial variability, we analyzed multiple data sets (Fyhn et al., 2004; Sargolini et al., 2006; Fyhn et al., 2007; Stensola et al., 2012; Domnisoru et al., 2013; Pérez-Escobar et al., 2016). By z-scoring the spike counts, we found that the observation by Fenton et al. (2010) on excess variability in hippocampal CA1 spike trains extends to recordings in medial entorhinal cortex (MEC) of rats and mice, as shown in Fig. S1. The average z-score across all grid cells in rat was $\sigma_z^2 = 6.62 \pm .27$ ($n = 199$), whereas mouse grid cells had a $\sigma_z^2 = 7.63 \pm .95$ ($n = 41$). The difference of the average variability across species was not significant (Two-tailed Welch's-test for equal population means: $w = -1.02$, $p = .31$).

To assess the spike-count statistics of MEC grid cells in mice at a more detailed level, we segmented mouse trajectories such that the expected number of spikes in each segment was the same (cf. Methods). To reach a fixed number of spikes, these trajectory segments covered time intervals of varying lengths. Such a segmentation requires no assumption that the firing fields have a grid-like structure, or even that distinct firing fields exist. We then constructed histograms of the spike counts on the trajectory segments.

If the spike-count statistics are determined solely by a time-varying firing rate, then the resulting distribution should be Poisson, for which the spike count variance is equal to the mean spike count (Fig. 1C).

For grid cells recorded in two-dimensional arenas, only 5% ($n = 138$) had spike-count distributions that were consistent with Poisson spiking (χ^2 test: $df = k - 2$, where k is the number of categories with expected counts larger than 5, $p < 0.05$). Instead, these distributions were frequently bimodal and strongly skewed, as shown in Figs. 1D and E. In fact, many times the spike count was zero; cells remained completely inactive much more frequently than expected from the Poisson null hypothesis.

Bimodal spike-count distributions can be described by mixture models, in which several distributions are combined. A particularly simple mixture, which is widely used outside of neuroscience (Lambert, 1992; Greene, 1994), invokes the concept of zero-inflation (ZI). In short, ZI draws a Bernoulli variable for every sample with probability α to decide whether the spike count s is set to zero or whether the count will be drawn from some standard spike-count distribution with expected value $\mu = \langle s \rangle \geq 0$. In this case, observing a spike count of zero can have two causes: the stochastic state

corresponds to the “zero regime” or the standard distribution produced a zero count. Under zero-inflation, a grid cell might not fire at all even though the animal is at the center of a grid field for this cell. On the other hand, between firing fields, the firing rate will be close to zero, so the spike count will likely be zero, regardless of whether zero-inflation is present.

Intermittent firing of grid cells also occurs when the animal runs on a linear track. Fig. 2 shows the same grid cell recorded in two-dimensional and one-dimensional environments. Passages through a firing field in which, contrary to expectations, the cell failed to fire are highlighted by dashed black lines (Figs. 2C and G). In both environments, the bin for zero spikes stands out in the spike-count histogram (Figs. 2D and H).

The degree of zero inflation in the grid-cell spike count can be measured by an index ZI_{idx} , as introduced by Puig and Valero (2006):

$$ZI_{idx} = 1 + \ln(p_0)/\mu \quad (1)$$

where μ is the empirical mean spike count and p_0 is the frequency of observing zero counts.

This heuristic measure supposes that the null hypothesis for the spike-count statistics is Poisson. Given that the Poisson probability of observing zero spike counts is $\exp(-\mu)$, the logarithm in the equation above is $\ln(p_0) = -\mu$, and hence $ZI_{idx} = 0$ for Poisson neurons. . If there are more zeros, then ZI_{idx} becomes positive.

131 of 136 grid cells recorded by Pérez-Escobar et al. in the open field had a positive ZI_{idx} (mean= 0.52 ± 0.01 , which deviates significantly from zero based on a t -test: $t = 44.66$, $p < 10^{-10}$, $N = 136$).

The ZI_{idx} is a sensitive, model-independent, empirical measure for the presence of zero inflation, which is calibrated against the Poisson null hypothesis. It makes no assumptions about the statistics of spiking (other than that the frequency p_0 of zero spikes should not be zero).

Given the finding of zero-inflation, we specifically tested whether zero-inflated Poisson (ZIP) models describe the spike count data better than Poisson models. Therefore, we fitted the Poisson and the ZIP model to the spike-count distributions and compared these models using a likelihood-ratio test.

The likelihood-ratio test penalizes ZIP for having one more degree of freedom ($df = 1$) in the comparison of likelihoods: this additional parameter is the probability α of being in the zero regime (see supplemental methods S1)). Unlike ZI_{idx} , the probability α is estimated using the maximum likelihood method and is specific to the underlying model for the spike count statistics. The ZIP model had a higher likelihood than the Poisson model for 52 of 65 pure grid cells (80%, $p < 0.001$, $df = 1$) in the data set, and for 67 of the 76 grid cells with conjunctive speed, border or head direction tuning (88%, $p < 0.001$, $df = 1$). The ZIP model is also a better match for 69% of non-grid, but still spatially modulated neurons ($p < 0.001$, $df = 1$). To confirm the significance of these results, we numerically simulated inhomogeneous Poisson processes based on the animals' trajectories through the firing rate maps. For these simulated data, we asked whether ZIP could accidentally have a higher likelihood than Poisson. ZIP, though, almost never had a higher likelihood—Poisson was preferred with ($p < 0.001$, $df = 1$).

In rodents, the spike trains of spatially modulated cells in the entorhinal

cortex and hippocampus are coupled to a 6-12 Hz theta rhythm. When the animal stops moving, the theta rhythm ceases, and cells reduce their rate of firing. To test whether pauses in the animal’s movement could explain pauses in grid cell spiking, we divided the data into segments corresponding to different speeds of the animal. Significant zero inflation was found in all speed ranges (Suppl. Fig. S2A).

Zero inflation on linear tracks is predicted by grid-cell behavior in open fields

In the experiments of Pérez-Escobar et al. (2016), mice were put on an 80 cm long linear track after the grid cells had first been recorded in a 70 cm² square box. The linear track experiments permit corroboration of zero inflation as a phenomenon in a context with less variation in the animal’s speeds and one less degree of freedom for movement. We, therefore, estimated the cells’ probability of being in the “zero-firing” state on the linear track and compared these estimates to ones derived from the open-field experiments. Indeed, cells that had a high probability α of being in the “zero-firing” state in the open field tended to have a high value of α on the linear track, too, independently of the running direction and light/darkness context of the track experiments (see Fig. 3A).

The value of α on the linear track was also correlated across different running directions and light contexts (Fig. 3B, $r = 0.68 \pm 0.03$, $p < 0.005$). Figure S3 details the correlations in α on a context-by-context basis.

The zero-inflation probability α did not differ between runs towards the right and the left end of the track (median = 0.13 ± 0.01 vs. 0.14 ± 0.01 , Wilcoxon signed rank test, two tailed: $U = 2625$, $p = 0.09$, $n = 115$). Furthermore, whether the mouse ran in the dark or in the light had no effect

on the strength of the zero inflation (Kruskal-Wallis test, $p > 0.1$).

These correlations exist across a range of time scales. On short time scales, motifs such as bursts or theta-rhythm modulated spiking might explain some part of zero inflation. While estimating the zero-inflation probability from longer time windows yields lower values for α , zero inflation is still significant on the time scale of seconds (Fig. S4A). Indeed, the number of pauses on the linear track, measured on short time scales, is predicted by a cell's behavior in the open field on much longer time scales (Fig. S4B).

The average firing rates did not correlate with the probability α of excess zeros (average Pearson $r = -0.05 \pm 0.01$, $p > 0.6$ for all light conditions and running directions on the linear track), even though the firing rates themselves were correlated across light conditions (average Pearson $r = 0.94 \pm 0.01$, $p < 10^{-10}$).

To check whether a grid cell's propensity to fire bursts of spikes correlates with the amount of zero inflation, we divided grid cells into bursty and non-bursty neurons as described by Latuske et al. (2015). The excess-zero probability α in the open field is similar for bursty and non-bursty grid cells ($\langle \alpha \rangle_{\text{non-bursty}} = 0.09 \pm 0.01$, $n = 40$) ($\langle \alpha \rangle_{\text{bursty}} = 0.10 \pm 0.01$, $n = 97$). The difference was not significant (Kruskal-Wallis for equal medians: $K = 0.88$, $p = 0.35$; t-test for equal means: $t = 0.80$, $p = 0.43$ and Mann-Whitney $U = 1742$, $p = 0.175$). Taken together, neither the firing rate nor the burst behavior has a discernible relationship to a grid cell's zero-inflation property.

Conjunctive grid cells are tuned not only to the location of the animal, but are modulated by additional factors, such as speed or head direction. Such conjunctive tuning to multiple signal dimensions will cause neurons to fire

more rarely if the different input streams interact in a multiplicative manner. We therefore tested whether cells that were previously labelled as conjunctive and were recorded both in the open field and on the linear track ($n = 115$) tended to have higher rates of zero-inflation. The results are summarized in Table 1. Indeed, conjunctive neurons in the open field had higher scores for zero-inflation.

Two sources of trial-to-trial variability: Missed and shifted firing fields

Domnisoru et al. observed events in which a grid cell fired no spikes during a passage through a grid field, which were called missed fields (see Suppl. Fig. 22 in Domnisoru et al. (2013)). Under Poisson statistics, such events should be rare. Examining intracellular recordings reveals trial-to-trial variations in the average membrane potential relative to the spiking threshold (Fig. 4A and C)

Moreover, previous reports indicate that the positions of firing fields are not fixed (Barry et al., 2007; Hardcastle et al., 2015; Keinath et al., 2018). Figure 6(A,D, top row, red) displays a grid cell’s spikes on single passages through two firing fields, which were recorded on the real linear track (Pérez-Escobar et al., 2016). Averaging across trials yields the usual estimate of the grid cell’s firing rate as a function of position (middle row, blue). For comparison, we use the average firing rate to draw Poisson surrogate spike trains (bottom row, black). On each passage through a firing field, we compute the distance covered by the mouse from the first to the last spike within that field (shown as solid lines connecting the spikes in the raster plot of Fig. 6A,C. At the single-trial level, the measured spike trains within a firing field tended to cover a shorter distance than the surrogate trains, even

though the (trial-averaged) firing-rate map was identical (Fig. 6B,D).

The finding that spike trains within a firing field cover shorter distances than expected could be the result of the fields shifting position from trial to trial. Figure 4A,B shows the firing fields for a mouse grid cell recorded intracellularly (Domnisoru et al., 2013). Both the intracellular voltage and the firing rate are highly variable from trial to trial. We now align the trials to undo the effect of positional jitter of the firing fields (Fig. 4C,D). Alignment sharpens the profile of the average voltage ramps (Fig. 4E) and the firing fields (Fig. 4F). Alignment also permits the recovery of more spatial information than from the time-averaged firing-rate maps (Fig. S5).

Narrower fields with higher firing rates, reduced variability, and increased information in the spike trains are also found after realignment of extracellularly recorded grid cells (Pérez-Escobar et al., 2016), as shown in Fig. S5.

Variability: Field positions versus field-to-field distances

As the firing fields shift, we wanted to test whether multiple fields in a single cell maintain their relative distances, given that the grid-cell population is thought to support distance estimation in addition to encoding the actual physical location of an animal (Huhn et al., 2009).

We quantified the variability in the relative distances by selecting cells that had at least two well-defined and well-separated firing fields (cf. Methods). For these cells, we computed the spatial autocorrelations of the firing rates on each trial. The distance between fields is reflected in a peak in the autocorrelation. We defined the grid cell's spatial period on a trial-by-trial basis as the position of the first maximum in the autocorrelation. For the spatial period to be well-defined, the height of the local maximum in the

Pearson correlation (which is normalized to lie between -1 and +1) had to lie at least 0.5 above the surrounding minima and occur at a spatial lag larger than 4.8 cm (equivalent to two spatial bins). Given that some fields might be missed on some trials, we furthermore insisted that an autocorrelation peak matching these criteria be present on at least 25% of the trials.

There is a simple prediction for the variance σ_{ptp}^2 in the spatial period across trials. Provided that the peak-to-peak distances are large compared to the magnitude of typical shifts, when each firing-field shifts its position independently across trials and then $\sigma_{\text{ptp}}^2 = 2\sigma_{\text{shift}}^2$. Figure 4 plots σ_{ptp} against σ_{shift} .

On average, in $67 \pm 4\%$ of the cells, the variance σ_{ptp}^2 was smaller than expected from the null hypothesis of random jitter (see Table 3). One example of a simultaneous recording is shown in Fig. 5. For these neurons (marked by orange points in Fig. 4), σ_{shift} was more than two times as large as σ_{ptp} . Thus the position of the fields was less consistent than the first peak in the spatial auto-correlation.

A cell with a high value of σ_{ptp} in one light or running condition tended to also have a high value of σ_{ptp} under other conditions (average correlation of σ_{ptp} across conditions was $\langle r \rangle = 0.53 \pm 0.03$ (sem), all correlations were significant (at level 0.01) after Bonferroni correction for $N = 15$ tests.

Independence versus coherence of firing-field shifts across grid cells

Next, we sought to distinguish two possible scenarios that could give rise to drifts in the firing fields: noise at the single-cell level or a population-level drift in the internal representation of the animal's location. In the second scenario, fields shift across the whole population of spatially tuned neurons

in a coherent fashion, so that the fluctuations in activity will be shared across the population as correlated noise. For this purpose, we studied simultaneous recordings from spatially modulated cells in the data from Pérez-Escobar et al. (2016).

We tested whether a field shift in one cell is mirrored by an identical field shift in other cells (see Fig. S7). First, the rate maps for each cell were normalized by the respective average peak firing rate, in order to eliminate trends in the firing rate over time. We then aligned the *trial-averaged* rate maps so that the average spatial phase of the grid-pattern was the same across cells. This procedure preserves relative drifts of the firing fields, if such are present. In the last step, we compared the Pearson correlations of the firing rate maps across cells for simultaneous trials to the correlation for randomly shuffled trials. If the firing fields jointly drift, then the Pearson correlation will be higher for the original than for shuffled trials. The results were deemed to be statistically significant for p -values below $0.05/(n_{\text{bins}}n_{\text{cells}})$ (Bonferroni correction for multiple testing).

We ran this analysis for runs towards the right and left end of the track under the three different lighting conditions used by Pérez-Escobar et al. (2016), which they had labeled l1 for light condition 1, l2 for light condition 2 (with a different pattern of lighted stripes on the apparatus) and d for darkness. The correlation analysis revealed a joint component of field position drift across all six contexts (Fig. S7 and Fig. S8A). In 41 out of 44 recordings of multiple cells, coherent drifting was detected in at least two contexts. In total, 3238 cell pairs were considered. Pairs with significant drift in all six contexts (5%) showed up 10^6 times more often than expected

for independent realizations of the null hypothesis of independent drift.

For grid-cell pairs and conjunctive-grid-cell pairs, higher ratios of consistent field displacement were observed than for non-grid spatial-cell pairs, see Table 2 and Fig. S8B for a comparison of the fractions of significant common shifts in pure-grid-cell pairs and other spatially-modulated cell pairs. In some recordings, neither cell in a pair exhibited significant firing-field drift, which meant that the correlation analysis had to remain inconclusive.

Drifts in one running direction were not correlated with drifts of the same grid cell in the successive runs in the opposite direction. Such a lack of correlation between the two run directions is consistent with the hypothesis that field positions in the two run directions are not only different (they occur at different spatial locations, as observed experimentally), but also functionally independent (Pröll et al., 2018).

Figure 5 shows a striking example of five simultaneously recorded cells that exhibit joint firing-field shifts across trials. We were able to undo these displacements by applying *identical* spatial shifts on each trial to the five firing rate histograms. The population-wide alignment sharpened the firing-field profiles, as was true for the alignment of trials on a single-cell basis (Fig. 4). Once again, but now on a population-wide basis, alignment permits the recovery of more spatial information from the MEC population of neurons. Moreover, the distance between the firing fields of two cells within the same trial was less variable than the positions of the fields across trials.

No evidence for error accumulation in grid cells recorded on linear tracks

Hardcastle et al. (2015) report that grid fields drift in open fields as path-integration errors accumulate. The amount of drift increases with the dis-

tance to the boundary and the time since the last boundary contact. Whenever the animal encounters a wall, these authors argue that the drift is reset to zero.

To test whether grid fields on the linear track are also subject to cumulative drift, we analyzed the jitter in the position of individual firing fields, treating runs towards the left and right end of the track separately. If drift accumulates until it is corrected by encountering a boundary, then the jitter should be greatest for the fields farthest from the most recently visited end. We used data from 28 grid fields that were reliably detected in both left and right runs and under all light conditions on the track (l1, l2 and d). Fields on the two end platforms were not considered. The boundary-driven error-correction hypothesis predicts that the size of the jitter grows with distance ran on the track. Therefore, we computed a linear regression between field position and jitter width for both running directions. Jitter width and field locations were measured as described in the Methods. No significant correlations were found between the magnitude of the jitter and the distance from the most recent boundary reached by the animal, however. Error accumulation would also predict that jitter causes the firing fields to become wider in the right half of the track for rightward runs, and in the left half for leftward runs. We therefore computed the difference in the jitter in left- and rightward runs for firing fields present in both directions. Under all conditions, the null hypothesis (identical jitter for both running directions) could not be rejected, based on a one-sided Mann-Whitney-U-test. For the three light contexts, this test yielded: (l1, fields in the left half (L), Mann-Whitney-U: $U = 178$, $p = 0.94$; l1, right half of track (R): $U = 239$, $p = 0.429$, l2-L:

$U = 74$, $p = 0.96$, l2-R: $U = 143$, $p = 0.51$, d-L: $U = 90$, $p = 0.99$, d-R: $U = 119$, $p = 0.16$). The median jitter-width was around 3.5 cm in each condition, compared to an average field-width of 6.7 cm. Again, we found no indication of error accumulation on the linear track. Not even in the dark, when the mice would most likely path-integrate, did there seem to be significant error accumulation; it cannot be ruled out, though, that other sensory cues were present on the apparatus that calibrated the field positions. It is also conceivable that error accumulation is too small over the time-span of a few seconds to be measurable in these experiments.

Trial alignment reduces variability

We hypothesized that alignment would make the spike-count statistics more Poisson-like and reduce the amount of zero inflation towards a small positive level determined by “missed-field” events. We fit a linear-nonlinear-Poisson (LNP) model to the spike trains of each cell, based on the animal’s position as the covariate, before and after alignment across trials (for details see Fig. S9). For every cell, we then compared the two fits using a likelihood ratio test. We penalized the aligned fits for having more free parameters; each trial’s shift was treated as an additional parameter.

Based on a likelihood-ratio test, in 368 out of 475 cells, the alignment improved the match of the spike-train statistics to an LNP model ($p < 0.001$). Surprisingly, not all of the cells that were better fit by an LNP model after alignment met the criteria of Pérez-Escobar et al. (2016) for spatially modulated cells: only 256 of the 368 cells were spatially modulated, classified either as border cells ($n=23$), grid cells ($n=116$) or other spatial cells ($n=117$). Only 33 spatially modulated cells did not significantly improve

their fit to an LNP model under alignment.

We estimated the shifts from the cross correlations across trials, not from a maximum-likelihood procedure to optimize the likelihood of the LNP model. Indeed, in some cases, alignment reduced the likelihood of the LNP model (n=96).

We then fit zero-inflated variants of the LNP model as described in Giles (2010). After alignment, the maximum-likelihood estimate of the zero-inflation probability α dropped in value by 54 to 70 percent (Wilcoxon tests, $p < 0.001$ for all six settings, N=61). Despite the drop, α did not vanish after alignment in most cases. More precisely, in only 8% of the cases did α drop to a value smaller than 0.01 after alignment, see Figure S6: not all unexpected firing pauses can be explained by firing fields shifting along the track.

While the shifts are not correlated across running directions, the overall amount of shift σ_{shift} of individual grid cells was highly correlated across light contexts (l1-l2, left: Pearson correlation coefficient $r = 0.86$, right: $r = 0.74$; l1-d, left: 0.72, right: 0.71; l2-d, left: 0.71, right: 0.77, all $p \ll 10^{-5}$) and less correlated across running directions (left-right: d: $r = 0.55$; l1: $r = 0.48$; l2: $r = 0.56$, all $p < 0.01$).

Coherence of the “zero-firing” state across cells

Firing field displacements contribute to the observed phenomenon of zero-inflation, as was seen by the drop in the zero-inflation probability α after alignment. Nevertheless, some degree of zero-inflation is preserved. For the data set of simultaneously recorded cells shown in Fig. 4, we asked whether the probability of being in the “zero-firing” state is correlated across these

cells, not just before alignment, but *after* alignment.

In Fig. 7A, we compute the probability of a cell being in the “zero-firing” state, given a ZIP model for the spike-count statistics. The Bayesian estimate of the “zero-firing” probability in each 200 ms long time window requires two parameters: (1) the cell’s firing rate in that time window, averaged over aligned trials; (2) the zero-state probability α for that particular cell, as estimated from the entire experiment. This probability is particularly high when, contrary to expectations, a cell does not fire even when the firing rate predicts it should. This probability is not correlated with the time-dependent firing rate itself (Fig. 7B).

To show that the “zero-firing” state is more strongly correlated across cells than would be predicted by chance, we drew independent ZIP surrogate spike trains for each of the five cells. Figure 7C marks the cases with asterisks in which the true correlations are significantly stronger than expected by chance.

Trial-to-trial displacement of firing fields induces noise correlations

We found that the spatial tuning of MEC cells changed from trial to trial, but did so jointly across cells. When firing fields shift together, cells with overlapping fields will exhibit correlated changes in their firing rates. In short, coherent shifts result in noise correlations. We therefore asked whether firing-field shifts could explain the noise correlations across cells at a quantitative level. For this purpose, we employed statistical models of cell ensembles. We created five artificial cells with multiple firing fields on a simulated linear track to mimic the experiments of Pérez-Escobar et al. (2016). For each cell, the fields represented a random slice through a two-dimensional

hexagonal grid (Yoon et al., 2013; Pröll et al., 2018). The firing fields were modeled by von Mises tuning curves with a concentration parameter $\kappa = 2.1$, as suggested in Herz et al. (2017) and a maximal firing rate of 12.5 Hz. Both the spatial phase (in the range from -30 cm to +30 cm) and the period of the lattice (10, 14, or 19.6 cm) were chosen randomly for every cell. We drew 100 such ensembles of 5 cells. Fields drifted from trial to trial by anywhere from 0 to 0.4 cm. Finally, we generated Poisson spikes based on the displaced tuning. Each ensemble was simulated 5 times for $N_{\text{trials}} = 100$ trials.

Noise correlations were computed as

$$r_{ij} = \frac{1}{N_{\text{trials}} N_{\text{bins}}} \sum_{x,k}^{N_{\text{bins}}, N_{\text{trials}}} \frac{(f_{ixk} - \overline{f_{ix}}) \cdot (f_{jxk} - \overline{f_{jx}})}{\sigma_{f_i} \sigma_{f_j}} \quad (2)$$

where f_{ixk} denotes the firing rate of a cell i in the spatial bin x on trial k and $\overline{f_{ix}}$ is the trial-averaged firing rate of cell i in bin x . Figure 8B plots the noise correlation against the spatial offset of the firing fields (as measured by the peak position of the spatial cross-correlations). In the model with coherent shifts of firing fields, the measured noise correlations were highest for zero offset and decay to either side as the spatial phase difference increased. In contrast, in the absence of shifts, the same graph was flat (Fig. 8A). To compare the model to data, we calculated the noise correlations for 903 co-recorded cell pairs across 138 grid cells measured on the linear track (Fig. 8C). The tuning offset between the grid fields in a pair was measured in 2.4 cm bins, while the number of trials that were averaged to estimate the noise correlations ranged from 13 to 110, with a median number of 64. The measured noise correlations were qualitatively similar to the model's noise-correlations.

A similar dependence of the noise correlations on the grid-phase offset has been reported for grid cells in rats (Mathis et al., 2013; Dunn et al., 2015; Tocker et al., 2015).

Discussion

In this study, a detailed analysis of the trial-to-trial variability in the responses of mouse grid cells revealed new insight into their coding properties: roughly half of the grid-cell excess variability can be explained as a result of the grid-cell population’s shifting representation of space. Despite such shifts, the peak-to-peak separation of firing fields is largely preserved. In many instances, we were able to realign repeated trials and recover spatial information in the firing-rate map, which revealed firing fields that had been partially “buried in the noise”.

Various authors have sought to explain the origins of neural variability; for a review, see Renart and Machens (2014). For instance, a conjunctive (mixed) tuning of neurons to additional sensory, task, or state variables unrelated to spatial location will increase neuronal variability. Hardcastle et al. (2017) find that some grid cells are jointly tuned to combinations of place, head-direction, and speed. In our analysis, conjunctive location- and head-direction tuned cells were, in fact, more variable than “pure” grid cells. Such conjunctive tuning can coexist with persistent variations in a cell’s grid field properties from field to field, which would allow grid cells to convey additional spatial information; in particular, the peak firing rate varies across fields (Diehl et al., 2017; Dunn et al., 2017; Ismakov et al., 2017).

Yet conjunctive tuning only provides a partial explanation for the excess

variability exhibited by grid cells. Other phenomena, such as variations in the sub-threshold voltage or firing threshold, intrinsic bursting behavior, or other non-stationarities must play a role. We focused on two factors: changes in the firing state and shifting field positions from trial to trial.

The failure of a grid cell to discharge in a firing field is corroborated by calcium imaging studies of entorhinal cortex (Low et al., 2014; Heys et al., 2014; Gu et al., 2018). Interspersed amongst trials with vigorous calcium dye responses, one will generally find traces with no response. Such behavior is consistent with a fluctuating spike threshold, which led us to model grid-cell discharge using a particularly simple version of a Hidden Markov Model (Gat et al., 1997): a zero-inflated Poisson process. Cells thus had two states: a silent state and an active state. In the intracellular data on grid cells (Domnisoru et al., 2013) that we reanalyzed, the silent state was associated with lower sub-threshold membrane potentials. The transitions between the silent and active state as the mice ran on the linear track could simply reflect natural fluctuations in the membrane potential. In addition, a silent state might correspond to the firing field’s position transiently shifting away from the track. As we observed field shifts parallel to the track’s direction from trial to trial, it is reasonable to suppose that fields also drift in the transverse direction. Were grid cells to integrate time since trial onset, not spatial distance, one would also measure apparent field drifts (Kraus et al., 2015; Eichenbaum et al., 2016; Tsao et al., 2018). As the animal’s speed is fairly constant on the linear track, grid-cell firing will show the same regularities in time as in space. However, if grid cells encode time, trial-to-trial speed variations would lead to a rescaling of the grid fields measured in

space. All fields would move either closer together or further apart. Moreover, fields should become larger or smaller, depending on speed. We found no evidence for such effects in the data-sets we analyzed. De Almeida et al. (2012) observe that grid-cell spiking falls into two modes: either the cell fires on an inbound trajectory and not when leaving the field, or vice versa. In the interpretation of these authors, grid cells switch between prediction and retrospection. Moreover, simultaneously recorded grid cells are likely to operate in the same mode—either prediction or retrospection. Correlated field displacement reflecting an uncertainty or error in the absolute position estimate could lead to similar effects, though: in one shift direction, the fields would be active behind the animal’s current position, whereas in the other direction, fields would be active ahead of it.

For the majority of mouse grid-cell recordings in darkness, the firing fields, measured in terms of absolute positions, are no longer discernible in the averaged firing-rate map (Pérez-Escobar et al., 2016; Chen et al., 2016). Nevertheless, when one measures the spatial autocorrelation over short time intervals, a dominant length scale emerges, which deviates only slightly from the one observed in light. This observation is consistent with the hypothesis that the firing fields continually shift, but in a manner that maintains the relative distance between firing fields. Our analysis shows that not only grid cells, but also other spatially modulated cells exhibit such shifts. Such coherent firing-field drifts are compatible with continuous-attractor models of grid-cell firing (Burak and Fiete, 2009). These models envision a hexagonal bump-attractor state in a recurrent neural network; the state moves coherently across the network in response to a velocity signal. Any source of

noise, barring some form of recalibration of the firing-rate map, will lead to coordinated drifts in the measured firing fields.

But the question remains why firing fields on the linear track shift not only in darkness, but also when lights are on and visual landmarks could anchor the firing-rate map. In the grid-cell recordings we analyzed, the mice never needed to pay attention to visual cues during traversals of the linear track. Little is known about the modulation of entorhinal grid cells by attention, though the presence of goals can change the grid pattern (Boccara et al., 2019; Bray, 2019; Butler et al., 2019). In the absence of an explicitly spatial task for the animals, it is hard to say whether firing pauses in grid cells and coherent firing field drift have behavioral consequences. We can only speculate that mice should be better at tasks that require them to estimate relative distances, rather than tasks that require them to locate objects in absolute coordinates.

Much more is known about the role of attention on place and time cells in hippocampus. In fact, place fields in mice are often unstable unless the mice are explicitly engaged in a spatial task (Kentros et al., 2004; Muzzio et al., 2009); in virtual reality experiments, only 15% of place fields are stable over several days (Ziv et al., 2013). The set of active hippocampal cells changes over the time-scale of minutes (Mau et al., 2018); in addition, multiple spatial and temporal representations alternate on time-scales of several tens of milliseconds to seconds (Kelemen and Fenton, 2016). Kinsky et al. find that the hippocampal map in mice rotates coherently when the arena is rotated, but not necessarily in the same direction as the external rotation (Kinsky et al., 2018). These map rotations explain some of the place-field instabil-

ity. Head-direction and grid cells presumably rotate their representations in concert with hippocampal cells.

Rats, however, seem to have stable grid-cell representations, even when the animals are forced to reorient (Weiss et al., 2017). There could well be species-specific differences in the amount of drift grid fields exhibit, even on the time-scale of minutes that affect measures of variability. It is conceivable, as well, that drift accumulates more in mice than in rats. Accumulated drift would manifest itself in the long-term loss of place-field stability for mice; rats, on the other hand, might pay more attention to distal visual and spatial geometry cues that help maintain the hippocampal map in register over longer time periods. We note that both rat and mouse grid cells exhibit similar amounts of excess variability (Fig. S1).

Computational models propose that grid cells give rise to place fields (Solstad et al., 2006; Rolls et al., 2006). So it is not surprising that CA1 and CA3 place fields shift, too (Mehta and McNaughton, 1997; Lee and Knierim, 2007; Roth et al., 2012). Given the reciprocal connections between hippocampus and entorhinal cortex (Marozzi et al., 2015), an unanswered question is whether the MEC grid-field drift drives hippocampal place-field shifts, or vice versa. Models of the entorhinal-hippocampal interaction might permit quantitative predictions of how much “noise” is transferred from one area to the other, and how much of the apparent noise is due to correlated shifts in the spatial maps (Rolls et al., 2006; Monsalve-Mercado and Leibold, 2017).

While we found that grid-cell discharge is highly variable in both 1D and 2D environments, we have only established a direct link between firing-field

dynamics and variability for the linear-track recordings. Several approaches have been used to measure the dynamics of grid spatial phases and single grid-field locations in 2D recordings in the open field: For example, Hardcastle et al. (2015) suggest a “spike-distance metric” to compare inbound and outbound spiking on firing-field crossings. Krupic et al. (2018) and Hägglund et al. (2019) use flow maps to study the dynamics of grid fields in response to changing the shape of the arena. Such attempts are fraught with difficulties, however. Grid cells often fire at low rates, even at the center of a firing field. As a consequence, the firing fields only become apparent after the animal’s trajectory has passed through each location multiple times. Moreover, unlike in experiments on virtual linear tracks, recording from animals foraging in 2D arenas does not permit dividing the data into well-defined “trials”. Nevertheless, stochastic mixture models could reveal whether the encoding of space is static or not. As drifting firing fields provide a generative model of grid-cell variability, one can fit such models to spike-count data and thereby deduce the range and variance of the field shifts, even if it is impossible to measure the shifts directly. Therefore, further research into variability could give us insight into the dynamics of two-dimensional firing fields. Moreover, understanding how the excess variability in grid-cell discharge arises will allow future studies to address how the nervous system might compensate for this variability.

Materials and methods

Data sources

We analyzed multiple data sets (Fyhn et al., 2004; Sargolini et al., 2006; Fyhn et al., 2007; Stensola et al., 2012; Domnisoru et al., 2013; Pérez-Escobar et al., 2016), focusing on tetrode recordings from Pérez-Escobar et al. (2016) and whole-cell data from Domnisoru et al. (2013).

The tetrode recordings from Pérez-Escobar et al. (2016) are available at <https://datadryad.org/resource/doi:10.5061/dryad.c261c>. For details on the experimental setting, see the original publication. The data set contains recordings entorhinal neurons from male wild-type C57BL mice. From these neurons grid cells were identified by letting mice run in a 70 cm \times cm arena (Pérez-Escobar et al., 2016). The mice were then put on a 80 cm \times 5.6 cm linear track. Positions were projected onto the long axis of the linear track. Coordinates were measured from the center of the track. Each trial yielded of a trajectory segment that extends ± 30 cm from of the track's midpoint. To eliminate the animal's running direction as a potential factor contributing to the trial-to-trial variability (Pröll et al., 2018), we split the data into runs toward the right and the left end of the track.

Whole-cell data from Domnisoru et al. (2013) were kindly provided by Cristina Domnisoru and David Tank. For experimental details, see the original publication. In these virtual-reality experiments, male wild-type C57BL mice were head-fixed and ran only in one direction. To start a new trial the animal was reset within the virtual-reality environment. The data consisted of 27 grid-cells recordings. Three cells were excluded from the analysis because they had a low number of trials ($n < 10$) or low firing rate ($f < 0.5$

Hz).

Calculation of firing-rate maps

We estimated a spatially discretized 2D firing-rate map F_{x_i, y_i} for each cell by counting the number of spikes in 2.4 cm^2 bins and dividing this number by the total time the animal spent in the spatial bin. In 1D, spikes and trajectory points were binned for each trial. For the real linear track, we used a bin size of 2.4 cm , and 8 cm for the VR data. These different bin sizes were used as mice ran faster in VR, as described in (Domnisoru et al., 2013), and the VR tracks were longer than the real tracks. The ratio of spike count to dwell time in each bin yielded an estimate of the spatial firing rate for individual trials. The firing-rate map was then calculated as the trial-averaged firing rate.

Spike-count distributions

To study the discharge variability, we discretized the trajectory (x_t, y_t) (or, in 1D, (x_t)) in steps of $\Delta t = 20 \text{ ms}$ and discarded trajectory points for which the rate $F(t) < F_{\min} = 5 \text{ Hz}$. Next, trajectory segments were chosen such that the expected number of spikes, $n = \sum_t F_{x_t, y_t} \cdot \Delta t$, was $n = 5$, unless otherwise stated. The actually measured number of spikes on such trajectory segments yielded a spike-count distribution, which was compared with the predicted spike count. The constraint $F(t) \geq 5 \text{ Hz}$ induced no statistically significant change in the count distributions for any of the 106 cells, but it restricted the duration of the trajectory segments to less than one second.

Field detection and jitter on linear tracks

To identify and follow individual firing fields across trials, we first detected peaks in the tuning curves. A peak was defined as a local maximum whenever the difference to the next closest local minimum was larger than 20% of the global maximum. The region between the two adjacent minima was defined as a firing field; regions outside these areas were not considered further. Each peak marked one firing field. Next, for each peak we computed

$$\langle x_{\text{field}} \rangle = \frac{\sum_i f_i x_i}{\sum_i f_i}$$

$$\langle \sigma_{\text{field}}^2 \rangle = \frac{\sum_i f_i (x_i - \langle x_{\text{field}} \rangle)^2}{\sum_i f_i},$$

where x_i are the bin centers and f_i is the firing rate on each trial. The field centers and widths were computed for each individual trial using the formulae above. If a cell did not spike as a field was crossed, then the corresponding firing field was undefined on that particular trial.

This approach assigns the field location to the “center-of-mass” of the firing field, rather than the trial-by-trial peak in the firing rate. The distance between field locations from one trial to the next trial can vary, with the limitation that these positions cannot fall outside the field boundaries (demarcated by the two local minima surrounding each trial-averaged peak), nor occur at the track boundaries, which correspond to positions more distant than 30 cm away from the center of the track.

A firing field was considered to be reliable whenever it was identified on left and right runs and the distance between the respective field boundaries was larger than 9.6 cm (4 bins).

Single-cell trial alignment

To align the spatial firing-rate profiles across trials, we first detected the location that elicited the largest number of spikes for each trial. We then computed the firing rate's centers-of-mass in a window of ± 7 bins around the detected peak. The firing rate-maps were then shifted to bring these centers-of-mass into alignment, while zero-padding the maps as necessary. This alignment procedure is illustrated in Fig. S9.

Software

All analyses were performed in the Python scripting language (Python Software Foundation. Python Language Reference, version 3.4, available at <http://www.python.org>) including the packages numpy (version 1.14.3), scipy (version 1.0.1) and sklearn (version 0.19.1).

Author contributions

This study was designed by JN, MS, and AVMH. Research, programming and data analyses were performed by JN. The manuscript was jointly written by JN, MS, and AVMH.

Acknowledgements

The authors would like to thank K. Allen, C.Domnisoru, J. Pérez-Escobar and D.W.Tank for making data from Domnisoru et al. (2013) and Pérez-Escobar et al. (2016) available. In particular, we thank Kevin Allen for his advice and helping us understand the differences in firing field properties

between rats and mice.

Funding information

This work was supported by the German Federal Ministry for Education and Research Grant 01GQ0440.

Declaration of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Figures

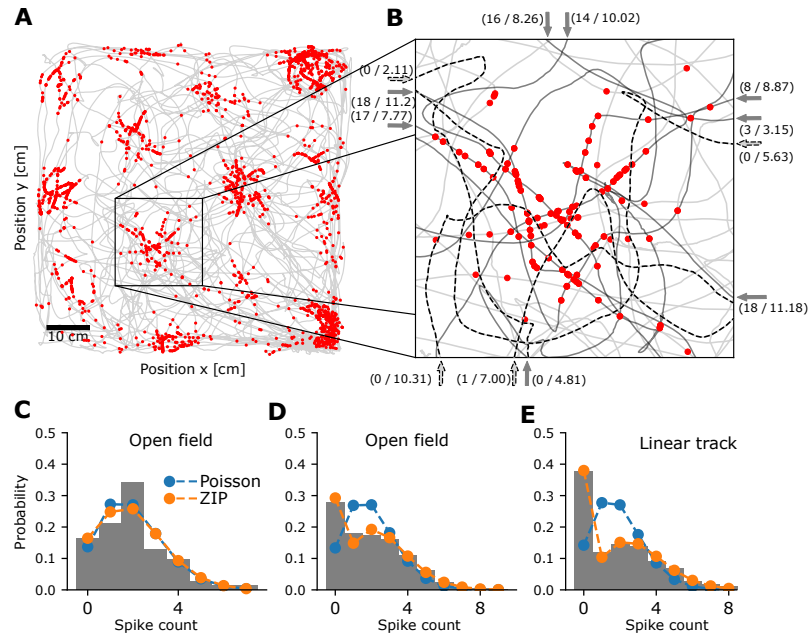


Figure 1: Variation in the grid-cell spike counts on single firing-field crossings recorded by Stensola et al. (2012) (panels A and B, cell T6C3) and Pérez-Escobar et al. (2016) (panels C-E). **A:** Trajectory of the rat (grey) with spike positions (red). **B:** Expanded view of one firing field. Spikes elicited during individual crossings are indicated in red. Selected paths longer than 50 cm (thick lines) are labelled with the observed number of spikes s and the number of spikes n expected from the firing rate map. Dashed black lines denote trajectories without any spikes. (**C-E**) Histograms of spike counts for one grid cell recorded in an open field (**C**) and another cell recorded in an open field (**D**) and on a linear track (**E**). For these histograms the cell's firing rate is integrated along segments of the animal's path until the expected number of spikes reaches a fixed threshold, here $n = 2$. The count distributions (grey) are fit by a Poisson (blue) and Zero-Inflated Poisson (ZIP) model (orange). A likelihood ratio test favors the Poisson model for the distribution in **C**, and the ZIP model for the examples in **D** and **E**.

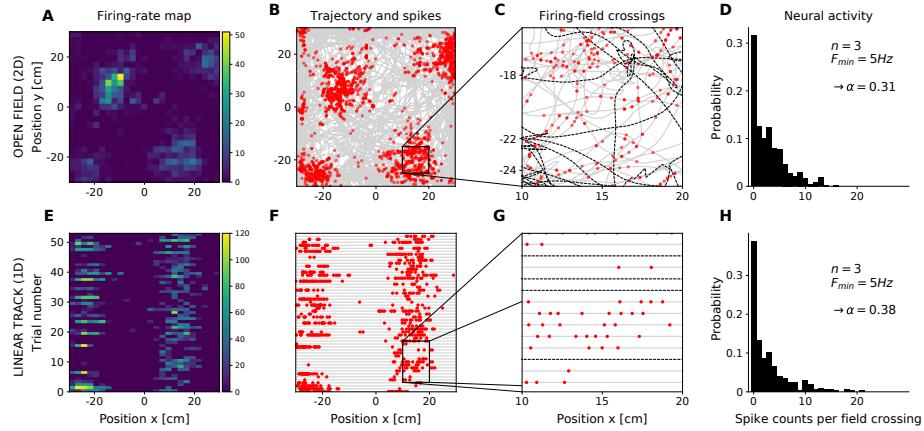


Figure 2: Example recordings of the same grid cell in the open field (top) and on the linear track, as measured by Pérez-Escobar et al. (2016) (bottom). **A** and **E** show the color-coded firing-rate histograms, obtained from spikes depicted in red in **B** and **F**, respectively. As shown by the expanded view in **C**, some crossings in the open field are associated with no or only few spikes. The same phenomenon is observed on the linear track in **G**. **D** and **H**: Spike-count histograms collected over trajectory segments for which the firing-rate map predicts that 3 spikes will be elicited. Nevertheless, spike counts of zero are frequently recorded. For these two examples, the ZIP-framework suggests that the cell is in its inactive (“zero”) state for 31% (2D) and 38% (1D) of the recorded time. Data shown for recording jp4312-27032016-0107, cell T8C3.

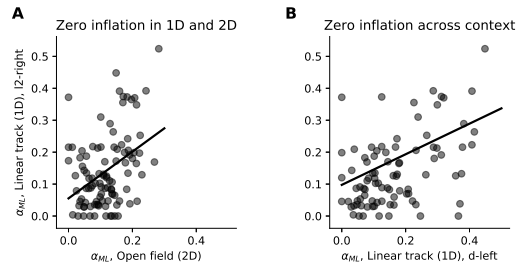


Figure 3: In the experiments of Pérez-Escobar et al. (2016), mice first navigated in an open field (2D) before they ran back and forth on a linear track in light or darkness. **A:** The zero-inflation parameter α in 1D is predicted by its value measured in the preceding 2D session. The solid line depicts the linear regression for the data ($r = 0.40$, $p < .001$). **B:** The proportion of excess zero counts remains similar across contexts and running directions. In particular, runs in darkness lead to similar zero-inflation measures as runs in light. The solid line depicts the linear regression for the data ($r = 0.47$, $p < .001$). Linear-track data for other combinations of lighting context and running direction are shown in Suppl. Fig. S3.

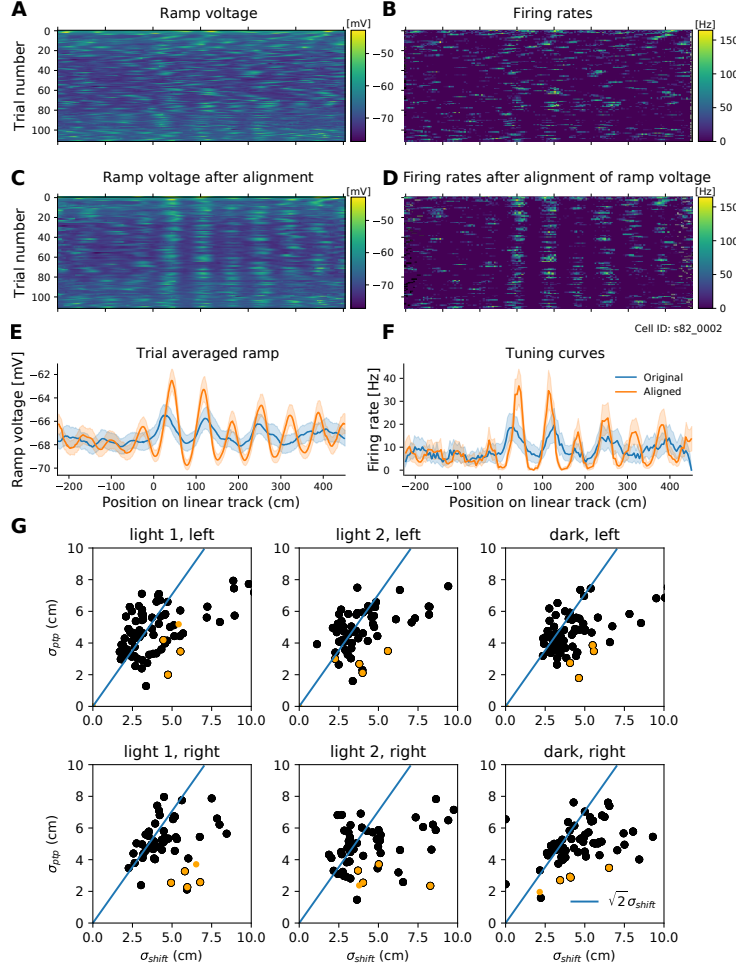


Figure 4: Trial alignment of data recorded on virtual linear tracks (A-F) by Domnisoru et al. (2013) (cell s82.0002) and on real linear tracks (G) by Pérez-Escobar et al. (2016). Ramp voltage (A) and firing rates (B) before alignment. C: Ramp voltage after alignment. D: Firing rates calculated after ramp-voltages were aligned. In E and F the trial-averaged voltage ramps and tuning curves, respectively, are shown before (blue) and after alignment (orange). G: Standard deviation of shifted trials σ_{shift} versus standard deviation of firing-field periods σ_{ptp} . Black and circled black: Cells that passed the selection criteria (two firing fields and robust autocorrelation, cf. Table 3). Orange: The simultaneously recorded cells shown in Fig. 5. The blue line depicts the relation $\sigma_{ptp} = \sqrt{2} \cdot \sigma_{shift}$ predicted for independent jitter. For the majority of the cells σ_{ptp} is smaller than expected from uncorrelated field shifts.

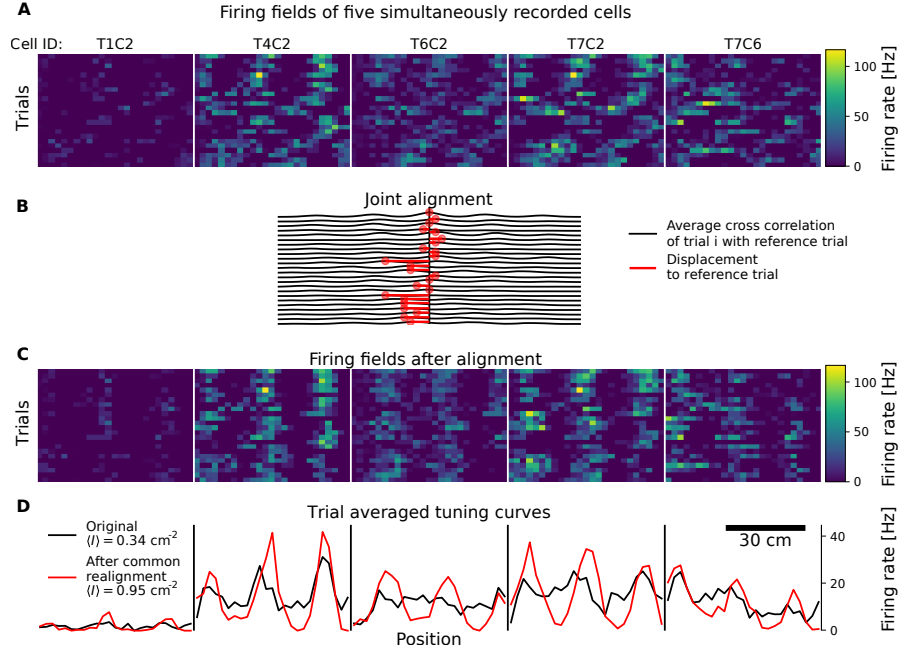


Figure 5: **A**: Firing rates of simultaneously recorded cells while the animal repeatedly ran along a real linear track. Pérez-Escobar et al. (2016) classified these cells as follows: T1C2: “spatially modulated”; T4C2 and T6C2: “speed-modulated grid”; T7C2: “unclassified”; T7C6: “speed-modulated head direction”. **B**: Cross-correlation between the first trial and subsequent trials. The location of the first peak in the cross-correlation (relative to $x = 0$) sets the spatial shift to bring each subsequent trial into alignment with the first trial. These shifts are displayed as red horizontal lines. The procedure uses concatenated spike trains of the five cells and does not optimize the alignment on a cell-by-cell basis. **C**: Firing rate maps after alignment, to be compared to the unaligned firing rates in **A**. **D**: After alignment, the firing-fields became more prominent and some fields emerged that were originally below detection threshold. Zero-inflation is observed both before and after alignment: the ZIP model is preferred over the Poisson model ($p < 0.001$, $df = 1$) for all five cells. Measured in terms of the aligned fields, the Fisher information $\langle I \rangle = \langle \mathbf{f}'(x)^T \Sigma^{-1} \mathbf{f}'(x) \rangle_x$ increases by 180%. Data shown for five cells of recording jp2098-03042016-0107, right runs in light condition 1 (11).

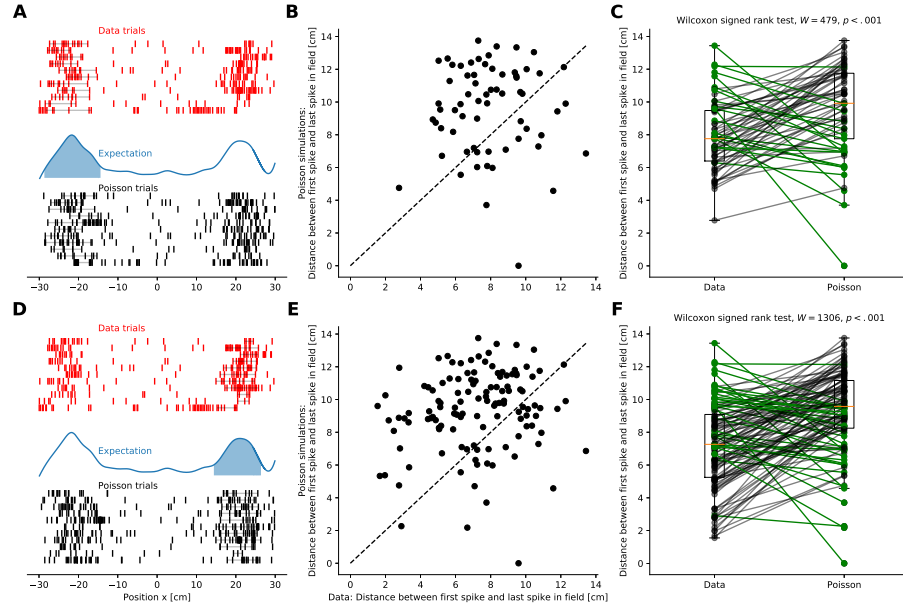


Figure 6: Spike trains within a firing field cover shorter distances than expected from firing-rate maps. In this example, a grid cell recorded on the linear track displayed two grid fields (Pérez-Escobar et al., 2016). Field boundaries were determined by the time-points when the firing rate reaches 25% of the peak firing rate. **A** Top: Raster plot of grid cell spikes for the first 11 trials. Horizontal lines connect the first to last spike within the field boundary of the first field on the left. Middle: Trial-averaged firing rate is shown in blue. The left-hand field area is shaded in light blue. Bottom: 11 simulated spike trains of Poisson spikes with the same time-varying firing rate and the same speed of the animal (Poisson process with expectations computed from the animal trajectory and tuning curve). **B**: Comparison of the distances between the first and last spike for all 125 recorded trials (both real data and surrogate trials). **C**: Same as in (A-C), but now for the right-hand firing field.

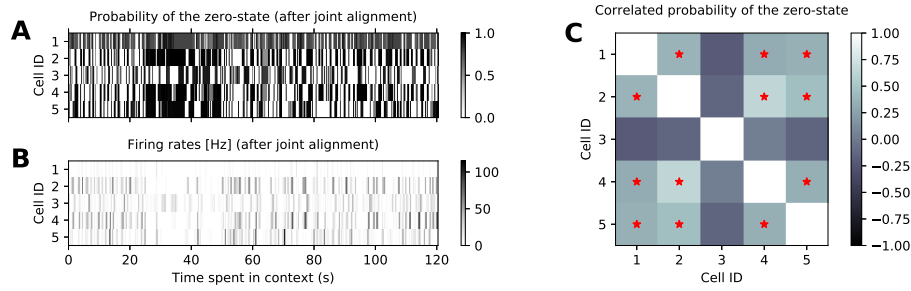


Figure 7: Coherence of the state transitions across the five simultaneously recorded cells shown in Fig. 5. Spike trains were first jointly aligned across all cells to counteract the coherent firing field drift. The phenomenon of zero-inflated spike counts still existed after alignment. **A**: Probability of the zero-state based on the expected spike counts $n_i(t)$ and the average zero-state probability α_i for each cell (cf. SI). Expected spike counts were taken from the jointly-aligned trial averages, measured in 200 ms long time windows. **B**: For comparison, the estimated firing rate for each time point, using the trial averaged tuning curves after joint-alignment. **C**: For each of the jointly-aligned cells, we estimated the time-resolved firing rate and the zero-inflation parameter α from an LNP model (cf. SI). With these parameters, we then simulated 40 independent, surrogate ZIP models for the five cells. By comparing the real data to the surrogates, we asked whether the zero-state transitions were more highly correlated across cells than expected by chance. To be deemed significantly correlated, the real correlation value had to be greater than the 95% percentile of the simulated correlations (marked by red stars).

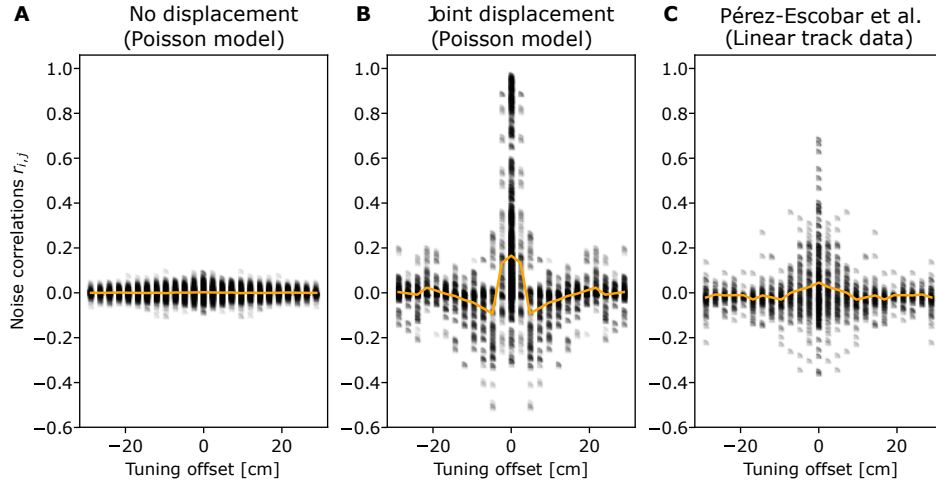


Figure 8: Shared trial-to-trial variations in the spatial phases increase noise correlations between otherwise independent Poisson grid cells. Firing-rate profiles for model cells were simulated as slices through 2D grids. The correlations are plotted against the tuning offset measured from the spatial cross-correlation between cells. **A**: Noise correlations in a Poisson simulation with static tuning curves across trials. **B**: Poisson simulation with jointly displaced tuning curves on individual trials across cells. **C**: Noise correlations measured from experimental grid-cell data (linear track, $N=138$, data shown for runs to right end of the track from all light conditions).

Tables

	N_{cells}	mean α open field	mean α linear track
grid	48	0.11 ± 0.01	0.14 ± 0.01
grid conjunctive	67	0.16 ± 0.01	0.17 ± 0.01
spatial	42	0.14 ± 0.02	0.12 ± 0.01
spatial conjunctive	50	0.17 ± 0.01	0.14 ± 0.01

Table 1: Conjunctively tuned cells recorded in Pérez-Escobar et al. (2016), which respond not only to position, but also to speed or head direction, tended to exhibit more zero inflation. The table lists the average zero-inflation probability α and their standard errors for different classes of spatially modulated cells, which were recorded in the open field and on the linear track. For the linear track data the average of α across all six experimental conditions (left/right running direction; l1/l2/d lighting context) is displayed and for statistical testing all the α measurements were used. One hypothesis is that the difference in zero-inflation between conjunctive and non-conjunctive cells should be less in 1D, as the head direction and speed vary less on the linear track; this expectation is borne out for grid cells. For grid cells and other spatially selective cells we tested the likelihood that a randomly selected estimate of α from a conjunctive cell was larger than a randomly selected estimate from a non-conjunctive cell; specifically, the Mean-Whitney U-Test asks whether one can reject the null hypothesis that this likelihood is 1/2. Levene tests for variance homogeneity showed no significant differences between the groups. On the linear track, grid cells had a U-statistic for conjunctive vs. non-conjunctive α of $U = 38279$, $p < .001$; whereas in the open field $U = 753$, $p < .001$. In the open field, non-grid, but spatially modulated cells had significantly different α if they were conjunctive ($U = 760$, $p = .012$); on the linear track, though, the differences between conjunctive and non-conjunctive spatially modulated cells were not significant ($U = 29034$, $p = .379$).

	Number of significant contexts:		
	1	4	6
grid	48/70 (69%)	21/70 (15 %)	3/70 (4 %)
grid conjunctive	65/83 (67%)	23/83 (28%)	14/83 (14%)
spatial	32/61 (52%)	3/61 (5%)	1/61 (2%)
spatial conjunctive	43/67 (64%)	7/67 (10%)	1/67 (1 %)

Table 2: Pairs of spatially modulated neurons with firing fields that drifted did so with a significant degree of coherence. In some experiments, the total amount of drift measured was small, in which case it was more difficult to determine that significant common field displacement had occurred. The table shows the numbers (in brackets, the percentages) of spatially modulated neuron pairs with significant common field displacement for the experiments of Pérez-Escobar et al. (2016) (see Fig. S7 for details of how significant coherence in the field shifts was assessed). Each combination of lighting context and running direction was treated as a separate experiment.

Context	$\langle \sigma_{\text{shift}} \rangle$ [cm]	$\langle \sigma_{\text{ptp}} \rangle$ [cm]	H1	U	p
left l1 (N=75)	3.94 ± 0.23	4.77 ± 0.17	52%	1041	0.04
left l2 (N=60)	3.95 ± 0.22	4.79 ± 0.17	58%	625	0.03
left d (N=71)	4.33 ± 0.23	4.65 ± 0.15	77%	404	< 0.001
right l1 (N=49)	4.61 ± 0.21	5.10 ± 0.21	73%	249	< 0.001
right l2 (N=56)	4.61 ± 0.30	4.80 ± 0.19	68%	369	< 0.001
right d (N=61)	4.82 ± 0.26	5.04 ± 0.16	75%	240	< 0.001

Table 3: Comparing the firing-field shifts to the trial-by-trial distances between firing fields. On individual trials, less variability in the distances between firing fields is observed than would be expected for independent jitter in the firing-field positions. We quantified the variability in the peak-to-peak distances and firing-field centers in terms of the standard deviations σ_{ptp} and σ_{shift} . H1: Fraction of cells with $\sigma_{\text{ptp}} < \sqrt{2} \sigma_{\text{shift}}$. N: total number of cells considered, U and p : test statistics and p -value of the Wilcoxon signed-rank test for testing the null hypothesis that the differences between σ_{ptp} and $\sqrt{2} \sigma_{\text{shift}}$ are symmetrically distributed around zero.

Supplementary Figures

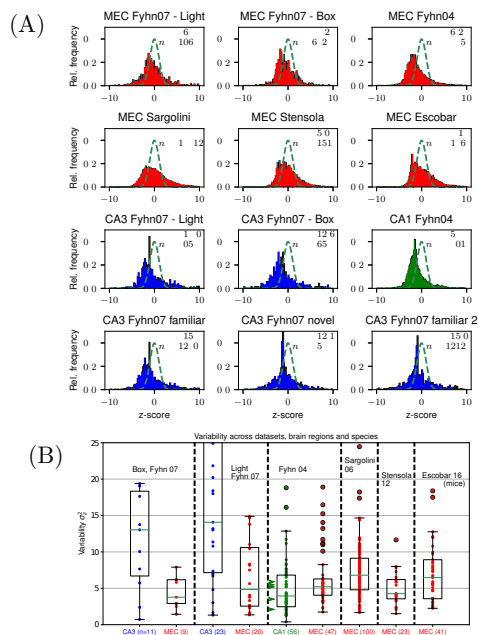


Figure S1: **A:** Pooled spike-count data from different experiments and areas within the hippocampal formation, collected in 2.4 cm^2 bins and then z-scored (color histograms); for comparison, the dashed green curve depicts a standard normal distribution with unit variance. The data sets are: (Fyhn04: Fyhn et al. (2004), Sargolini: Sargolini et al. (2006), Fyhn07: Fyhn et al. (2007), Stensola: Stensola et al. (2012), Escobar: Pérez-Escobar et al. (2016), familiar: Box and light condition, novel: Cylinder and dark condition). The data from Pérez-Escobar et al. (2016) were recorded in mice; all other data in rats. Note that the variance of the distributions of CA3 place cells dropped after remapping (Cylinder and dark sessions, see also Suppl. Fig. S1). **B:** Consistent with previous findings (green arrow heads, values taken from Fenton et al. (2010)) CA1 z-score variances (green dots) are higher than expected from Poisson. The same is true for MEC grid cells (red) and CA3 place cells (blue). The number of cells for each distribution is given in parentheses on the x-axis labels.

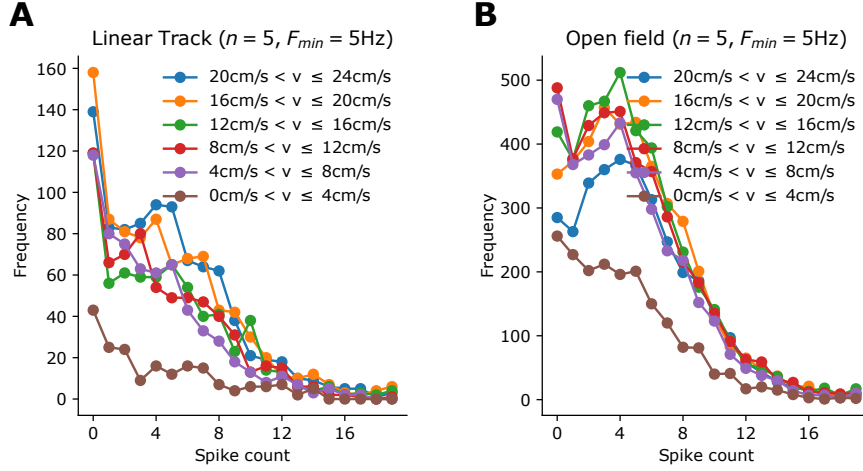


Figure S2: Zero inflation is observed for all speed ranges on the linear track (**A**) and in the open field (**B**). Pooled distributions of $N=61$ non-conjunctive grid cells are shown. Likelihood-ratio tests preferred zero-inflated models over Poisson for all ranges that were tested. The excess-zero probabilities α of zero-inflated negative binomial (ZINB) fits did not correlate with the speed range (ZINB, linear track: $r = 0.5$, $p = 0.31$; open field: $r = 0.12$, $p = 0.82$, $N = 6$). In comparison to the ZIP model, for which the additional variance is introduced through an excess of zero spike counts, the ZINB model also adds variance to the spiking mode: the variance of the negative binomial distribution is larger than the mean (unlike in the Poisson model), so that the spiking regime in the ZINB model intrinsically exhibits overdispersion.

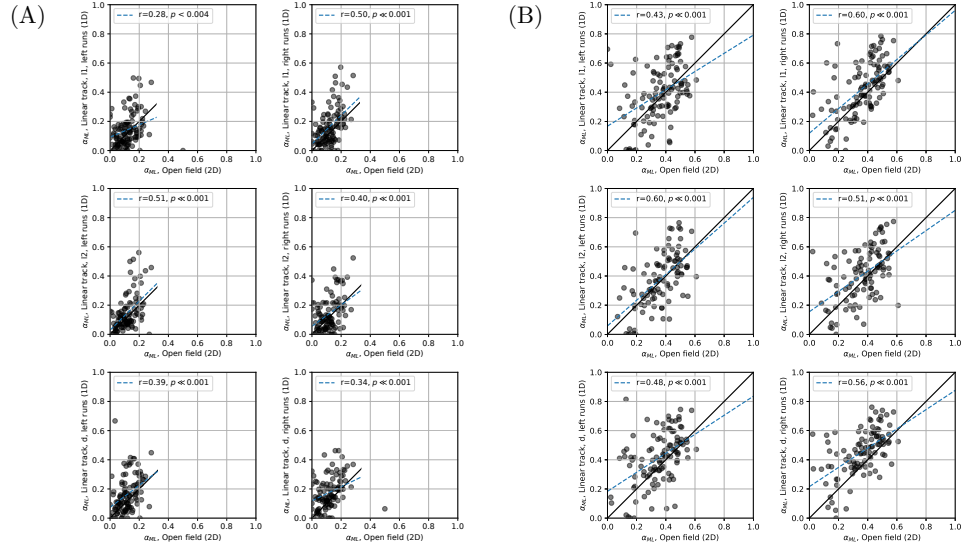


Figure S3: Scatter plots and regression lines for the zero-inflation parameter α of grid cells measured separately in 1D and 2D (Pérez-Escobar et al., 2016). Each panel shows a different lighting context. **A:** Estimated α -values for spike count distributions with expected spike count $n = 5$ and trajectory segments with a minimum firing rate of $F_{min} = 5$ Hz. **B:** For comparison, the α -estimates for spike counts with a lower expected number $n = 1$ of spikes (with the same $F_{min} = 5$ Hz constraint for the trajectories)

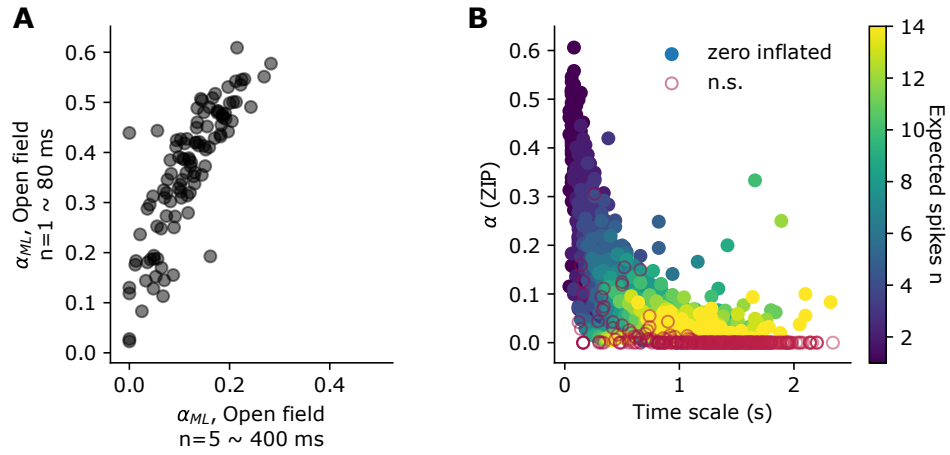


Figure S4: **A:** Zero inflation when the expected spike-count number is $n = 1$ vs. the same measure for $n = 5$. **B:** Significant zero inflation is found across time scales, which denote the average integration windows needed to accumulate an average of n spikes, as given by the firing-rate maps. For all filled points, the ZIP model was preferred in a likelihood ratio test over the Poisson model at a significance level $p < 0.001$.

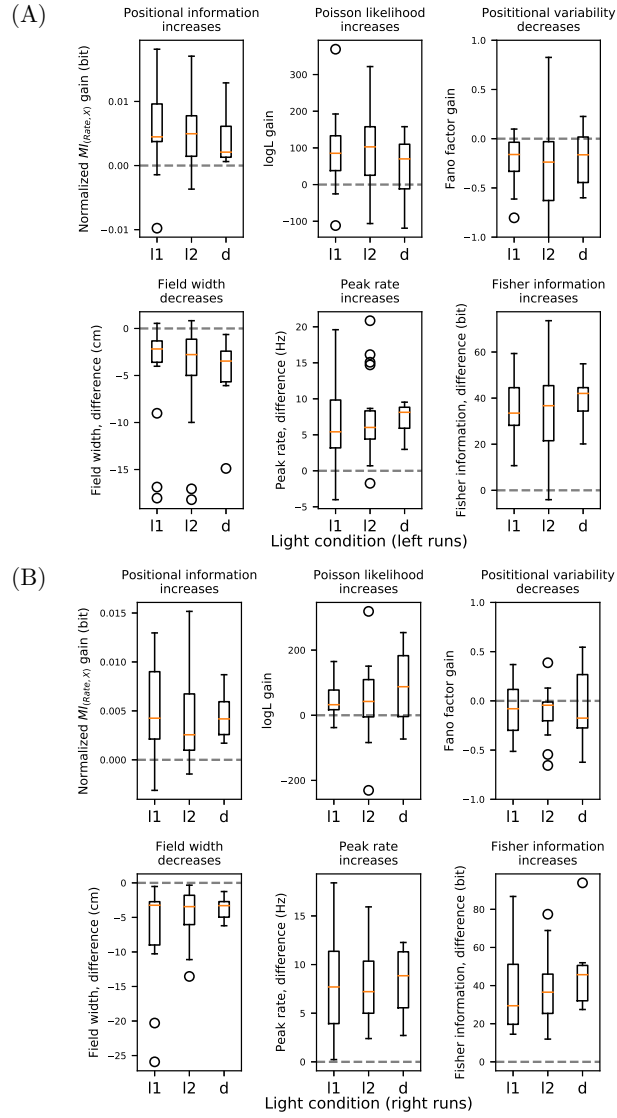


Figure S5: Effects of trial alignment on grid-cell data measured on the linear track. Box-plots are shown for various measures of firing field properties under three light conditions (l1: Lighting condition 1, l2: Lighting condition 2 and d: Light off, for details see (Pérez-Escobar et al., 2016)). **A**: runs towards the left end of the track **B**: runs to the right end of the track.

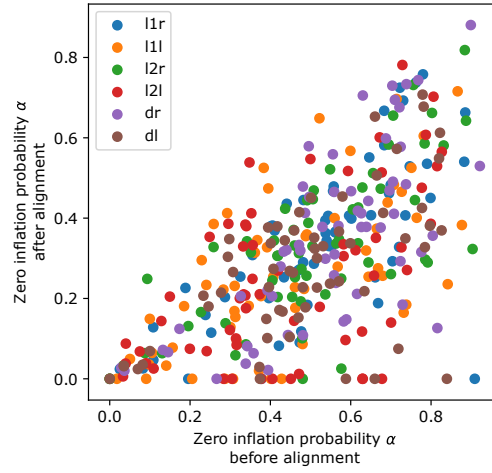


Figure S6: Maximum-likelihood estimates of the zero-inflation probability α are reduced after grid-field alignment. Data shown for grid cells across light contexts (l1, l2 and d indicated by different colors, see legend) on runs towards the right (r) of left (l) end of the track. The probability α for being in the zero-state was estimated from a zero-inflated LNP model (see Fig. S9).

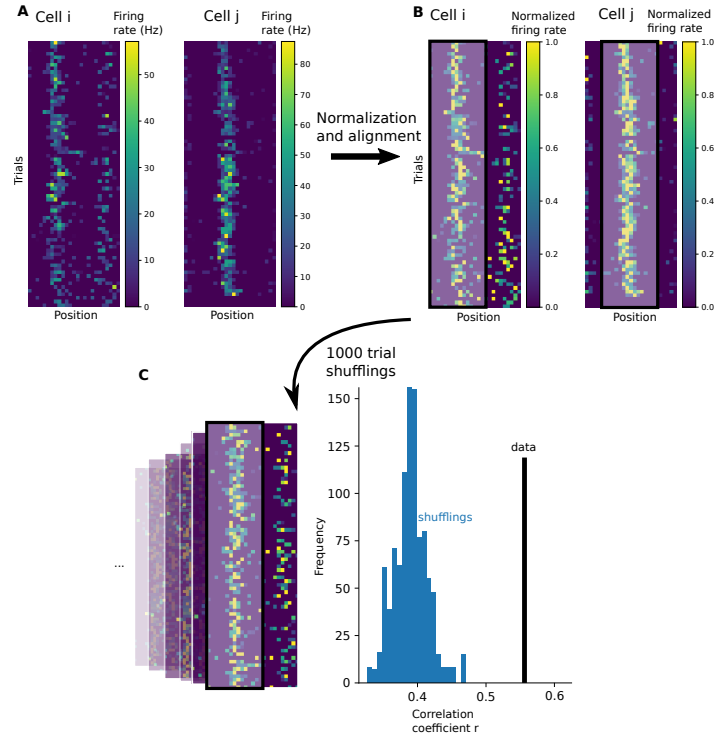


Figure S7: If grid fields shift coherently across cells over time, a field shift in one cell should be mirrored by comparable field shifts in other cells. To test this hypothesis, we normalized the rate maps of each cell (as shown in **A**) by the respective average peak firing rate, so as to eliminate trends in the firing rate over time. The *trial-averaged* firing fields were then aligned to achieve maximal overlap. To this end, we shifted the data on every trial for a given cell by a common amount to bring the ensemble of rate maps into alignment (**B**). This procedure preserves relative drifts of the firing fields. In the last step, we compared the Pearson correlations of the firing-rate maps across cells for simultaneous trials to the correlation for randomly shuffled trials (**C**). If the firing fields drifted jointly, then the Pearson correlation will be higher than for shuffled trials, as seen on the right. To determine significance of the differences in the correlations, we performed a Bonferroni correction for multiple testing: A cell-pair was assumed to show significant firing consistency whenever the p-value of the linear regression was smaller than $0.05/(n_{bins}n_{cells})$. If the correlation coefficient of the original trial ordering exceeded the 95 percentile of the shuffled distribution, the cells were considered to have undergone common shifts over different trials.

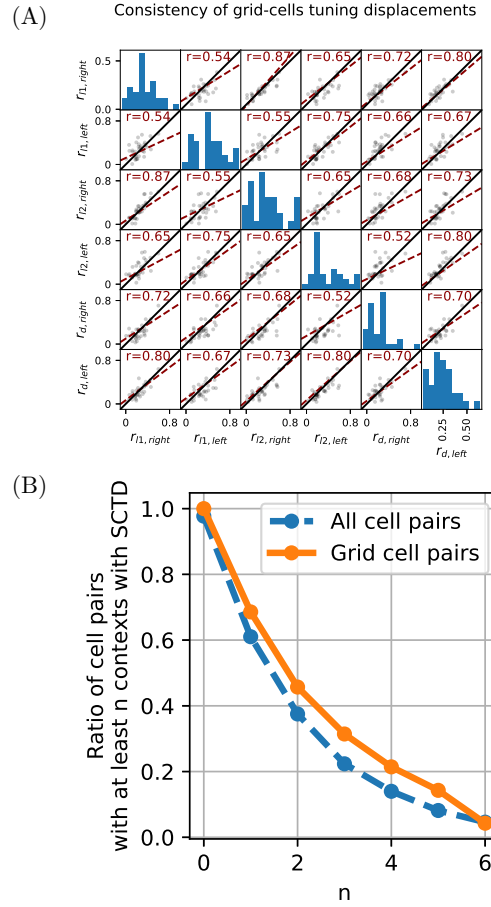


Figure S8: Pairs of grid cells exhibited consistent field jitter across contexts in the experiments of (Pérez-Escobar et al., 2016). Cell pairs with consistent tuning displacement in one context are likely to also have consistent tuning displacement in other settings. **A:** Scatter matrix of rate correlations (r with subscripts denoting the context) shown for pairs of grid cells that exhibit significant consistency in at least two contexts. Linear regression shown as dashed line with the corresponding r -value in dark-red (all regressions were highly significant, $p < 0.01$). On the diagonal the distributions of the rate correlations are displayed. **B:** About 40% of all cell pairs show significant consistency under tuning displacement (SCTD) in at least 2 ($n=2$) contexts. Grid cells (orange) show a slightly increased ratio.




Figure S9: To evaluate the alignment we computed linear-nonlinear Poisson (LNP) fits using only the animal's position as input signal. Output signals were spike counts collected in 200ms time windows. Based on the alignment parameters (green arrow) the spatial design matrices (25 bins, 2.4 cm bin width) were transformed (red arrow) to bring them into alignment. Following this shift, firing fields were fitted as before. This approach is similar to the time-shifting described by Hardcastle et al. (2015), except that the data is shifted in space, not in time, and separately for each trial.

Supporting information

Notation.

- $\{x\}$: Set with elements x_i
- P : Probability density function or probability mass function on a set
- $P(s|n)$: Probability of s given n
- $\langle x \rangle = \sum x P(X = x)$: the expected value of a random variable X , where the sum is taken over all possible values of X
- $\sigma_x^2 = \langle (x - \langle x \rangle)^2 \rangle$: Variance of x
- $\mathbb{N}_{0,+}, \mathbb{N}_+$: Positive natural numbers with and without zero

S1. The Zero-Inflated Poisson (ZIP) model

When counts have higher variance than expected from a Poisson model (variance $>$ mean), we refer to these data as exhibiting *excess variability* or *overdispersion*. We note that the equality of variance and mean counts in the Poisson model also holds when the mean firing rate changes with time.

The negative binomial model, zero inflation, and hurdle models are widely used in economics and other fields to describe an over-abundance of zero counts (Giles, 2010). Adding excess zeros to a Poisson model increases the variance of the counts and so leads to overdispersion. Such models are called *zero-inflated Poisson* (ZIP) models.

The stochastic process underlying ZIP begins with a random binary variable B , which has probability α of being zero, and probability of $1 - \alpha$ of being one. If $B = 0$, the spike count is automatically zero. Otherwise, a

spike count is drawn from a Poisson distribution. If the excess-zero probability $\alpha = 0$, the ZIP distribution reduces back to the Poisson distribution. Consider a sequence of count measurements $\{s\} \in \mathbb{N}_0^N$ that represent N independent realizations of a random variable S drawn from a ZIP distribution. We normalize the rate parameter of the Poisson distribution $n \rightarrow n(1-\alpha)^{-1}$, so that the expected value of S always remains

$$\langle S \rangle = n, \quad (3)$$

regardless of the value of α . Thus, the probability distribution for a spike count $S = s$ is:

$$p(S = s | n, \alpha) = \begin{cases} \alpha + (1 - \alpha)e^{-n/(1-\alpha)} & \text{for } s = 0 \\ (1 - \alpha)e^{-n/(1-\alpha)} \left(\frac{n}{1-\alpha}\right)^s / s! & \text{for } s > 0 \end{cases} \quad (4)$$

The variance σ_S^2 of a homogeneous ZIP process is:

$$\sigma_S^2 = n + \frac{\alpha}{1 - \alpha} \cdot n^2 \quad (5)$$

This model displays overdispersion for $\alpha > 0$. The implicit assumption in this model is that the two random processes occur on different time-scales, so that they can be cascaded: the random variable B is set for a (possibly variable) time period, and then, provided $B = 1$, spikes are generated using Poisson statistics at a (sub-)millisecond time-scale for that time period. Importantly, the Poisson process need not be “homogeneous”, i.e., the Poisson rate can vary in time; it is the integrated rate that sets the expected spike count n underlying the spike count distribution.

Once α and n has been estimated, for any time point t we can compute

the posterior probability of being in the “zero-firing” state as

$$P(0\text{-state}, t | s_t = 0) = \frac{\alpha}{\alpha + (1 - \alpha) \exp(-n_t)},$$

which uses Bayes’ rule for probabilities. We use this equation to study the correlation in the state transitions across simultaneously recorded cells.

In the Zero-Inflated Negative Binomial (ZINB) model, the Poisson distribution is replaced by a negative Binomial distribution. The probability distribution for the ZINB can be worked out analogously to Eq. (4), as described by Giles (2010).

The zero-inflation index ZI_{idx} is not identical to α . It is defined as

$$ZI_{\text{idx}} = 1 + \ln(p_0)/\mu \quad (1)$$

in the main text, where p_0 is the frequency of zero counts and μ is the mean count. ZI_{idx} is a random variable that depends on the stochastic realization of p_0 and μ . It is defined independently of the underlying statistical model of spike counts.

The expected value of ZI_{idx} as a function of α for the ZIP model will behave as :

$$\langle ZI_{\text{idx}} \rangle = 1 + \left\langle \frac{\log(\alpha + (1 - \alpha)e^{-n_t})}{(1 - \alpha)n_t} \right\rangle_{\alpha, n_t}$$

For the ZIP model, $\langle ZI_{\text{idx}} \rangle \geq \alpha$, and is a saturating function of α .

S2. Model comparison via likelihood ratio tests

The ZIP model adds one parameter (zero-inflation probability α) to the Poisson model; therefore, Poisson is a special case of the more general case of ZIP and is “nested” within the ZIP class. A likelihood ratio test allows

for a fair comparison between the Poisson and ZIP models, as it penalizes ZIP for having one more degree of freedom (i.e., one additional parameter) and similarly it penalizes the shift parameters required to align trials (one additional parameter per trial).

The likelihood ratio (LR) test starts with a firing rate map, which yields predictions for the spike count in time windows of $dt = 200\text{ ms}$. We denote the predicted spike count in the i -th time window as n_i , whereas the observed spike counts are called s_i . Often, n_i is a function of the parameters that describe the statistical model for generating spike counts. We make one key assumption: the spike count in the i -th window is conditionally independent of the spike counts in previous time windows.

The likelihood L for observing a sequence of spike counts $\{s_i\}$ is then $L = \prod_i P(s_i|n_i)$, where $P(s_i|n_i)$ is the probability of a spike count s_i given the expected spike count $n_i = f_i \cdot dt$ at time point i . The $P(s_i|n_i)$ are either computed from the Poisson distribution or the ZIP distribution before or after trial alignment. The likelihood of the model with the lower number of parameters (Poisson without alignment here) is denoted by L_{H_0} and the model with more degrees of freedom is denoted by L_{alt} (ZIP model or trial alignment model). Then the likelihood ratio test statistic is given by the difference in the logarithms of the likelihoods $LR = 2 \cdot [\ln(L_{H_0}) - \ln(L_{\text{alt}})]$.

Wilks' theorem states that LR asymptotically follows a χ^2 -distribution under the null hypothesis H_0 as the sample size becomes large. As noted earlier, the two models being compared, in general, have different numbers of parameters. The difference in the number of parameters sets the degrees

of freedom for the χ^2 -distribution (Wilks, 1938). In practice, the LR -value is compared to the χ^2 value that corresponds to a desired significance level. LR has been shown to have higher statistical power than other competing approaches for model comparison (Solomon, 1975).

References

- C. Barry, R. Hayman, N. Burgess, and K. J. Jeffery. Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience*, 10(6):682–684, 6 2007. doi: 10.1038/nn1905.
- C. N. Boccara, M. Nardin, F. Stella, J. O’Neill, and J. Csicsvari. The entorhinal cognitive map is attracted to goals. *Science*, 363(6434):1443–1447, 3 2019. doi: 10.1126/science.aav4837.
- N. Bray. Grids cells go for a goal. *Nature Reviews Neuroscience*, page 1, 4 2019. doi: 10.1038/s41583-019-0170-0.
- Y. Burak and I. R. Fiete. Accurate Path Integration in Continuous Attractor Network Models of Grid Cells. *PLoS Computational Biology*, 5(2): e1000291, 2 2009. doi: 10.1371/journal.pcbi.1000291.
- W. N. Butler, K. Hardcastle, and L. M. Giocomo. Remembered reward locations restructure entorhinal spatial maps. *Science*, 363(6434):1447–1452, 3 2019. doi: 10.1126/science.aav5297.
- G. Chen, D. Manson, F. Cacucci, and T. J. Wills. Absence of Visual Input Results in the Disruption of Grid Cell Firing in the Mouse. *Current biology*, 26(17):1–8, 2016. doi: 10.1016/j.cub.2016.06.043.
- L. De Almeida, M. Idiart, A. Villavicencio, and J. Lisman. Alternating predictive and short-term memory modes of entorhinal grid cells. *Hippocampus*, 22(8):1647–1651, 8 2012. doi: 10.1002/hipo.22030.

- G. W. Diehl, O. J. Hon, S. Leutgeb, and J. K. Leutgeb. Grid and Nongrid Cells in Medial Entorhinal Cortex Represent Spatial Location and Environmental Features with Complementary Coding Schemes. *Neuron*, 94(1): 83–92, 2017. doi: 10.1016/j.neuron.2017.03.004.
- C. Domnisoru, A. A. Kinkhabwala, and D. W. Tank. Membrane potential dynamics of grid cells. *Nature*, 495(7440):199–204, 2013. doi: 10.1038/nature11973.
- B. Dunn, M. Morreaunet, and Y. Roudi. Correlations and functional connections in a population of grid cells. *PLoS Comput Biol*, 11(2):e1004052, 2015. doi: 10.1371/journal.pcbi.1004052.
- B. Dunn, D. Wennberg, Z. Huang, and Y. Roudi. Grid cells show field-to-field variability and this explains the aperiodic response of inhibitory interneurons. *arXiv preprint*, pages 1–33, 2017.
- H. Eichenbaum, D. G. Amaral, E. A. Buffalo, G. Buzsáki, N. Cohen, L. Davachi, L. Frank, S. Heckers, R. G. M. Morris, E. I. Moser, L. Nadel, J. O’Keefe, A. Preston, C. Ranganath, A. Silva, and M. Witter. Hippocampus at 25. *Hippocampus*, 26(10):1238–1249, 2016. doi: 10.1002/hipo.22616.
- A. A. Fenton and R. U. Muller. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences of the United States of America*, 95(6): 3182–7, 3 1998.
- A. A. Fenton, W. W. Lytton, J. M. Barry, P.-P. Lenck-Santini, L. E. Zinyuk, S. Kubík, J. Bures, B. Poucet, R. U. Muller, and A. V. Olypher. Attention-

like modulation of hippocampus place cell discharge. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(13): 4613–25, 3 2010.

M. Fyhn, S. Molden, M. P. Witter, E. I. Moser, and M.-B. Moser. Spatial representation in the entorhinal cortex. *Science*, 305(5688):1258–64, 8 2004.

M. Fyhn, T. Hafting, A. Treves, M.-B. Moser, and E. I. Moser. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132): 190–4, 3 2007. doi: 10.1038/nature05601.

I. Gat, N. Tishby, and M. Abeles. Hidden Markov modelling of simultaneously recorded cells in the associative cortex of behaving monkeys. *Network: Computation in Neural Systems*, 8(3):297–322, 1997. doi: 10.1088/0954-898X/8/3/005.

D. Giles. Notes on the Zero-Inflated Poisson Regression Model. *Department of Economics, University of Victoria*, 2010.

W. H. Greene. Accounting for Excess Zeros and Sample Selection in Poisson and Negative Binomial Regression Models. *NYU Working Paper No. EC-94-10*, pages 1–37, 3 1994. doi: 10.1007/BF00857937.

Y. Gu, S. Lewallen, A. A. Kinkhabwala, C. Domnisoru, K. Yoon, J. L. Gauthier, I. R. Fiete, and D. W. Tank. A Map-like Micro-Organization of Grid Cells in the Medial Entorhinal Cortex. *Cell*, 175(3):736–750, 10 2018. doi: 10.1016/j.cell.2018.08.066.

- M. Häggglund, M. Mørreaunet, M.-B. Moser, and E. I. Moser. Grid-Cell Distortion along Geometric Borders. *Current Biology*, 3 2019. doi: 10.1016/J.CUB.2019.01.074.
- K. Hardcastle, S. Ganguli, and L. Giocomo. Environmental Boundaries as an Error Correction Mechanism for Grid Cells. *Neuron*, 86(3):827–839, 2015. doi: 10.1016/j.neuron.2015.03.039.
- K. Hardcastle, N. Maheswaranathan, S. Ganguli, L. M. Giocomo, K. Hardcastle, N. Maheswaranathan, S. Ganguli, and L. M. Giocomo. Article A Multiplexed , Heterogeneous , and Adaptive Code for Navigation in Medial Entorhinal Cortex Article A Multiplexed , Heterogeneous , and Adaptive Code for Navigation in Medial Entorhinal Cortex. *Neuron*, 94(2):375–387, 2017. doi: 10.1016/j.neuron.2017.03.025.
- A. V. M. Herz, A. Mathis, and M. Stemmler. Periodic population codes: From a single circular variable to higher dimensions, multiple nested scales, and conceptual spaces. *Current Opinion in Neurobiology*, 46:99–108, 10 2017. doi: 10.1016/J.CONB.2017.07.005.
- J. Heys, K. Rangarajan, and D. Dombeck. The Functional Micro-organization of Grid Cells Revealed by Cellular-Resolution Imaging. *Neuron*, 84(5):1079–1090, 11 2014. doi: 10.1016/j.neuron.2014.10.048.
- Z. Huhn, Z. Somogyvári, T. Kiss, and P. Érdi. Distance coding strategies based on the entorhinal grid cell system. *Neural Networks*, 22(5-6):536–543, 2009. doi: 10.1016/j.neunet.2009.06.029.

- R. Ismakov, O. Barak, K. Jeffery, and D. Derdikman. Grid Cells Encode Local Positional Information. *Current Biology*, 27(15), 2017. doi: 10.1016/j.cub.2017.06.034.
- J. Jackson and A. D. Redish. Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus*, 17(12):1209–1229, 2007.
- K. Jezek, E. J. Henriksen, A. Treves, E. I. Moser, and M.-B. Moser. Theta-paced flickering between place-cell maps in the hippocampus. *Nature*, 478(7368):246–249, 10 2011. doi: 10.1038/nature10439.
- A. T. Keinath, R. A. Epstein, and V. Balasubramanian. Environmental deformations dynamically shift the grid cell spatial metric. *eLife*, 7, 10 2018. doi: 10.7554/eLife.38169.
- E. Kelemen and A. A. Fenton. Coordinating different representations in the hippocampus. *Neurobiology of Learning and Memory*, 129:50–59, 2016. doi: 10.1016/j.nlm.2015.12.011.
- C. G. Kentros, N. T. Agnihotri, S. Streater, R. D. Hawkins, and E. R. Kandel. Increased Attention to Spatial Context Increases Both Place Field Stability and Spatial Memory. *Neuron*, 42(2):283–295, 4 2004. doi: 10.1016/S0896-6273(04)00192-8.
- N. R. Kinsky, D. W. Sullivan, W. Mau, M. E. Hasselmo, and H. B. Eichenbaum. Hippocampal Place Fields Maintain a Coherent and Flexible Map across Long Timescales. *Current Biology*, 28(22):3578–3588, 2018. doi: 10.1016/j.cub.2018.09.037.

- B. Kraus, M. Brandon, R. Robinson, M. Connerney, M. Hasselmo, and H. Eichenbaum. During Running in Place, Grid Cells Integrate Elapsed Time and Distance Run. *Neuron*, 88(3):578–589, 11 2015. doi: 10.1016/J.NEURON.2015.09.031.
- J. Krupic, M. Bauza, S. Burton, and J. O’Keefe. Local transformations of the hippocampal cognitive map. *Science*, 359(6380), 2018. doi: 10.1126/science.aao4960.
- D. Lambert. Zero-Inflated Poisson Regression, With an Application To Defects in Manufacturing. *Technometrics*, 34(1):1–14, 1992. doi: 10.1080/00401706.1992.10485228.
- P. Latuske, O. Toader, and K. Allen. Interspike Intervals Reveal Functionally Distinct Cell Populations in the Medial Entorhinal Cortex. *The Journal of Neuroscience*, 35(31):10963–10976, 2015. doi: 10.1523/JNEUROSCI.0276-15.2015.
- I. Lee and J. J. Knierim. The relationship between the field-shifting phenomenon and representational coherence of place cells in CA1 and CA3 in a cue-altered environment. *Learning & memory (Cold Spring Harbor, N.Y.)*, 14(11):807–815, 2007. doi: 10.1101/lm.706207.
- R. J. Low, Y. Gu, and D. W. Tank. Cellular resolution optical access to brain regions in fissures: Imaging medial prefrontal cortex and grid cells in entorhinal cortex. *Proceedings of the National Academy of Sciences*, page 201421753, 12 2014. doi: 10.1073/pnas.1421753111.

- E. Marozzi, L. L. Ginzberg, a. Alenda, and K. J. Jeffery. Purely Translational Realignment in Grid Cell Firing Patterns Following Nonmetric Context Change. *Cerebral Cortex*, pages 1–9, 2015. doi: 10.1093/cercor/bhv120.
- A. Mathis, A. V. M. Herz, and M. B. Stemmler. Multi-Scale Codes in the Nervous System: The Problem of Noise Correlations and the Ambiguity of Periodic Scales. *Physical Review E*, 88(2):022713, 5 2013. doi: 10.1103/PhysRevE.88.022713.
- W. Mau, D. W. Sullivan, N. R. Kinsky, M. E. Hasselmo, M. W. Howard, and H. Eichenbaum. The Same Hippocampal CA1 Population Simultaneously Codes Temporal Information over Multiple Timescales. *Current Biology*, 28(10):1499–1508, 5 2018. doi: 10.1016/J.CUB.2018.03.051.
- M. R. Mehta and B. L. McNaughton. Expansion and Shift of Hippocampal Place Fields: Evidence for Synaptic Potentiation during Behavior. In *Computational Neuroscience*, pages 741–745. Springer US, Boston, MA, 1997. doi: 10.1007/978-1-4757-9800-5_115.
- M. M. Monsalve-Mercado and C. Leibold. Hippocampal Spike-Timing Correlations Lead to Hexagonal Grid Fields. *Physical Review Letters*, 038101 (1):1–4, 2017.
- I. A. Muzzio, C. Kentros, and E. Kandel. What is remembered? Role of attention on the encoding and retrieval of hippocampal representations. *The Journal of Physiology*, 587(12):2837–2854, 6 2009. doi: 10.1113/jphysiol.2009.172445.

- J. A. Pérez-Escobar, O. Kornienko, P. Latuske, L. Kohler, and K. Allen. Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. *eLife*, 5(JULY):1–21, 2016. doi: 10.7554/eLife.16937.
- M. J. Prerau, P. A. Lipton, H. B. Eichenbaum, and U. T. Eden. Characterizing context-dependent differential firing activity in the hippocampus and entorhinal cortex. *Hippocampus*, 24(4):476–492, 2014. doi: 10.1002/hipo.22243.
- M. Pröll, S. Häusler, and A. V. M. Herz. Grid-cell activity on linear tracks indicates purely translational remapping of 2D firing patterns at movement turning points. *The Journal of neuroscience*, 38(31):7004–7011, 7 2018. doi: 10.1523/JNEUROSCI.0413-18.2018.
- P. Puig and J. Valero. Count Data Distributions. *Journal of the American Statistical Association*, 101(473):332–340, 3 2006. doi: 10.1198/016214505000000718.
- A. Renart and C. K. Machens. Variability in neural activity and behavior. *Current Opinion in Neurobiology*, 25:211–220, 2014. doi: 10.1016/j.conb.2014.02.013.
- E. T. Rolls, S. M. Stringer, and T. Elliot. Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. *Network (Bristol, England)*, 17(4):447–65, 12 2006. doi: 10.1080/09548980601064846.
- E. D. Roth, X. Yu, G. Rao, and J. J. Knierim. Functional differences in the

backward shifts of ca1 and ca3 place fields in novel and familiar environments. *PLoS ONE*, 7(4):1–10, 2012. doi: 10.1371/journal.pone.0036035.

F. Sargolini, M. Fyhn, T. Hafting, B. L. McNaughton, M. P. Witter, M.-B. Moser, and E. I. Moser. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–62, 5 2006. doi: 10.1126/science.1125572.

D. L. Solomon. A Note on the Non-equivalence of the Neyman-Pearson and Generalized Likelihood Ratio Tests for Testing a Simple Null versus a Simple Alternative Hypothesis. *The American Statistician*, 29(2):101–102, 5 1975. doi: 10.1080/00031305.1975.10477383.

T. Solstad, E. I. Moser, and G. T. Einevoll. From Grid Cells to Place Cells : A Mathematical Model. *Hippocampus*, 1031:1026–1031, 2006. doi: 10.1002/hipo.

H. Stensola, T. Stensola, T. Solstad, K. Frøland, M.-b. Moser, and E. I. Moser. The entorhinal grid map is discretized. *Nature*, 492(7427):72–78, 2012. doi: 10.1038/nature11649.

G. Tocker, O. Barak, and D. Derdikman. Grid cells correlation structure suggests organized feedforward projections into superficial layers of the medial entorhinal cortex. *Hippocampus*, 00:n/a–n/a, 2015. doi: 10.1002/hipo.22481.

A. Tsao, J. Sugar, L. Lu, C. Wang, J. J. Knierim, M.-B. Moser, and E. I. Moser. Integrating time from experience in the lateral entorhinal cortex. *Nature*, 561(7721):57–62, 9 2018. doi: 10.1038/1586-018-0459-6.

- S. Weiss, G. Talhami, X. Gofman-Regev, S. Rapoport, D. Eilam, and D. Derdikman. Consistency of Spatial Representations in Rat Entorhinal Cortex Predicts Performance in a Reorientation Task. *Current Biology*, 27(23):3658–3665, 12 2017. doi: 10.1016/J.CUB.2017.10.015.
- S. S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, 3 1938. doi: 10.1214/aoms/1177732360.
- K. Yoon, M. A. Buice, C. Barry, R. Hayman, N. Burgess, and I. R. Fiete. Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature Neuroscience*, 16(8):1077–1084, 2013. doi: 10.1038/nn.3450.
- Y. Ziv, L. D. Burns, E. D. Cocker, E. O. Hamel, K. K. Ghosh, L. J. Kitch, A. E. Gamal, and M. J. Schnitzer. Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, 16(3):264–266, 3 2013. doi: 10.1038/nn.3329.

Spike afterpotentials shape the burst activity of grid cells during navigation

Caroline Fischer*, Dóra É Csordás*, Johannes Nagele*, Martin Stemmler and Andreas V.M. Herz

Bernstein Center for Computational Neuroscience Munich and
Faculty of Biology, Ludwig-Maximilians-Universität München
Großhaderner Straße 2, 82152 Martinsried-Planegg, Germany

*equal contribution

Abstract

Grid cells in the medial entorhinal cortex (MEC) of navigating rodents encode the animal's spatial environment based on lattices of hexagonally arranged firing fields. Their spike trains are organized on multiple time scales and include high-frequency bursts in the 150-300 Hz range. A mechanistic understanding of these burst sequences is, however, largely missing. In this study, we reanalyzed whole-cell recordings from male mice running in a virtual corridor (Domnisoru et al., 2013) and tetrode data obtained during movements in a real two-dimensional arena (Latuske et al., 2015). The membrane potentials of some grid cells recorded in virtual reality showed depolarizing afterpotentials (DAPs) known from in-vitro studies of MEC principal neurons. All such cells were located in Layer II, generated bursts, and their inter-spike intervals (ISIs) were typically between 5 and 15 milliseconds. The ISI distributions of all other Layer-II cells peaked sharply at ~4.1 milliseconds and varied only minimally across that group (standard deviation: 0.1ms). This dichotomy in burst behavior is explained by cell-group-specific dynamics of spike afterpotentials. Layer III neurons were only sparsely bursting and had no DAPs. The same two classes of bursting neurons also emerged when clustering extracellular spike train autocorrelations. Yet no difference in the spatial coding properties of cells within these two classes was discernible. As the ion-channels underlying DAPs can be modulated in various ways, our results suggest that temporal features of grid-cell activity can be altered to serve different functions without affecting the cells' spatial tuning characteristics.

This work was supported by the German Federal Ministry for Education and Research Grant 01GQ0440. We thank D.W. Tank and K. Allen for making data from Domnisoru et al. (2013) and Latuske et al. (2015), respectively, available; and S.Häusler for stimulating discussions.

Author contributions: All authors designed research. C.F., D.C., and J.N. performed research and analyzed data (equal contribution). A.V.M.H. wrote and edited the paper with support from M.S. and the other authors.

Introduction

The firing fields of grid cells in the medial entorhinal cortex (MEC) form hexagonal lattices spanning the explored environment (Hafting et al. 2005). Lattice scales follow a geometric progression so that discrete grid-cell modules emerge. Within each module, lattices are co-aligned and spatially phase shifted (Stensola et al., 2012). The detailed grid layout reflects the environment's shape, contextual information and goal locations (Krupic et al., 2015; Diehl et al., 2017; Boccara et al., 2019; Giocomo et al., 2019) but overall, the spatial coding properties of grid cells are surprisingly robust.

In contrast to this spatial stereotypy, grid cells show rich temporal behavior, from slow depolarization ramps (Domnisoru et al., 2013, Schmidt-Hieber and Häusser, 2013), theta-band (7-12Hz) spike locking and phase precession (Hafting et al., 2008; Reifenstein et al., 2012), to gamma-band activity (Chrobak and Buzsáki, 1998; Colgin et al., 2009) and burst sequences with instantaneous firing rates of up to 300 Hz (Latuske et al., 2015). Notably, not every grid cell participates in all these phenomena. In particular, there are two distinct subclasses of grid cells, those that burst frequently and those that do not or only rarely generate bursts (Latuske et al., 2015).

These observations raise various questions: Can grid-cell burst phenomena be explained by intrinsic membrane processes or are network interactions required? Are there distinct burst classes or should one rather think of one continuum? Is burst activity cell-type specific? Is it modulated in vivo? Why do grid-cell spatial coding properties not vary with burst propensity?

In this study, we tested the hypothesis that grid-cell bursting is shaped by cell-intrinsic membrane-potential dynamics. Two mechanisms come to mind. Bursts could be a resonance phenomenon reminiscent of action-potential (AP) sequences riding on theta-band membrane-potential oscillations in MEC Layer-II stellate cells (Alonso and Klink, 1993; Engel et al. 2008; see also Hasselmo, 2013; Newman and Hasselmo, 2104). This would, however, require strong electrical resonances far above the gamma band, which have not been reported. Alternatively, bursts could result from AP-triggered processes that increase the probability of further discharges.

Indeed, slice experiments have shown that depolarizing afterpotentials (DAPs) arise in a majority of principle cells in superficial MEC layers (Alonso and Klink, 1993; Canto and Witter, 2012). DAPs are at the center of triphasic deflections following an AP, sandwiched between fast and medium after-hyperpolarization (fAHP and mAHP). The DAP maximum occurs some five-to-ten milliseconds after the AP and peaks a few millivolts above the fAHP minimum. In stellate cells, DAPs become more pronounced when neurons are hyperpolarized, whereas the reverse is true for pyramidal neurons (Alessi et al., 2016). Not all cell types associated with spatial grid firing rate maps have DAPs in vitro, however. In particular, layer-III neurons are reported to have no DAPs (Canto and Witter, 2012).

DAPs do not only agree in their relevant time scale with intra-burst inter-spike intervals (ISIs), DAPs also play a causal role for bursting in vitro. Alessi et al. (2016) reported that during DAPs the AP current threshold was reduced such that the cells' average excitability increased by over 40%. Conversely, neurons without strong DAPs did not burst at the beginning of an AP train (Canto and Witter, 2012).

To test the functional relevance of DAPs under in-vivo conditions, we reanalyzed whole-cell recordings from mice moving on a linear track (Domnisoru et al., 2013) and could show that DAPs play a decisive role for burst firing in MEC Layer-II neurons: Cells with DAP were bursty and their intra-burst ISIs were compatible with the DAP mechanism. ISI distributions of the other Layer-II cells were highly uniform and had a sharp peak at 4.1 ± 0.1 ms (SD across

this cell group). All remaining neurons were sparsely bursting and those with known location resided in Layer III. These results are compatible with reanalyzed extracellular recordings from open-field arenas (Latuske et al., 2015); there, the two bursty cell groups did not differ in their spatial coding properties. As the ion-channels underlying DAPs can be modulated in many ways, these findings suggest that temporal features of grid-cell activity can be altered to serve different functions without affecting the cells' spatial tuning characteristics.

Materials and Methods

Data. We reanalyzed data from two separate grid-cell studies in navigating wild-type (C57BL/6) male mice. The data set "D" (Domnisoru et al., 2013) contained voltage traces (Sampling frequency: 20kHz) from whole-cell recordings in head-fixed animals running on cylindrical treadmills embedded in virtual corridors. The data set "L" (Latuske et al., 2015) contained tetrode data (Sampling frequency: 20 or 24 kHz) obtained during movements in a real square arena (70 x 70 cm).

Grid cell selection. As a first step, the same grid-cell criteria as in the original publications were used. Data set D: The original data set contained recordings from 27 cells. One recording (cell s066) was partially corrupted and therefore excluded. Two cells (s081 and s115_30) had mean firing rates above 10Hz and were removed to allow for an unbiased comparison with data set L, which contained only cells with firing rates below 10Hz to exclude interneurons. Data set L: The original data set contained recordings from 115 grid cells. To avoid cluster artefacts, 11 cells were removed whose spike-time autocorrelations had no data in more than 90% of the 1ms-bins.

Spike-train characterization. The firing rate of a cell was defined as number of spikes divided by the total duration of the recording. For graphical illustrations, spike-time autocorrelations and inter-spike interval (ISI) distributions were derived from binned data (bin width: 1ms). To calculate the peak location and width of ISI distributions, the recorded time difference between each pair of successive spikes was represented by a Gaussian kernel with a standard deviation of 1ms. These individual kernel density (KD) estimates were summed up across the entire recording. The analogous procedure was used for autocorrelations.

The location of the ISI peak was determined as the inter-spike interval for which the KD estimate was maximal. Similarly, the width of the ISI distribution was defined as full width at half maximum. The mean ISI and its standard deviation was calculated from all ISIs, the coefficient of variation (CV) was defined as the ratio between standard deviation and mean.

A burst was defined as a sequence of at least two spikes with ISIs shorter than 8ms. The fraction of ISIs smaller than 8ms was calculated relative to all ISIs below 200ms and serves as a measure for the cell's burstiness. The fraction of ISIs between 8ms and 25ms was computed in the same way. An event is a burst or an isolated spike. The fraction of single spikes was defined as the number of spikes that do not belong to a burst divided by the number of events.

Principal component (PC) analysis. For both data sets, autocorrelations were calculated for time lags τ_{\max} up to 50ms. For data set D the two cells that were excluded when the principal components were computed are nevertheless shown in the PC plot. To test the robustness of the PC analysis of the D data, the maximal time lag τ_{\max} was varied between 30ms and 100ms (see also Results).

Identification of neuron classes. For the D data set, visual inspection of the two-dimensional space spanned by the first two PCs suggested two main cell groups, whose arrangement was determined by k-means clustering with $k=2$ clusters (scikit-learn). To test the robustness of the

k-means clustering for the L data set, cluster analyses were performed on the 50-dimensional raw autocorrelations as well as in PC spaces with $N=2-4$ dimensions. The reliability of clustering was estimated using silhouette scores (Rousseeuw et al., 1987). To preserve the data's diversity and to minimize the risk of falsely assigning cluster structure to continuous data distributions, no data were excluded by using, e.g., linear discriminant analysis.

Membrane-potential dynamics. The whole-cell voltage traces contained sizeable fluctuations that reflected synaptic inputs and potential movement artefacts. To obtain reliable information about the membrane potential before and after an action potential, AP-triggered averaging had to be performed. The APs themselves varied in amplitude and width, both within and across the different recordings, suggesting that the recording quality fluctuated in time; the slowly decaying AP amplitudes of some cells indicated run-down effects. To guarantee a good recording quality and obtain reliable estimates of the subthreshold membrane-potential dynamics on the time scales relevant for fAHPs and DAPs, we studied well isolated APs (no further APs within 25ms before and after the trigger AP), and required the individual AP amplitudes to be larger than 40mV (measured relative to the membrane potential 10ms before the AP maximum) and APs width to be smaller than 1ms.

The pre-AP voltage slope was calculated from the cell's average AP-triggered voltage trace within the last 10ms before AP onset; AP onset was determined by a threshold crossing (15mV/ms) in the average AP-triggered voltage trace.

For cells with DAPs, the fAHP amplitude ΔV_{fAHP} was defined as the average voltage minimum during the fAHP relative to the voltage at AP onset. This means that ΔV_{fAHP} is negative for DAP cells (see Fig. 1). The DAP-deflection ΔV_{DAP} was defined as difference between the voltage level at the DAP peak and at the minimum of the preceding fAHP. It is positive for cells with DAPs. The time interval between the AP peak and the following fAHP minimum is denoted by Δt_{fAHP} , the time interval between the AP peak and the following DAP maximum is called Δt_{DAP} .

To compare the afterpotentials of different neurons, the definitions of ΔV_{fAHP} and ΔV_{DAP} had to be generalized to cells without DAP. To this end, we calculated the population averages and across all cells with clearly visible DAPs ($n=7$). We then used these mean time intervals, $= 1.8$ ms, $= 4.6$ ms, to determine voltage changes corresponding to ΔV_{fAHP} and ΔV_{DAP} values for cells without DAP. These are called ΔV_1 and ΔV_2 , respectively. For consistency, we also used these definitions for cells with DAP in Fig. 2A, so that the true fAHP and DAP amplitudes are underestimated.

Spatial coding properties. For data set L, grid score and head-direction score were calculated as proposed by Sargolini et al. (2006), the spatial information as in Skaggs et al. (1996).

Experimental design and statistical analysis. We reanalyzed data originally recorded by Domnisoru et al. (2013) and Latuske et al. (2015) and refer the reader to these two publications for details on the experimental design. All our analyses were performed in Python 2.7.6. Specific statistical tests used are stated throughout the text. The Kruskal-Wallis test, the Kolmogorov-Smirnov test, the Chi-square test and the median tests are taken from scipy.stats. The linear regression, the principal component analysis and the k-means clustering are taken from scikit-learn.

Bootstrapping. To assess the fAHP and DAP parameters, we bootstrapped the AP-triggered voltage traces of a cell by using sampling with replacement and repeated this procedure 10000 times to obtain mean values and standard errors.

Results

The temporal firing characteristics of grid cells in the medial entorhinal cortex (MEC) of behaving rodents vary strongly from cell to cell, even if their mean firing rates are almost identical (Latuske et al., 2015). Some neurons rarely fire with inter-spike intervals shorter than 8ms; their spike-time autocorrelations have a pronounced dip at short time lags (Fig. 1A, left). Other cells show an autocorrelation peak in the 5-15 ms range with broad flanks (Fig. 1A, middle) and yet other grid cells have autocorrelations that are sharply peaked at even shorter lags (Fig. 1A, right). The first group of neurons have been termed "non-bursty" by Latuske et al. (2015), the other neurons were summarized as "bursty" neurons by these authors. Since even "non-bursty" neurons generate bursts from time to time we will call them "sparsely bursting" (Simmonet and Brecht, 2019).

We wondered whether differences in the *in vivo* spike patterns of bursty neurons could be explained at a mechanistic level by differences in their single-cell dynamics and whether differences in the cells' temporal discharge patterns were reflected in their specific spatial tuning properties. To this end, we reanalyzed whole-cell recordings from mice moving on a linear track in virtual reality (Domnisoru et al., 2013) and extracellular recordings from mice navigating in two-dimensional environments (Latuske et al., 2015).

Grid cells differ in the voltage deflections following an action potential.

We first focused on the intracellular linear-track data as these provide information about both, spike times and membrane-potential dynamics.

The time courses of the membrane potentials recorded by Domnisoru et al. (2013) show striking cell-to-cell differences within the first ten milliseconds following an action potential (AP). Three types of behavior can be distinguished from the spike-triggered voltage traces:

- (a) a monotone repolarization that is gradually slowing down (Fig. 1B, left panel)
- (b) a fast hyperpolarization (fAHP) followed by a depolarizing afterpotential (DAP), as shown in the middle panel,
- (c) a short repolarizing phase that abruptly turns into a much slower voltage decay, which may include a flat shoulder (right panel).

To quantify these distinct behaviors, we used parameters that capture the two salient features of cells exhibiting DAPs – the voltage minimum during the fAHP and the voltage peak during the DAP (see inset of Fig. 1C). The "fAHP-depth" ΔV_{fAHP} measures the voltage minimum relative to the membrane potential at AP onset. This minimum occurs at some time Δt_{fAHP} after the AP peak. The "DAP-deflection" ΔV_{DAP} measures the difference between the membrane potential at the DAP peak and the fAHP minimum. The DAP peak is attained at some time Δt_{DAP} after the AP peak (Fig. 2A)

To extend the ΔV_{fAHP} and ΔV_{DAP} measures to voltage traces of cells with no detectable DAP, two time intervals corresponding to Δt_{fAHP} and Δt_{DAP} need to be defined. For concreteness, we used the population means across all cells with a clearly visible DAP ($n=7$), resulting in $\Delta t_{\text{fAHP}} = 1.8$ ms and $\Delta t_{\text{DAP}} = 4.6$ ms. We then determined the voltage differences corresponding to V_{fAHP} and ΔV_{DAP} at these two time points, and named them ΔV_1 and ΔV_2 , respectively. This means that cells with a shoulder or slow voltage decay (Fig. 1B, right panel) have zero or small ΔV_2 irrespective of their ΔV_1 value whereas large negative ΔV_2 values indicate a strong decline in membrane potential until around five milliseconds after the action potential. To treat all neurons on equal footing, we used ΔV_1 and ΔV_2 for cells with a DAP, too. This implies that the ΔV_1

and ΔV_2 values shown for these cells in Fig. 1C underestimate the full fAHP and DAP amplitudes.

Within the parameter space spanned by ΔV_1 and ΔV_2 (Fig. 1C), neurons fall into two distinct groups – cells with a pronounced DAP (negative ΔV_1 and positive ΔV_2) and cells with no detectable DAP (negative ΔV_2), which typically have also no fAHP (positive ΔV_1). All neurons with a DAP are located in Layer II, and five out of these seven cells show large theta-band oscillations in their membrane potential. They have been termed "large theta cells" by Domnisoru et al. (2013). One cell, whose type and anatomical location is not known, has an only slightly negative ΔV_2 (-0.1mv) but strong fAHP and is therefore included in the group of DAP neurons. Cells without a visible DAP (n=16) are located in Layers II and III and are also more diverse in their anatomy (four pyramidal, three stellate, nine unidentified cells).

Grid cells differ in their spike-train characteristics.

To resolve the high diversity within the non-DAP cells and to understand the relation between this group and the DAP-cells, we analyzed the discharge patterns of all 24 grid cells. Aiming at an informative low-dimensional representation, we carried out a principal component analysis of their spike-time autocorrelations (Fig. 2A). We found that for a maximal time lag τ_{\max} of 50ms, the first two principal components, PC1 and PC2, explain 66% and 18% of the cell-to-cell variability, respectively, whereas the contribution from PC3 adds only another 4%. Together, PC1 and PC2 thus account for 84% of the variability. This value changes by less than 3% if τ_{\max} is varied between 30 ms and 100 ms (data not shown) and starts to decrease for shorter or longer maximal lag. These findings suggest that a two-dimensional PC representation of the grid-cell autocorrelations in the 0-50 ms range describes the essence of the cell-to-cell variability in a reliable manner.

The mean autocorrelation is highly peaked at a lag τ of around 4 ms (Fig. 2B), and so are both principal components (Fig. 2C). This indicates that brief activity bursts in the 250 Hz range play an important role for both the mean grid-cell discharge patterns as well as their cell-to-cell variability.

Within this two-dimensional representation (Fig. 2A), neurons without a DAP have a negative or only small positive second principal component and strongly vary in their PC1. Cells with negative PC1 are sparsely bursting as the example in the left panel of Fig. 1A, marked with a blue arrow in Fig. 2A. Cells with positive PC1 are bursting as the cell in the right panel of Fig. 1A (yellow arrow in Fig. 2A). Cells with a DAP have positive PC2, only a small PC1, and are also bursting, though with a much broader peak in their autocorrelation (see the example in middle panel in Fig. 1A, marked with a red arrow in Fig. 2A).

This grouping based on intracellular information and visual inspection of the AC principal components might not properly distinguish between bursting and sparsely-bursting neurons with small PC1. To better discriminate between these two larger cell groups, we carried out a k-means clustering with k=2. This analysis suggested that 9 cells should be classified as sparsely bursting ("SB") neurons; based on their intracellular characteristics, the remaining 15 bursting ("B") cells are either DAP cells ("BD⁺") or cells without detectable DAP ("BD⁻"). The same clusters emerge if the spike data from the first and second half of each experiment are treated separately, and provide further evidence for the robustness of our approach.

All bursty neurons whose anatomical position was classified by Domnisoru et al. are located in Layer II, none in Layer III (two bursty cells were not assigned to a layer). Furthermore, bursty neurons are more likely to be stellate than pyramidal cells (6 versus 2 cells), in agreement with the larger abundance of stellate cells compared to pyramidal cells (Alonso

and Klink, 1993). From an anatomical point of view, there is no difference between bursty neurons with and without DAP: Within the BD^+ group ($n=8$), there are three stellate cells, one pyramidal neuron and four non-identified cells. Within the BD^- group ($n=7$), there are three stellate cells, one pyramidal neuron and three non-identified cells. In contrast, not a single sparsely-bursty cell was identified as a stellate cell (pyramidal and non-identified cells: 3 and 6 out of $n=9$, respectively) and non-bursty neurons tend to reside in Layer III (three cells versus one cell in Layer II; 5 cells were not classified). Finally, in the ΔV_1 - ΔV_2 representation, BD^- neurons overlap with SB cells, but tend to have less negative ΔV_2 values and the ΔV_1 and ΔV_2 values within the group of bursty neurons are correlated with a slope of -0.49 (standard error: 0.06). The three groupings are robust, as confirmed by bootstrapping and indicated by error bars in Fig. S1.

Post-AP dynamics explain the spike-train characteristics of bursty neurons.

Based on the cell classification made possible by the DAP measurements (without which the substructure of the autocorrelation data from bursty neurons could not have been resolved) we are now in the position to compare the cells' intracellular and spike-train characteristics at the group level (Fig. 3).

Pooling the membrane-potential traces within each of the three groups (Fig. 3A) confirms the impression gained from individual cells. Sparsely-bursting neurons show a smooth and monotone AP down-stroke (Fig. 3A, left panel), bursty cells with DAP exhibit a local voltage minimum followed by repolarization (Fig. 3A, middle panel), and bursty cells without DAP have a more pronounced kink separating the initial AP downstroke from the following repolarization (Fig. 3A, right panel), whose first phase is less steep than that of sparsely-bursting neurons.

To visualize these distinct features, the spike-triggered voltage traces were averaged for each neuron and then aligned to each cell's mean voltage at AP onset. Without such voltage alignment (see insets), differences in membrane potential at AP onset are apparent (SB: -59.42 \pm 1.22 mV, BD^+ : -57.12 \pm 3.09 mV, BD^- : -61.40 \pm 6.83 mV) ($p(BD^+, BD^-)=0.25$; $p(B, SB)=0.39$; $p(BD^+, SB)=0.15$; $p(BD^-, SB)=0.96$, Kruskal Wallis, no correction applied), as well as a somewhat larger variability of the afterpotentials within the two groups of bursty neurons. The voltage slope during the last 10ms before AP onset does not differ significantly between BD^+ and BD^- neurons (BD^+ : 0.50 \pm 0.03 mV/ms; BD^- : 0.48 \pm 0.08 mV/ms; $p=0.35$, Kruskal Wallis) but does so when sparsely-bursting and bursty neurons are compared (B: 0.49 \pm 0.06 mV/ms; SB: 0.34 \pm 0.05 mV/ms; $p=0.00015$, Kruskal Wallis). Finally, visual inspections suggests that there are no membrane-potential oscillations in the > 150 Hz regime that would be expected if the bursts resulted from electric resonances.

The averaged autocorrelations (Fig. 3B, left panel) and interspike intervals (Fig. 3C, left panel) of sparsely-bursting cells confirm that although these neurons rarely generate spike sequences with short ISIs – only 2% of all their ISIs are less than 8ms – if they do fire such bursts, however, the most probable ISI below 8ms is 4.30 \pm 0.81ms long (see the black arrows in Figs. 3A and B). For the majority of neurons, the intra-burst spike-count distributions are compatible with exponential decays and reveal that there is no preferred burst size or "unit of information", such as a spike doublet or triplet. In the spirit of defining hippocampal "complex spike bursts" (Ranck, 1973), a burst will be a sequence of two or more spikes with inter-spike intervals less than 8ms. The exact choice of cutoff threshold was not critical; qualitatively similar results were obtained using ISI thresholds up to 15ms (data not shown).

Both types of bursty neurons exhibit prominent ISI- and autocorrelation peaks at short time scales (Fig. 3B, C, middle and right panels). Population averages within each group show that the most likely ISI of cells without a DAP is significantly shorter than that of cells with a DAP (4.12 ± 0.12 ms vs. 6.96 ± 3.73 ms, $p=0.01$, Kruskal Wallis); the same is true for the autocorrelation peaks (4.13 ± 0.11 ms vs. 9.46 ± 4.41 ms, $p=0.001$, Kruskal Wallis). These differences are readily explained by the different time courses of the post-spike voltage deflections: The rapid fAHP time course of BD^+ cells strongly reduces the chance that a second AP is fired directly after the first AP, whereas in BD^- cells the down-stroke of the first AP stops abruptly at depolarized levels, often above the AP threshold (see Fig. 2A), resulting in a rather short absolute refractory period (mean of 10 % shortest ISI in BD^- cells: 3.45 ms versus 5.24 ms in BD^+ cells. Consistent with this picture, the DAP opens a wide "window of opportunity" for a second AP, resulting in broader ISI distributions (6.7 ± 3.27 ms vs. 3.66 ± 0.36 ms, $p=0.01491$, Kruskal Wallis) and autocorrelation functions (13.26 ± 7.34 ms vs. 10.38 ± 11.47 ms, $p=0.24716$, Kruskal Wallis) for BD^+ cells, compared to BD^- cells. Finally, a direct role of the post-AP dynamics in burst behavior is also suggested by the observation that for BD^+ cells, the most likely ISI mirrors Δt_{DAP} , the time interval between an AP and the succeeding DAP peak (no difference of median values, $p=0.61$, median=5.13 ms, median test).

These considerations do, however, not explain why sparsely-bursting cells have much broader ISI distributions and autocorrelation functions although their post-AP membrane-potential dynamics is only slightly steeper than that of BD^- cells (see Fig. 2C). This observation indicates that besides the intrinsic voltage dynamics, differences in cell-type (all identified SB cells are pyramidal neurons whereas only one of four BD^- cells is pyramidal) and anatomical position (four of five classified SB cells are in LIII, all BD^- cells are in LII) play an important role, too. In contrast, the cell-type composition and anatomical location of the BD^- and BD^+ groups are almost identical and cannot explain the observed differences in their spike trains.

Finally, the probability to see a certain number n of spikes per burst decays monotonically for all three cell groups, without any special role for spike doublets or triplets (Fig. 3D). More precisely, in 21 out of the 24 cells, the decay is consistent with an exponential distribution (linear fit in log. space; chi-square test for the correlation of the fit with the data; $p>0.05$).

Spike-train characteristics of bursty cells are conserved across experimental conditions.

So far, the analysis was based on a relatively small number of neurons recorded in head-fixed animals running in a virtual linear corridor (Domnisoru et al., 2013). To understand whether our conclusions generalize to other experimental conditions, we reanalyzed a complementary data set with 104 grid cells from mice that randomly foraged in a square environment (Latuske et al., 2015). Although these extracellular recordings do not offer direct access to the membrane-potential dynamics, they may still reveal signatures of the different post-AP dynamics. In particular, we expected that grid cells may not only show the bursty vs. non-bursty dichotomy revealed by Latuske et al. (2015), but that there would also be qualitative differences within the bursty subpopulation.

To facilitate the comparison between the two data sets we again focused on spike-time autocorrelations with a maximum lag of 50 ms and 1 ms-binning. To minimize any observer bias, k-means cluster analyses were performed on the 50-dimensional raw autocorrelations as well as in principal-component spaces with $N=2-4$ dimensions. To analyze the reliability of the k-means cluster algorithm for different k values, we calculated silhouette scores (Rousseeuw et al., 1987). Irrespective of the dimensionality of the data, separations into three clusters led to the best performance and the clusters resulting for different N were almost

identical. The same was true when the cluster analysis was based on autocorrelations computed from the first or second half of spike events for each cell: 93.3% of the cells kept their cluster identity (Fig. S2). For concreteness, we present cluster results for $N=3$. For comparison with the data from Domnisoru et al., we again plot the first two principal components against each other (Fig. 4A). The mean autocorrelation (Fig. 4B) closely resembles that from the virtual-track data (see Fig. 2B); the same applies to the principal components (cf. Fig. 4C and Fig. 2C).

There is also a high similarity between the corresponding group-averaged autocorrelations in both data sets (cf. Fig. 4D and Fig. 2B). This is remarkable as the cluster analysis of the open-field data (Latuske et al., 2015) just reflects the overall structure of the grid-cell autocorrelations without any fine tuning or insight from intracellular measurements. There is one important difference between both data sets, however. The sparsely bursting neurons recorded by Latuske et al. (Fig. 4D, left panel) fire hardly any spike within the first few milliseconds (so that the authors named them "non-bursty" neurons) and their autocorrelation has a pronounced peak at around 15ms. The autocorrelation function of the sparsely bursting neurons recorded by Domnisoru et al. (Fig. 3A, left panel) exhibits a local peak at around 4ms, and grows much more slowly, with a local maximum at 30-40ms. On the other hand, the average autocorrelations of the BD^+ and BD^- cells are almost identical under both experimental conditions.

As shown in Fig. 5, the qualitative similarities and differences between the cell groups extends to other spike-train characteristics. Some measures, such as the fraction of inter-spike-intervals below 8ms (Fig. 5B), the most likely ISI (Fig. 5D), or the width of the ISI histogram (Fig. 5E) could have been expected; others, such as the high similarity of mean firing rates across the three cell groups (Fig. 5A), could not have been predicted from the shape of the autocorrelation.

Although some p-values differ strongly between the virtual-track and open-field data (see, e.g., those for the location of the ISI peak), on average, the p-values for comparisons between corresponding cell groups are rather similar. This indicates that the increase in statistical power of the larger open-field data set roughly balances errors caused by the unsupervised cluster algorithm whose class-assignments have not been pruned by, e.g., linear discriminants.

In the virtual-track data, the autocorrelations of the BD^- neurons ($n=7$) peaked at 4.1 ms with a cell-to-cell variability of 0.1 ms (SD). The autocorrelations of the corresponding cells from the open-field recordings ($n=25$) show a peak location at 3.5 ms and a slightly larger standard deviation (0.22 ms). These deviations may reflect differences between the two experimental settings, which might also cause the larger grid-field sizes in virtual reality compared to open-field environments (c.f., Domnisoru et al., 2013, supp. Fig 9), or between the measured intra- and extracellular signals (Anastassiou et al., 2015). In addition, there might be age differences of the tested animals, strain differences between the two mouse lines or other influences.

Spatial response properties are shared across all three cell groups.

In the next step of our analysis, we asked whether the pronounced differences in the temporal response characteristics of the three cell groups translate into differences in their spatial firing patterns. The study of Latuske et al. (2015) had shown that this was *not* the case when one compares bursty with sparsely-bursting grid cells. However, the two groups of bursty neurons might nevertheless differ in their spatial behavior. To obtain reliable field estimates, we used the open-field data for this analysis. We tested various measures, including grid score (Sargolini et al., 2006), spatial information (Skaggs et al., 1996), and head-direction score

(Sargolini et al., 2006), but could not detect any significant differences between BD^+ and BD^- cells (Fig. 5F-H).

Slice experiments show that depolarizing afterpotentials of stellate cells can be modulated; if the holding potential is decreased, the amplitude ΔV_{DAP} of the following DAP increases, and it decreases whenever the holding potential is increased (Alessi et al., 2016). As shown by Domnisoru et al. (2013), a grid cell is depolarized when the animal is located in a firing field of that cell and hyperpolarized in the out-of-field regions. We therefore wondered whether a BD^+ cell might preferentially generate DAP-mediated bursts when one of its grid field is entered, as the membrane-potential ramp might facilitate larger DAPs and thus make DAP-mediated burst firing in these neurons more likely.

To test this hypothesis, we took open-field data from Latuske et al. (2015) and investigated in detail whether spikes belonging to the bursts of a BD^+ cell had an above-chance probability to occur at the edges of its firing fields and, more generally, whether those spikes differed in their spatial statistics from other spikes of the same neuron, in the spirit of a place-cell study by Harris et al. (2001). In particular, we analyzed the distribution of spike distances from the respective firing-field centers as well as topological features of the discharge patterns of BD^+ cells, with special focus on interspike-intervals expected for DAP-triggered bursts. Despite extensive efforts, we could not find any significant differences between the spatial firing characteristics of BD^+ versus BD^- cells. As a complementary check, we used the data from Domnisoru et al. (2013) to search whether DAP deflections measured in the firing fields of BD^+ cells were smaller than the DAP deflections of out-of-field spikes but did not find any obvious changes either.

These unexpected findings suggest that despite the striking differences in the spike-train patterns of bursty cells with and without DAP, these differences have no consequences for the cells' spatial tuning properties. Temporal variations in the membrane potential, in particular the large theta-oscillations observed in some bursty grid cells, are uncorrelated with the animal's trajectory and may easily mask less prominent spatial dependencies.

Such decoupling of spatial and temporal tuning characteristics might endow the system with added plasticity and computational flexibility.

Bursty grid cells: One continuum or two clusters?

Since BD^+ and BD^- cells showed indistinguishable spatial tuning properties, we reconsidered their partition into two distinct groups based on their temporal firing characteristics. Could it be that the data are better described as a single group with continuously varying parameters?

To answer this question, we went back to the data from Domnisoru et al. (2013) and analyzed how the cells' salient spike-train characteristics depended on the two biophysical parameters ΔV_1 and ΔV_2 (Fig. 6). For the mean firing rate (Fig. 6A,B) no correlations with ΔV_1 and ΔV_2 are apparent. As the firing rates do not vary significantly between all three cell groups (Fig. 5A), this might have been expected but there are also no significant trends within each cell group (ΔV_1 : $p_{BD^+}=0.88$, $p_{BD^-}=0.24$; ΔV_2 : $p_{BD^+}=0.58$, $p_{BD^-}=0.90$, as tested by shuffling the x, y values independently for all points and computing the Pearson correlation for each new sample. The p-value is the number of samples for which the correlation value was larger than in the original sample.).

A different picture emerges when the fraction of short ISIs (below 8 ms) is considered (Fig. 6C,D). Visual inspection suggests a joint trend for BD^+ and BD^- cells; the larger ΔV_1 , the more frequent are short ISIs (Fig. 6C), at the expense of larger ISIs (Fig. 6E). With p-

values of 0.02 and 0.04, respectively, both trends are statistically significantly different from the null-hypothesis (no increase/decrease as a function of ΔV_1).

This is in agreement with our earlier functional interpretation of post-AP voltage deflections: For negative ΔV_1 (i.e., BD^+ cells), cells quickly hyperpolarize, making short ISIs rare. As ΔV_1 takes more negative values or as ΔV_2 grows, the likelihood of DAP-associated ISIs (in the range of 8-25 ms) therefore grows, as suggested by Fig. 6E and Fig. 6F (p-value: 0.012). On the other hand, as ΔV_1 approaches zero, the fAHP influence decreases so that short ISIs become more and more likely, at the expense of longer ISIs (cf. Fig. 6C and 6E). Consistent with these trends, the location of the ISI peak tends to grow for increasingly negative ΔV_1 (Fig. 6G) and increasingly positive ΔV_2 (Fig. 6H) if the entire population of bursty neurons is considered (ΔV_1 : $p=0.04$, ΔV_2 : $p=0.01$). Within the BD^- population, however, the ISI peak does hardly vary at all, as emphasized before.

These results indicate that there is no clear-cut answer to the question whether the population of bursty neurons forms one joint though under-sampled cloud or contains two distinct sub-populations. More importantly, however, is the observation that in either case, the spike-train characteristics do depend on the cells' individual DAP properties, which supports the view that DAPs do not only exist under in-vivo conditions but also play a functional rule.

Discussion

Tetrode recordings in freely moving rats (Mizuseki et al., 2009, Ebbersen et al., 2016) and mice (Latuske et al. (2015) have shown that principal neurons in superficial MEC layers come in two functional subclasses, cells that burst frequently and others that do not or only rarely burst. Our reanalysis of whole-cell data from mice running on linear virtual tracks (Domnisoru et al., 2013) suggests that the "sparsely bursting" (SB) grid cells are located in Layer-III and that they do not generate depolarizing afterpotentials (DAPs), in agreement with previous slice studies (Canto and Witter, 2012).

Bursty neurons varied strongly in their burstiness (Fig. 5B) and the overall shape of their inter-spike interval (ISI) distributions and autocorrelations (Fig. 3A). This diversity can be understood in terms of the cell-specific shapes of spike afterpotentials: Neurons without a DAP (" BD^- cells") had inter-spike interval (ISI) distributions that peaked sharply at around four milliseconds and varied only minimally across that group of cells whereas the ISIs of neurons with a DAP (" BD^+ cells") were most frequent between 5 and 15 ms.

At first sight, the void in the lower left quadrant of the ΔV_1 - ΔV_2 diagram (Fig. 2A) speaks against a continuum of bursty neurons and rather points to the existence of two separate subgroups. This impression might, however, be due to a sampling artefact; there are only 15 bursty grid cells with intracellular recordings in the data set from Domnisoru et al. (2013). We therefore investigated the dependencies of various spike-train characteristics on ΔV_1 and ΔV_2 (Fig. 6). The smooth behavior of some measures, such as the burstiness, i.e. the fraction of ISIs below 8ms, or the ISI-peak location, and the lack of any sharp transitions in the other measures, support the assumption of one single, though under-sampled group of neurons. Although based on small numbers, the equal stellate-to-pyramidal-cell ratio (3:1) of the BD^+ and BD^- subgroups is pointing in the same direction.

Consistent with this alternative hypothesis of one single group of bursty neurons, the physiological properties of individual cells could either be fixed or undergo plastic changes that move the biophysical cell parameters between the BD^+ and BD^- regions. In the ΔV_1 - ΔV_2 space (Fig. 1C), a transition from BD^+ to BD^- corresponds to an increase in ΔV_{fAHP} .

accompanied by a somewhat smaller decrease in ΔV_{DAP} . Such a parameter change can be achieved through modifications of the AP-threshold, fAHP-minimum and/or DAP-maximum, as illustrated by the arrows in Fig. 2A. Various ion channels have been implicated in DAP generation, from sodium and calcium channels (Alessi et al., 2016), to potassium (Eder et al., 1991) and HCN channels (Dickson et al., 2000), which also play a key role for slower grid-cell rhythms (Giocomo and Hasselmo, 2009). These channels could be regulated, e.g., by cholinergic stimulation, which has been shown to induce DAPs and after discharges in MEC-Layer-II neurons (Magistretti et al., 2004). Such modulations would have a direct impact on the precise temporal characteristics of bursting grid cells.

Modulations of the biophysical parameters governing the afterpotentials might even occur at the time scale of single runs through the animal's environment. Indeed, circumstantial evidence suggests that BD^+ cells do not generate a DAP after every AP, and that conversely, some action potentials of BD^- cells are followed by a DAP. One might even postulate that a large fraction of bursty cells are capable of generating DAPs – slice experiments suggest 85% of Layer-II stellate cells and 73% of Layer-II pyramids have DAPs (Canto and Witter, 2012) – but that this mechanism is under external control so as to switch cells between BD^+ and BD^- behavior.

Irrespective of this speculation, there remains the fact that the ISI distributions of BD^- cells have ultra-sharp peaks, whose location varies only minimally within that group. Notably, the same short ISIs are elicited by the sparsely bursting neurons in Layer III (see also Mizuseki et al., 2009) and could be mediated by specific couplings between Layer-II BD^- cells and Layer-III SB neurons. The precise function of burst sequences in the 250-300 Hz regime remains an open question. Similarly, it is not obvious how cells with highly distinct firing characteristics can be orchestrated to create one joint grid-cell network (but see Pastoll et al., 2013), in which all principal neurons have roughly the same spatial tuning properties (Fig. 5E). But with their high burst activity BD^- neurons might be ideally suited to drive other neurons in the network during steady-state operation whereas the DAPs of BD^+ cells might trigger synaptic plasticity, similar to their function in CA3 pyramidal neurons (Mishra et al., 2016), and thus play a critical role for network reconfiguration when the animal learns about new environments (Krupic et al., 2018) or goals (Boccara et al., 2019; Giocomo et al., 2019).

Switching on the DAP mechanism (without interfering with the preceding fAHP) would then increase the probability of additional APs (Alessi et al., 2016) as well as provide a trace for LTP (Mishra et al., 2016) at incoming synapses. Once these are strengthened and the DAP mechanism has been turned off again, the cell can fire precisely tuned bursts with short ISIs. These cell-intrinsic processes could be complemented by precisely wired and timed synaptic inputs (Varga et al., 2010; Couey et al., 2013; Pastoll et al., 2013; Buetfering et al., 2014; Fuchs et al., 2016; Schmidt et al., 2017; Winterer et al., 2017). Through short-term plasticity and integrative postsynaptic processes (Lisman, 1997; Izhikevich et al., 2003) such reorganization could result in a stronger influence on downstream neurons.

In contrast to what one might have expected, the strong dependence of DAPs on the neuron's recent history (Alonso and Klink, 1993; Canto and Witter, 2012; Alessi, 2016) does not seem to translate into a spatial burst code. For example, one might have hypothesized that the DAP of a Layer-II stellate cell should be particularly large when the animal is moving into one of the cell's firing fields, as this corresponds to raising the membrane potential away from its previous out-of-field hyperpolarization. However, we could not find any signature for the ring-like burst-field structure expected in this scenario. In fact, we could not find *any* spatial dependencies despite vigorous search. This came as a surprise, given the role of burst firing for spatial coding in the hippocampus (Harris et al., 2001) or subiculum (Simmonnet and Brecht, 2019). Similarly, spike doublets do not seem to play any special role for burst coding.

Together, these findings suggest that grid-cell bursts are either not utilized for spatial coding, apart from their contribution to theta-phase precession (Hafting et al., 2008, Reifenshtein et al., 2012), or that the spatial coding is masked by temporal fluctuations that are uncorrelated with spatial coordinates.

It is well known that after-spike potentials play a critical role in the control of AP firing patterns. For example, medium afterhyperpolarizations control theta-band clustering of action potentials in MEC stellate neurons (Fransen et al., 2004; Pastoll et al., 2013). Our study extends these and related findings to the 250-300 Hz range and provides a novel mechanistic explanation of grid-cell burst-firing in navigating rodents.

References

- Alessi C, Raspanti A, Magistretti J (2016) Two distinct types of depolarizing afterpotentials are differentially expressed in stellate and pyramidal-like neurons of entorhinal-cortex layer II. *Hippocampus* 26:380–404.
- Alonso A, Klink R (1993) Differential electroresponsiveness of stellate and pyramidal-like cells of medial entorhinal cortex layer II. *J Neurophysiol* 70:128–143.
- Anastassiou CA, Perin R, Buzsáki G, Markram H, Koch C (2015) Cell type- and activity-dependent extracellular correlates of intracellular spiking. *J Neurophysiol* 114:608–625.
- Boccaro CN, Nardin M, Stella F, O’Neill J, Csicsvari J (2019) The entorhinal cognitive map is attracted to goals. *Science* 363:1443–1447.
- Buetfering C, Allen K, Monyer H (2014) Parvalbumin interneurons provide grid cell-driven recurrent inhibition in the medial entorhinal cortex. *Nature Neuroscience* 17:710–718.
- Canto CB, Witter MP (2012) Cellular properties of principal neurons in the rat entorhinal cortex. II. The medial entorhinal cortex. *Hippocampus* 22:1277–1299.
- Chrobak JJ, and Buzsáki G (1998) Gamma oscillations in the entorhinal cortex of the freely behaving rat. *J Neurosci* 18:388–398.
- Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser MB, Moser EI (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* 462, 353–357.
- Couey JJ, Witoelar A, Zhang SJ, Zheng K, Ye J, Dunn B, Czajkowski R, Moser MB, Moser EI, Roudi Y, Witter MP (2013). Recurrent inhibitory circuitry as a mechanism for grid formation. *Nat Neurosci* 16:318–324.
- Domnisoru C, Kinkhabwala AA, Tank DW (2013) Membrane potential dynamics of grid cells. *Science* 495:199–204.
- Diehl GW, Hon OJ, Leutgeb S, Leutgeb JK (2017) Grid and nongrid cells in medial entorhinal cortex represent spatial location and environmental features with complementary coding schemes. *Neuron* 94: 83–92.
- Ebbesen CL, Reifenshtein ET, Tang Q, Burgalossi A, Ray S, Schreiber S, Kempter R, Brecht M (2016) Cell type-specific differences in spike timing and spike shape in the rat parasubiculum and superficial medial entorhinal cortex. *Cell Reports* 16:1–11.
- Eder C, Ficker E, Gündel J, Heinemann U (1991) Outward currents in rat entorhinal cortex stellate cells studied with conventional and perforated patch recordings. *European Journal of Neuroscience* 3:1271–1280.

- Engel TA, Schimansky-Geier L, Herz AVM, Schreiber S, Erchova I (2008) Subthreshold membrane-potential resonances shape spike-train patterns in the entorhinal cortex. *J Neurophysiol* 100: 1576–1589.
- Fransén E, Alonso AA, Dickson CT, Magistretti J, Hasselmo ME (2004) Ionic mechanisms in the generation of subthreshold oscillations and action potential clustering in entorhinal layer II stellate neurons. *Hippocampus* 14:368–384.
- Fuchs EC, Neitz A, Pinna R, Melzer S, Caputi A, Monyer H (2016) Local and distant input controlling excitation in layer II of the medial entorhinal cortex. *Neuron* 89:194–208.
- Giocomo LM, Hasselmo ME (2009) Knock-out of HCN1 subunit flattens dorsal-ventral frequency gradient of medial entorhinal neurons in adult mice. *Journal of Neuroscience* 29:7625–7630.
- Hafting T, Fyhn M, Molden S, Moser M-B, Moser EI (2005) Microstructure of a spatial map in the entorhinal cortex. *Nature* 436:801–806.
- Harris KD, Hirase H, Leinekugel X, Henze DA, Buzsáki G (2001) Temporal Interaction between Single Spikes and Complex Spike Bursts in Hippocampal Pyramidal Cells. *Neuron* 32:141–149.
- Hasselmo ME (2013) Neuronal rebound spiking, resonance frequency and theta cycle skipping may contribute to grid cell firing in medial entorhinal cortex. *Phil Trans Royal Soc B*. 369:20120523.
- Izhikevich EM, Desai NS, Walcott EC, Hoppensteadt FC (2003) Bursts as a unit of neural information: selective communication via resonance. *Trends Neurosci* 26:161–167.
- Krupic J, Bauza M, Burton S, Barry C, O’Keefe J (2015) Grid cell symmetry is shaped by environmental geometry. *Nature* 518:232–235.
- Krupic J, Mauza M, Burton S, O’Keefe J (2018) Local transformations of the hippocampal cognitive map. *Science* 359:1143–1146.
- Latuske P, Toader O, Allen K (2015) Interspike intervals reveal functionally distinct cell populations in the medial entorhinal cortex. *Journal of Neuroscience* 35:10963–10976.
- Lisman JE (1997) Bursts as a unit of neural information: making unreliable synapses reliable. *Trends Neurosci* 20:38–43.
- Magistretti J, Ma L, Shalinsky MH, Lin W, Klink R, Alonso A (2004) Spike patterning by Ca^{2+} -dependent regulation of a muscarinic cation current in entorhinal cortex layer II neurons. *J Neurophysiol* 92: 1644–1657.
- Mishra RK, Kim S, Guzman SJ, Jonas P (2016) Symmetric spike timing-dependent plasticity at CA3–CA3 synapses optimizes storage and recall in autoassociative networks. *Nat comm* 7:11552.
- Mizuseki K, Sirota A, Pastalkova E, Buzsáki G (2009) Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* 64:267–280.
- Newman EL, Hasselmo ME (2014) Grid cell firing properties vary as a function of theta phase locking preferences in the rat medial entorhinal cortex. *Front Syst Neurosci* 8:193.
- Pastoll H, Solanka L, van Rossum MC, Nolan MF (2013) Feedback inhibition enables theta-nested gamma oscillations and grid firing fields. *Neuron* 77:141–154.

- Ray S, Naumann R, Burgalossi A, Tang Q, Schmidt H, Brecht M (2014) Grid-layout and theta-modulation of layer 2 pyramidal neurons in medial entorhinal cortex. *Science* 343: 891–896.
- Reifenstein ET, Kempter R, Schreiber S, Stemmler MB, Herz AVM (2012) Grid cells in rat entorhinal cortex encode physical space with independent firing fields and phase precession at the single-trial level. *Proc. Natl. Acad. Sci. USA* 109: 6301–6306.
- Rousseeuw PJ (1987) Silhouettes: a Graphical Aid to the Interpretation and Validation of Cluster Analysis. *Computational and Applied Mathematics*. 20: 53–65.
- Rowland DC, Obenhaus HA, Skytøen ER, Zhang Q, Kentros CG, Moser EI, Moser M-B (2018) Functional properties of stellate cells in medial entorhinal cortex layer II. *eLife* 7:e36664.
- Sasaki T, Leutgeb S, Leutgeb JK (2014) Spatial and memory circuits in the medial entorhinal cortex. *Curr Opin Neurobiol* 32:16–23.
- Schmidt H, Gour A, Straehle J, Boergens KM, Brecht M, Helmstaedter M (2017) Axonal synapse sorting in medial entorhinal cortex. *Nature* 549:469–475.
- Schmidt-Hieber C, Häusser M (2013) Cellular mechanisms of spatial navigation in the medial entorhinal cortex. *Nature Neuroscience* 16:325–331.
- Simonnet J, Brecht, M (2019) Burst firing and spatial coding in subicular principal cells. *J Neurosci* 39:3651–3662.
- Stensola H, Stensola T, Solstad T, Frøland K, Moser M-B, Moser EI (2012). The entorhinal grid map is discretized. *Nature* 492:72–78.
- Sun C, Kitamura T, Yamamoto J, Martin J, Pignatelli M, Kitch LJ, Schnitzer MJ, Tonegawa S (2015) Distinct speed dependence of entorhinal island and ocean cells, including respective grid cells. *PNAS* 112:9466–9471.
- Tang Q, Burgalossi A, Ebbesen CL, Ray S, Naumann R, Schmidt H, Spicher D, Brecht M (2014) Pyramidal and Stellate Cell Specificity of Grid and Border Representations in Layer 2 of Medial Entorhinal Cortex. *Neuron* 84:1191–1197.
- Varga C, Lee SY, Soltesz I (2010). Target-selective GABAergic control of entorhinal cortex output. *Nat. Neurosci.* 13:822–824.
- Winterer J, Maier N, Wozny C, Beed P, Breustedt J, Evangelista R, Peng Y, D’Àlbis T, Kempter R, Schmitz D (2017) Excitatory Microcircuits within Superficial Layers of the Medial Entorhinal Cortex. *Cell Reports* 19:1110–1116.

Figure 1.

Spike afterpotentials of MEC cells from mice moving in virtual corridors.

A, typical examples of grid-cell burst behavior. Left panel: autocorrelation function of a sparsely bursting cell; middle panel: a bursting cell with broad autocorrelation flanks; right panel: a bursting cell with sharply peaked autocorrelation.

B, grid cells differ in their spike afterpotentials. Left panel: a monotone repolarization that is gradually slowing down; middle panel: fast hyperpolarization (fAHP) followed by a depolarizing afterpotential (DAP); right panel: a short repolarization that abruptly turns into a much slower voltage decay, which may include an initially flat shoulder.

C, characterization of spike afterpotentials. Inset: Definition of parameters. Main panel: group data.

Figure 2.

Spike-time autocorrelations of MEC cells from mice moving in virtual corridors.

A, principal component analysis of spike-time autocorrelations suggests a separation in two major groups, sparsely bursting (blue) and bursting, as supported by k-means clustering. Based on the intracellular measurements (see Fig. 1), the group of bursting neurons can be subdivided in cells with DAP ("BD⁺"), which are shown in red, and cells without detectable DAP ("BD⁻"), shown in yellow. The arrows mark the example neurons shown in Fig. 1.

B, mean autocorrelation.

C, the first two principal components of the spike-time autocorrelations. The pronounced peaks in B and C demonstrate that inter-spike intervals of around 4ms are indicative of both the mean grid-cell discharge patterns and their cell-to-cell variability.

Figure 3.

Group-level analysis of MEC cells from mice moving in virtual corridors.

A, average spike-triggered membrane potential for isolated action potentials (no further AP within 25ms before and after the AP). The main plot shows data that were aligned to AP onset before the group average was taken, the inset illustrates the absolute membrane potential values.

B, averaged autocorrelation functions.

C, averaged inter-spike intervals distributions. Arrows in B and C highlight bursts of sparsely bursting cells, inset in C with logarithmic time scale emphasizes theta-band activity.

D, averaged intra-burst spike count distributions.

Figure 4.

Spike-time autocorrelations of MEC cells from mice moving in open arenas.

A, k-means cluster analysis (k=3) of spike-time autocorrelations.

B, mean autocorrelation across the entire dataset.

C, the first two principal components of the spike-time autocorrelations. The sharp peaks in B and C again demonstrate the prevalence of short inter-spike intervals (here around 3.3ms) in the mean grid-cell discharge patterns as well as their cell-to-cell variability.

D, autocorrelations averaged across all neurons from one cell group reveal a striking similarity between cells recorded on virtual tracks and in open fields. Strongest deviations are shown by non-bursting / sparsely bursting cells in the Latuske et al. vs. Domnisoru et al. data.

Figure 5.

Comparison of spike-train characteristics and spatial coding across data sets.

A, mean firing rates.

B, fraction of ISIs below 8ms ("burstiness").

C, shortest ISIs (first decile)

D, ISI peak, i.e., most likely inter-spike interval.

E, width of the ISI distribution.

F, grid score.

G, spatial information.

H, head direction score.

Despite strong differences in temporal spike-train differences (**A-E**), the spatial coding properties (**F-H**) of grid cells are largely conserved across all three cell groups.

Figure 6.

Cluster structure of bursty grid cells

Dependence of key spike-train parameters on fast hyperpolarization (ΔV_{fAHP}) and depolarization (ΔV_{DAP}).

A and **B**, mean firing rates.

C and **D**, fraction of ISIs below 8ms ("burstiness"). Across the entire population of bursty neurons, the larger ΔV_{fAHP} , the more frequent are short ISIs

E and **F**, fraction of ISIs between 8 and 25ms.

G and **H**, location of ISI peak.

While the firing rates do not exhibit a trend, neither within the two cell groups nor across the groups, the other quantities depicted show trends that are statistically different from the null-hypothesis (no increase/decrease as a function of ΔV_{fAHP} or ΔV_{DAP}) that are consistent with the proposed functional role of post-AP voltage deflections. In addition, the data suggest that the population of bursty neurons either forms one joint though under-sampled cloud or contains two distinct sub-populations. In both cases, however, the cells' spike-train characteristics do depend on their DAP properties.

Figure S1.


Stability of afterpotential parameters.

To quantify the reliability of the parameters characterizing the spike afterpotentials, we carried out a bootstrapping analysis. The error bars indicate s.e.m. as obtained from 1000 samples. The three arrows show how a cell's position in the ΔV_1 - ΔV_2 space changes under variations of the parameter indicated. Such variations could arise due to measurement errors or modulations of cell parameters.

Figure S2.

Robustness of cluster analysis.

To test the robustness of the PCA-based class assignment of the data from Latuske et al. (2015), we separately considered the first and second half of all spikes for each neuron. We then computed the autocorrelations within these two sets and projected the results into the PC space of the full data. K-means clustering ($k=3$), carried out in the same way as had been done for the full data set, results in some neurons switching group identity. Only 6.7% of cells switch identity, when 1st and 2nd halves are compared, underscoring the robustness of the cluster analysis.



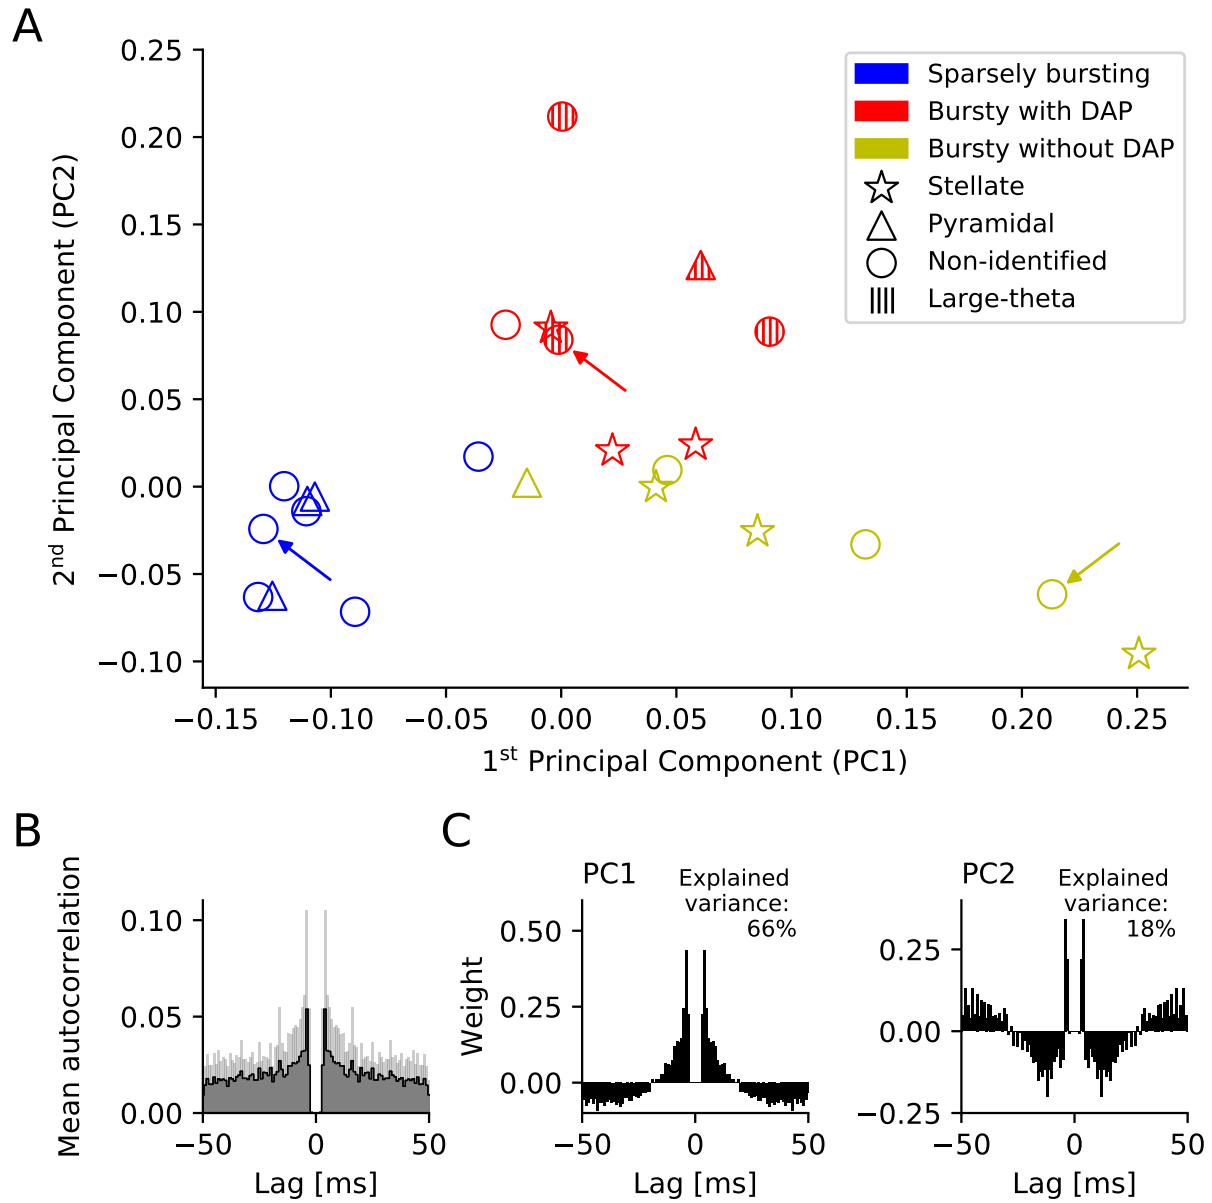


Figure 2. Spike-time autocorrelations of MEC cells from mice moving in virtual corridor.

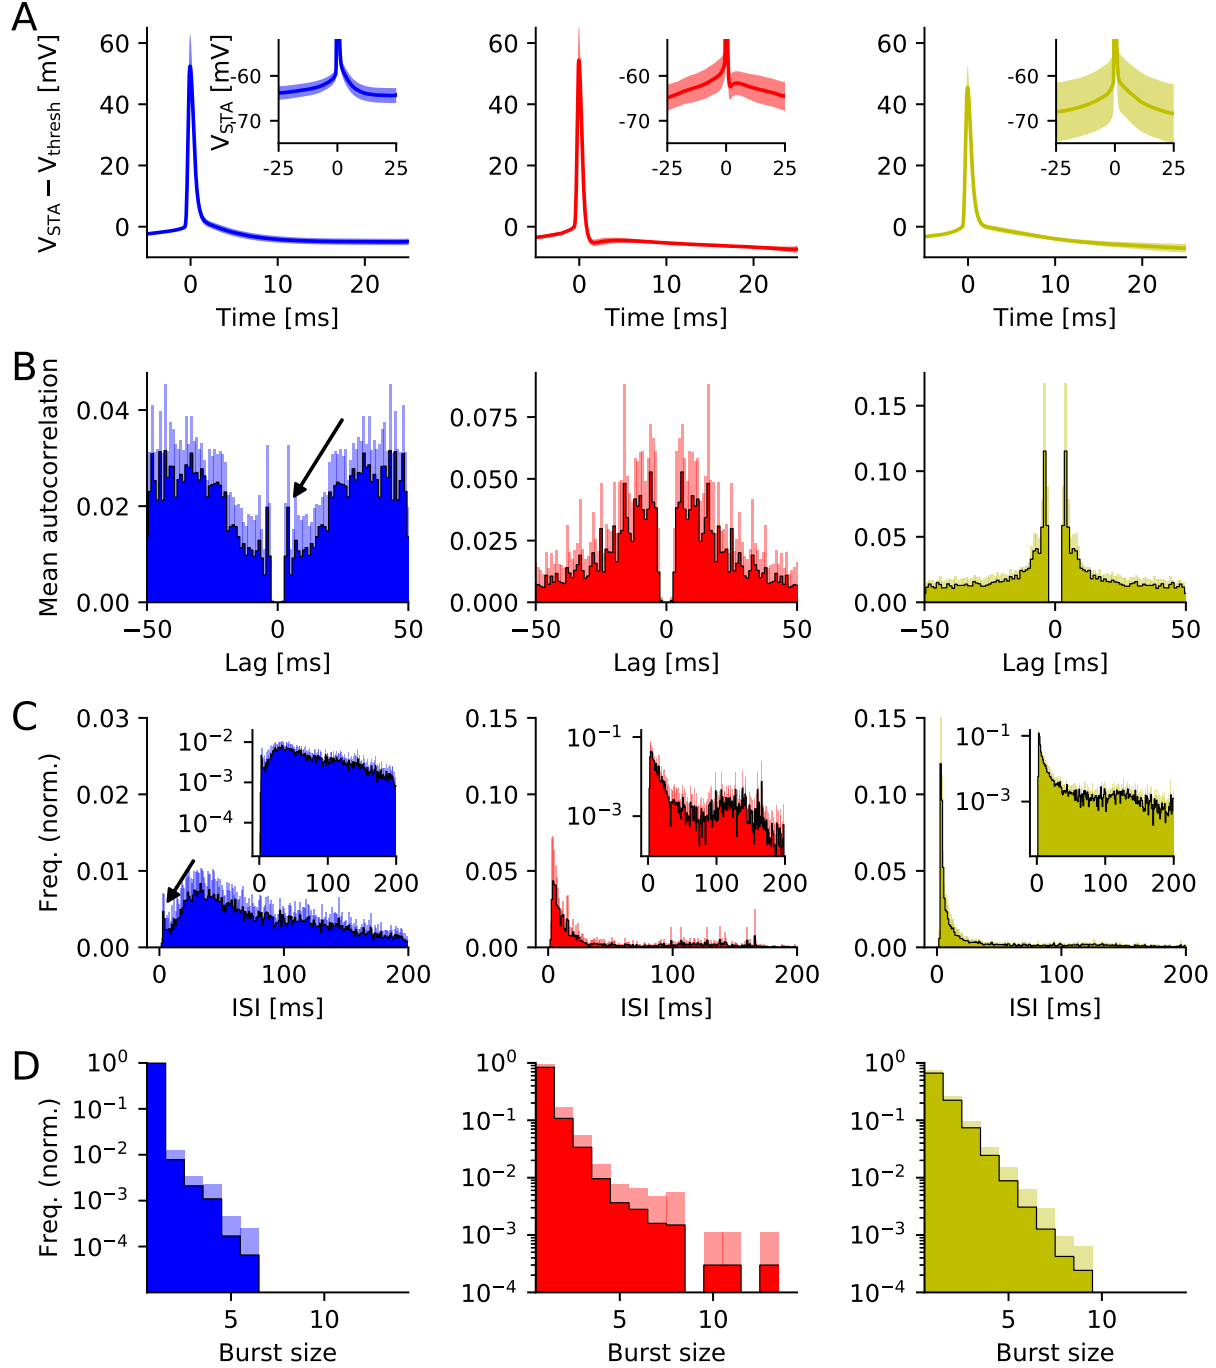


Figure 3. Group-level analysis of MEC cells from mice moving in virtual corridors.

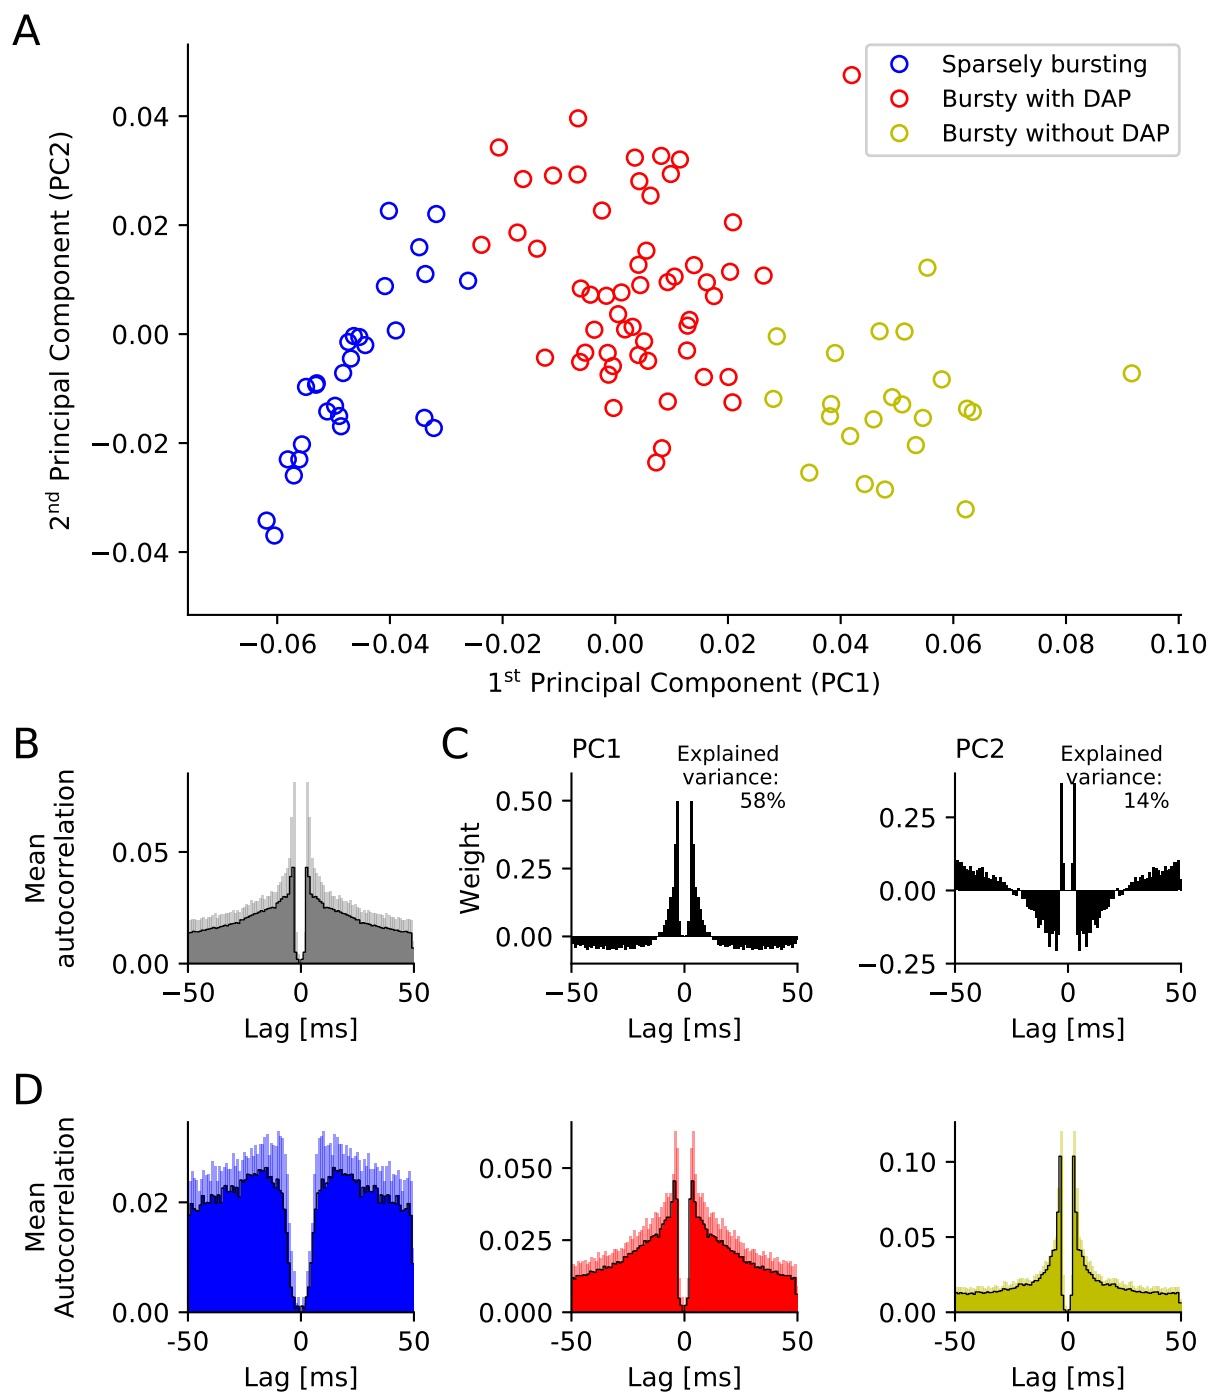


Figure 4. Spike-time autocorrelations of MEC cells from mice moving in open arenas.

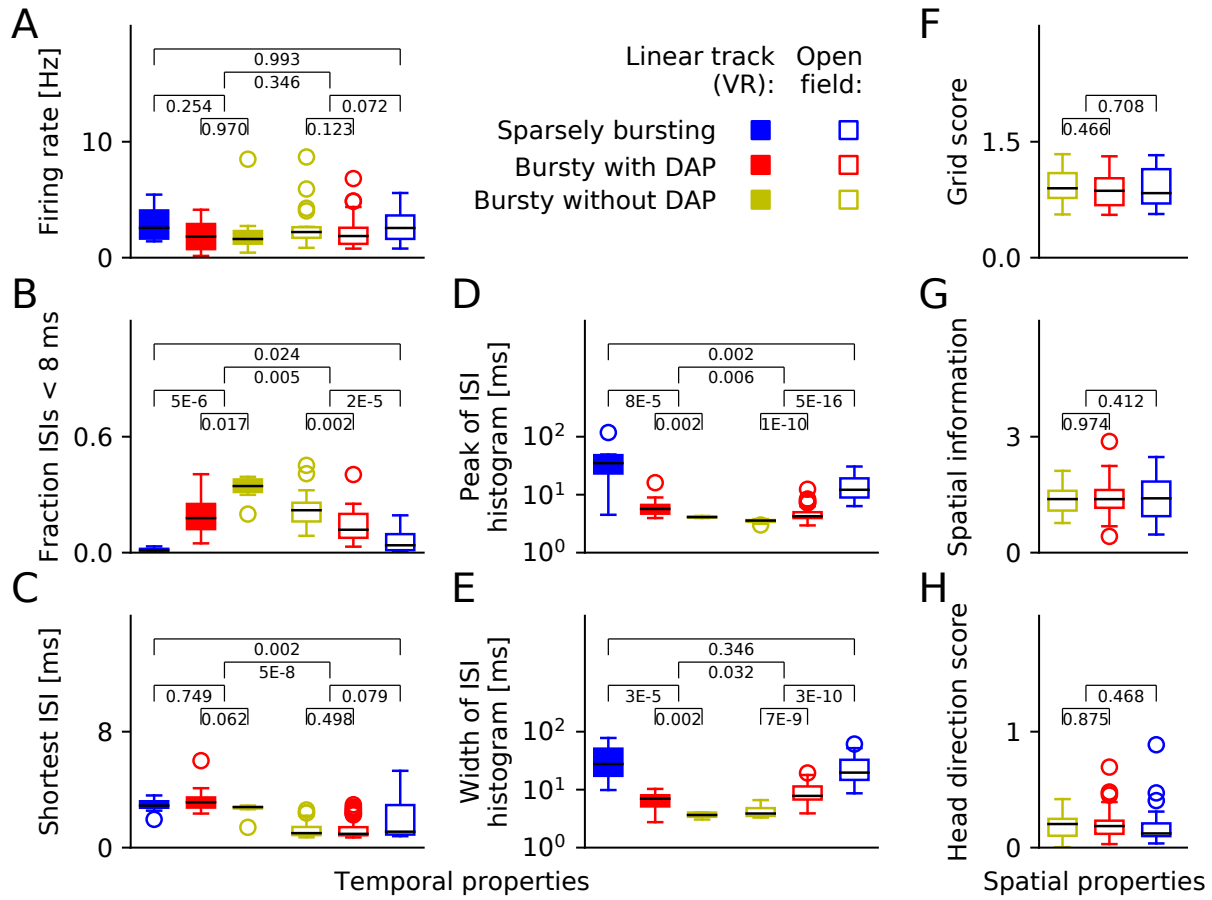


Figure 5. Comparison of spike-train characteristics and spatial coding across data set.

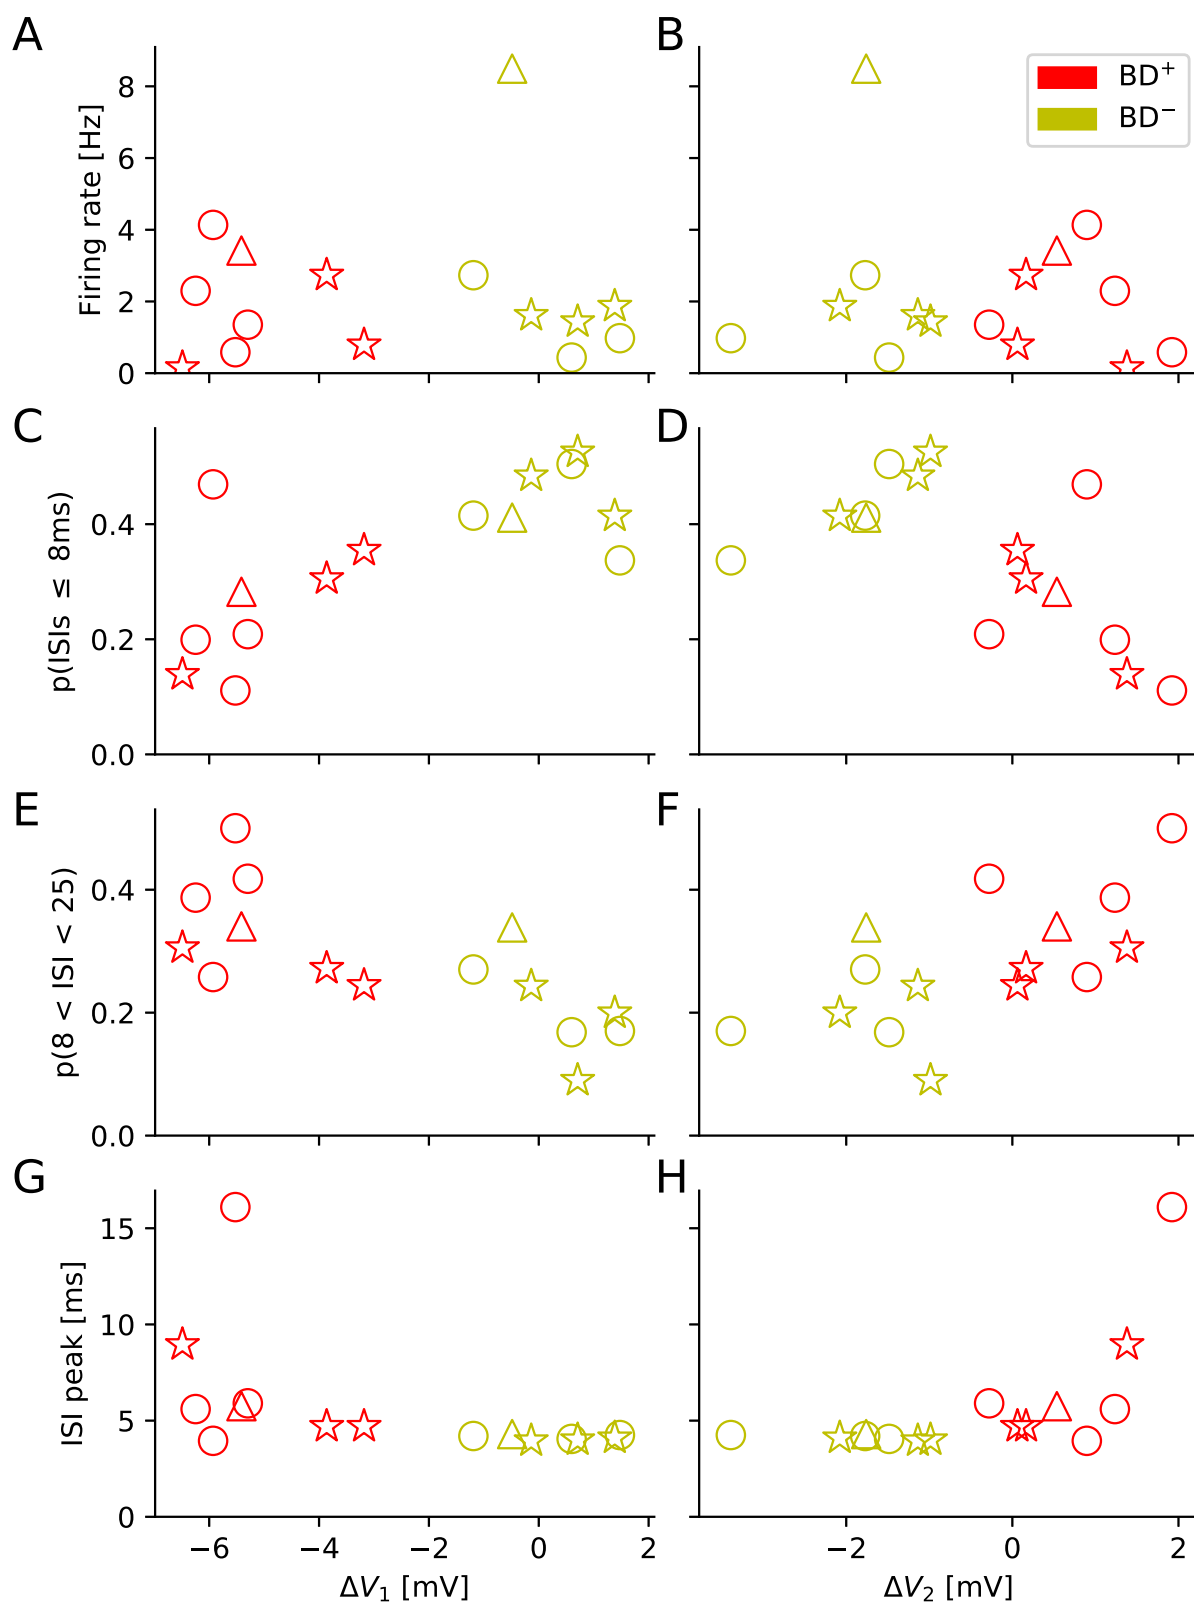


Figure 6. Cluster structure of bursty grid cell.

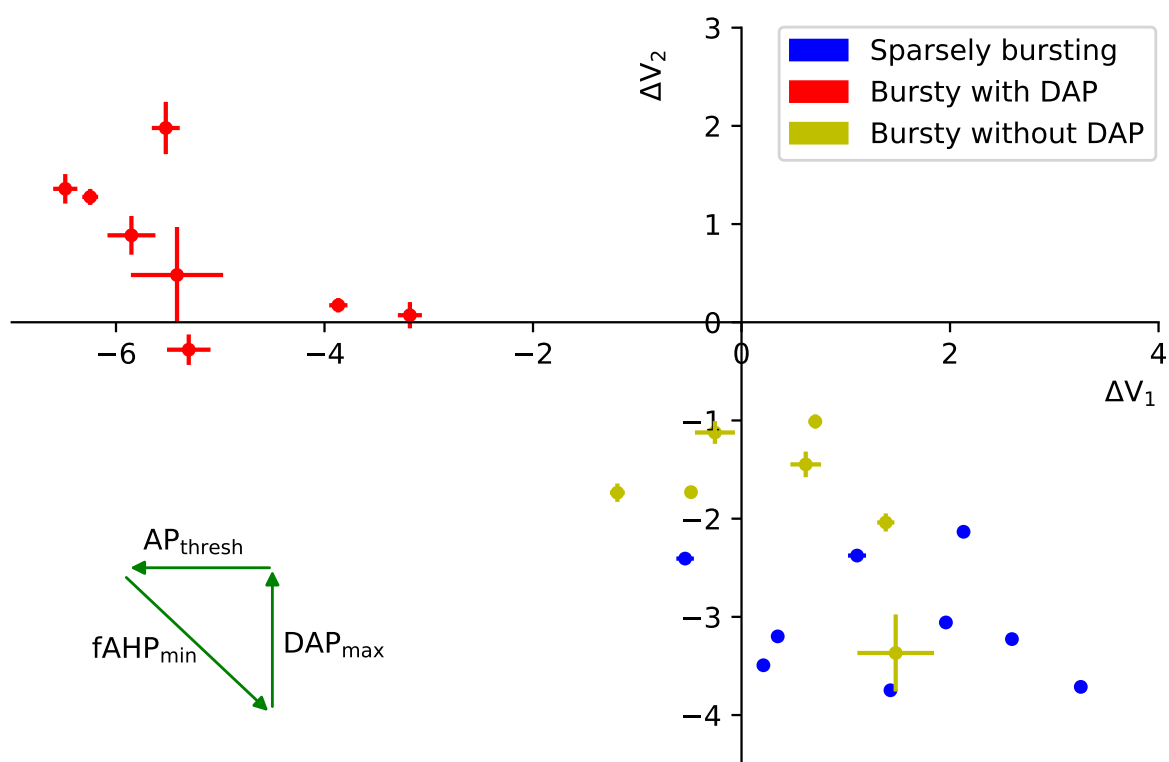


Figure S1. Stability of afterpotential parameters.




Figure S2. Robustness of cluster analysis.

Chapter 4

Discussion

In my dissertation I studied the variability of neuronal representations of space. In particular, I investigated spatially modulated cells in the temporal lobe with a focus on the grid cell population in the medial entorhinal cortex (MEC) of rodents. These cells are characterized by being highly active only when the animal crosses certain regions in space. These regions, called firing fields, are arranged in a hexagonal grid. I wanted to understand how reliable grid cells signal the location of the animal and whether variability in their activity is just noise or can be attributed to their response to unknown factors. While this question simply reads, its answer is not trivial.

In Chapter 2 of this thesis we showed that under assumption of stable spatial tuning curves grid and place cells exhibit excess variability in their spiking on crossings of firing fields in the open field. This phenomenon, called *overdispersion*, is well known for place cells in the CA1 of hippocampus [111]. By reanalyzing data from the Moser lab we extended these findings to other neurons involved in spatial-navigation tasks: For place cells in the CA3 we measure even higher variability, presumably due to differential functional roles of CA1 and CA3 place cells [116, 117, 118], whereas simultaneously recorded MEC grid cells vary comparably to CA1 cells in their spiking.

A detailed analysis of spike count distributions of grid cells reveals that such cells frequently have a mode at zero, even when the overall firing rate is high. In other words unexpectedly long inter spike intervals (ISIs) are present within firing fields. This motivated us to investigate how much of the variability is captured by introducing a zero inflation (ZI) to the otherwise Poisson spiking process. Initially no assumptions on the temporal structure of this inflation were made. Such relatively simple model turns out to be superior to the Poisson model in fits to the data and is able to capture large amounts of the excess variability. This finding was extended to the negative binomial spiking model and its zero-inflated variant but is not reported in this thesis.

The conclusion is the same: the specific mode at zero spiking is the main contribution to trial-to-trial variability of grid cells, but not exclusive to grid cells and found in irregular spatially modulated cells in the MEC as well. The zero-inflated model variants as used in the manuscript assisted the quantification of excess zeros but the inflation was not linked to external factors and modelled as a stochastic component. In other words, moments

of neural quiescence could not be predicted (see next section for some details). However, we found that the chance of extra zeros in a grid cell's spike counts was correlated across experimental conditions even for non-conjunctive cells. The large variability in the occurrence of ZI across the cell population and the consistency across settings is compatible with several hypothesis, i.e., that the ZI might be a cell-specific feature or result from a long lasting state at the network level. Drifting firing fields as reported earlier in very different experimental circumstances [119] could, in principle, lead to the observed variability as well. We therefore looked at linear track recordings [5, 26] where the animal's trajectories are more constrained and enable more direct tracking of firing field locations over trials.

In fact, we found drifts in the firing field locations. The drifts were largely correlated across simultaneously recorded grid cells. Such non-stationarities induce noise correlations when comparing the data to stationary statistical models. Corrected for the drifts, the spiking is less variable in space and provides more information about the location. We concluded that the internally represented location, which is encoded in the grid-cell network, is drifting in these recordings. Translational shifts of grid-cell firing fields due to contextual changes have been reported by Marozzi et al. [120]. This provides an interesting perspective on our findings: the relevance of certain aspects of a context to the mouse might fluctuate on small timescale. Thereby, small translational shifts could be induced even within the same environment. Such motion of the firing fields is difficult to follow in the open field due to low sampling. However, the fact that the drifts of grid-cell firing fields are correlated across neurons offers the possibility to track field shifts by Bayesian decoding. For this purpose it is important to limit the training data to snapshots of the recording where a stationary tuning can be assumed with high certainty (e.g. in situations with low zero inflation). On this basis, drifting fields could be tracked in other moments.

One example would be to use grid cell population activity in recordings during light to decode the positions that are represented by the population when the light was turned off. The deviation of the decoded position from the true position as well as jumps in the decoded position can help to understand the drifting. However, detecting field jitter of a few cm requires high precision and therefore many co-recorded grid cells.

As mentioned in the discussion in Chapter 2, in recent years, there has been an interesting turn in the consideration of neural networks, partly also due to the increasing computing power: a topological reconstruction of the underlying geometry embedded in the neural firing does not require knowledge about the signal, i.e. the location of the animal, because it is solely based on similarity metrics applied to the population vector of neurons [121, 122, 123, 124, 125, 126, 127]. If a low dimensional cognitive map in the sense of Tolman [14] is embedded in the activity of neurons in the hippocampus and the entorhinal cortex then dimensionality-reduction techniques and manifold learning might recover that map [128, 129]. Trial-to-trial variability of grid cells that is highly correlated across the population in the way we observed has no effect on a low dimensional manifold in the neural code that is related to space but might expand along extra dimensions independent of location. There is a chance that such dimensions capturing the variability can be correlated to other relevant, yet unobserved, aspects.

During my doctoral work I also analyzed the other side of the ISI spectrum: My interest

was not limited to the origin of the surprisingly long ISIs at about 400ms, but I also studied very small ISIs below 50ms and their effects on neuronal variability. On that scale *bursting* takes place in the MEC [87]. Burst firing not only induces temporal correlations in the spike trains and thereby affects the coefficient of variation (CV), but also contributes to trial-to-trial variability on small time scales: If a cell on some trials triggers bursts within some spatial bins and not on others as the mouse crosses a firing field, typically a large variability was measured. Bursts of many grid cells are usually followed by periods of quiescence which can be seen by a peak at around 100ms in the ISI return maps (even when pooled across neurons, see Fig. E.1 for an example).

Therefore, in Chapter 3 we focused on relatively large timescales ($\approx 400ms$) but we also tested even larger scales. On that scales we could not find differences in the ZI across grid cells with different spiking characteristics on small temporal scales (bursty versus non-bursty, as suggested by Latuske et al. [87]). However, the cell classification into bursty and non-bursty might dismiss important aspects of the variability in the grid cells' burst characteristics (see Chapter 3 for a detailed discussion of this aspect).

In fact, further grid-cell simulations indicate that pauses on such time scales are unlikely even in scenarios where all spikes contribute to burst events (data not shown). That is, because there is limited space (i.e. the firing fields) for the bursts to occur while maintaining the overall firing rate map. Moreover, the correlated shifts of firing fields across cells on different trials suggest that such bursts needed to be correlated as well. Correlated bursting could be orchestrated by theta-gamma coupling [130]. Based on our findings, I suggest that high firing rate periods and long spiking pauses are reflecting *in-field* and *out-of-field* states [26] of single cells that are orchestrated via fix couplings within the circuit.

It is not clear whether in-vivo bursts in the MEC reflect a special spiking regime. In particular, it is unknown whether there are differences in the high firing-rate regime as compared to expectations based on renewal processes given the grid cells' ISI distributions. However, an analysis in the Appendix (Fig. E.1C) shows that spikes with small ISIs come in packages and are stronger clustered in time than expected from the renewal hypothesis.

Lately, bursts were shown to define sharper firing fields and carry more spatial information than other spike events in a subset of principal cells in the subiculum [131]. However, it is unknown whether bursts play a special functional role in the MEC. We hypothesized that the firing of bursts might highlight special events: Previous reports focused on bursts of Layer 2 stellate cells in vitro [93, 94]. They found increased bursting rates following hyperpolarization periods. That suggests an over-abundance of short ISIs in vivo whenever the animal is entering a firing field as in such a moment the membrane voltage of grid cells is ramping up [26].

To test this hypothesis, among other related ideas, we treated each burst event as a single "spike-like" event and computed burst-rate maps, grid-scores and spatial information. Similarly, we estimated the Poisson-likelihoods of any individual ISI in a spatially resolved manner based on the firing rate maps. Short ISIs with a low likelihood were expected to appear at firing-field edges. However, while we were not able to clearly proof or falsify that hypothesis despite extensive effort, the findings presented in Chapter 2 offer a possible explanation for our result: The trial-to-trial shifts we observed - not only in

the firing rates, but specifically in the locations of firing - induce errors in the analysis of spatial events. While we expected a ring structure to emerge around the field center when restricting analyses to bursts instead of single spikes (considering stationary firing fields), a jitter of the field position on a moment-to-moment base smears such a ring. As a consequence, dependent on the distribution of the jitter, the ring can be transformed into a plateau or a large Gaussian bump and other shapes. Another possible functional role of bursting in grid cells that was not tested might be synchronization of brain rhythms [80]. Grid-cell spikes [132] as well as gamma oscillations [133] tend to emerge at specific theta phases, but the interaction of cell firing and theta oscillation in the LFP is not completely clear yet [130]. However, computational models suggest that such cross-frequency couplings enable information transfer across brain regions [134]. Experiments indicate that theta-gamma cross-frequency couplings in the entorhinal-hippocampal circuit play a critical role in memory formation [135] and are impaired in an Alzheimer's disease mouse model [136].

Similarly, little is known about the cellular mechanisms of the burst generation in MEC: Are bursts only the consequence of abrupt reduction of inhibition [93] and is the leading spike of a burst elicited by rebound spiking [137]? Are grid cells specifically predisposed to fast and rhythmic spiking due to gamma modulations [130, 136, 138] or is the repetitive spiking triggered and maintained after onset *internally*? In Chapter 3 we studied how bursting in grid cells and spike afterpotentials (DAP) are related. The DAPs in the MEC have only been observed in-vitro [93, 94]. To this end, we analyzed the in-vivo whole-cell recordings of Domnisoru et al. [26] and showed that grid cells show DAPS in-vivo as well, but not as frequent as compared to in-vitro slice recordings. Furthermore, we were able to identify bursty and non-bursty cell clusters in the whole-cell recordings. Despite small differences in the detailed temporal correlation structure the clusters could be validated by a larger dataset of in-vivo tetrode recordings from Latuske et al. [87]. As discussed in the manuscript it remained unclear whether such clusters reflect diversity in principal cell types in the MEC or rather indicate that physiological properties of neurons are modulated dynamically. Furthermore, it could be speculated, that the DAP as well as bursting are the results of dendritic backpropagation [139, 140, 141] or reflecting interactions on a network level: bursting might be a resonance phenomena between interneurons and principal cells that are known to build tightly connected microcircuits [142, 143, 144].

DAPs could be an artefact of such dynamics in situation in which feedback inhibition kicked in right after a spike and excitation was not sufficient to reach the AP threshold. However, the variability of the mode in the ISI distribution of bursty grid cells without DAP was found to be surprisingly small. Therefore, it seems unlikely that such precision can be acquired without tuning of cell-internal properties. Nevertheless, the question remains open, whether variability in function and physiology reflects different classes or situations.

4.1 Complementary work

The study in Chapter 3 was complemented by extensive in-vitro patch-clamp recordings of neurons in the MEC of rats and gerbils by Dr. Franziska Kümpfbeck (her dissertation will be made available at <https://edoc.ub.uni-muenchen.de/>). That data was analyzed using the mecPhysio Toolbox described in the appendix E.2. Our results supplemented the results of Alessi et al. [94] and above all questioned the identification of cell types via frequently used immunohistological-staining approaches in the MEC. Overall, it became clear to us that the neurons in the entorhinal cortex can hardly be classified into well defined groups.

Together with Yulia Kostina, we fitted generalized linear models with history and stimulus filter to the grid cells from the Domnisoru dataset. This approach is described in Pillow et al. [145] and was suggested for cell classification based on some simulations but it has never been applied to real recordings. As a proxy for the input current, the membrane potential (spikes removed) of the neurons was fed into to the GLM. We then performed a clustering on the resulting parameters. The clusters were largely identical to the one described in Chapter 3.

The firing field jitter we reported in Chapter 2 lead to spiking pauses that were unexpected with respect to classically used firing rate maps. This essential finding was backed up by many parallel analysis using general HMM models, constrained versions thereof (continuous ZI models, see appendix), zero inflated generalized linear models (ZIGLM) and comparing to alternative observation models, which can increase the probability for zero counts as well (like the generalized Poisson model [146, 147]) and are not described in the manuscript. Moreover, in the CA1 and CA3 recordings from Fyhn et al. [77, 148] we found indications for zero inflation as well. However, the number of cells and field crossings was very limited.

I would like to point out that zero-inflation is a very robust phenomenon beyond to what is worked out in Chapter 2: In extensive surrogate modelling I found that none of a large bank of modifications to inhomogeneous-Poisson spiking models of grid cell activity could explain the correlation lengths found in grid cell autocorrelations on scales of hundreds of milliseconds. Such modifications included bursting states, theta and gamma modulations, conjunctive tuning and various manipulations of the field shapes while preserving the overall firing rate of the real neuron during the experiment. The only way I was able to reproduce such temporal correlation structure was by introducing long lasting down states. These could be quantified using the zero-inflated models and were found in all grid cell data at hand to us, whether recorded from rat or mice and independent of the shape of the arena and the laboratory.

By carefully reanalyzing a dataset of Stensola et al. [36] in a earlier study we found evidence for stronger correlations across co-recorded neurons in the unexpected down-states than in unexpected up-states (corresponding to bursting periods) [149]. We wondered whether these states are associated to features of the local field potential (LFP), specifically amplitude, power and frequency of delta ($\approx 1\text{Hz}$), theta ($7 - 12\text{Hz}$) and gamma ($> 50\text{Hz}$) oscillations. We could not identify obvious relations but some data indicated negative

correlations of gamma power and off-states. Such correlations would not be in contradiction with shifted firing fields. Further analyses are needed to test this observation in detail.

Initially it is was not clear whether the factors leading to zero inflation are fix properties of single cell, dynamic in time or network phenomena. Therefore we performed extensive cross-validated analysis of zero inflated linear-nonlinear Poisson models for MEC neurons. Such models are a variant of the more general ZIGLM. This enabled us to investigate whether neurons tend to exhibit quiescence in some moments and not in others. These models included the factors position, head direction, speed, time since start of the trial and all combinations of them testing their influences in both regimes, the Poisson regime as well as the zero inflation regime. With that, we were able to test whether zero inflation changes across the duration of a session.

We found that for nearly all cells that we studied (Escobar et al. data [150] and Stensola et al. data [36]) a single tuning dimension for the Poisson in combination with a constant zero inflation or a spatially modulated zero inflation was preferred. Typically, none of the covariates considered, except the position, was able to predict the sequence of the neuronal-silence phases better than chance. This is in alignment with the observations of field jitter on the linear track as field displacements can either induce rather uniform distribution of excess zero counts along the track (for large jitter) or spatially confined to the edges of firing fields (small jitter).

The resulting zero-inflation parameters for the Escobar et al. data were correlated across light contexts and environments (linear track and open field) and with the ones reported in Chapter 2 where spikes were collected in time bins of varying length and fix expectations.

Furthermore, in alignment with the jitter idea, the pauses tended to be more synchronous than strong bursting events [149]. This was also true after controlling for the animal's running speed. It remains open how much of the zero inflation can be attributed to jitter and what fraction could be explained by additional state modulation. After trial alignment in 1D, however, zero-inflation often remained significant and therefore such state modulations cannot be excluded. Other sources of variability might contribute to the overdispersion as well.

In case of a conjunctive grid cell that fires only reliably in the grid-like fashion whenever the animal is heading towards a certain direction it is clear that sometimes this neuron will not be active within a grid field. Such inactivity will be reflected in a high number of zero counts. I would like to emphasize that in neuroscientific experiments often not all signal dimensions are known and can be controlled and therefore are latent. If, for example, the researchers who discovered the grid cells [16, 17] had only analyzed the tuning conditioned to one spatial dimension instead of two they would occasionally still have found striking periodicity or significant spatial modulation. However, the spike count distributions would have shown even higher excess variability and excess zeros as compared to the case of not being ignorant to the second spatial dimension. In line with that also the coding of the third spatial dimension could explain some of the observed noise. Indeed, for place cell firing it was already shown that they can get inactive when the animal is elevating its head [151]. Similar could be true for grid cell firing. Another example illustrates the

significance of the principle that is sketched here as well: Excess zeros do show up even if a grid cell was simulated 2D/3D with a *deterministic* spiking process and then analysed in 1D/2D. However, restricting the analysis solely to internal covariates (as the activity of other neurons) and ignoring external signals might remove the zero inflation. In fact, the robust correlation structure in the MEC suggests a grid cell’s activity to be reliable with respect to its functional embedding in the network. Being ignorant about the external world a grid cell is mostly doing what it is expected to do. Similarly, the in-vitro patch-clamp slice recordings in the MEC of rats and gerbils performed by Franziska Kümpfbeck suggested highly stereotypical spiking patterns for repeated current injections of the same kind. In summary, we find that that most of the variability in grid-cell spiking is due to network phenomena and is not generated within single neurons.

As an extension to the work presented in chapter 2 we provide a novel and flexible statistical approach to model the effects of field shifts onto spike count distributions towards the end of the Appendix: the Beta-Poisson mixture (BP) model. Fitting the corresponding distribution to count data allows to roughly estimate the jitter distribution of firing fields as well as the shape of the fields without any direct euclidean measures. It is flexible in capturing ZI as well as unimodal excess variability. We found that this model was superior to Poisson, negative binomial and their zero-inflated variants to explain variations in the spike-count distributions of grid cells. These results will be made available to the public in a separate manuscript.

4.2 Time scales and spatial scales of analysis - the impact of binning

We like to mention that generally zero inflation could be detected in all grid-cell recordings at hand to us, whether recorded from rat or mice and independent of the shape of the arena and the laboratory. Quantitatively, the zero-inflation parameter α depends on the time scales on which spikes are pooled but is correlated across parameter ranges and cells.

Generally, zero inflation seems to be a very robust phenomenon as discussed before. Whenever a spike-count or firing-rate have been collected in spatial or temporal bins, the typical sign of zero inflated count distributions emerges: a bimodal distribution with one peak at zero. This general motif becomes especially prominent, whenever we look at bins, in which high firing rate are expected and multiple samples can be collected.

Similarly, and not mentioned in the manuscript as well, the results on the variability σ_z^2 (variance of the standardized firing rates) were qualitatively robust across a wide range of parameters and different firing rate models (in particular smoothed variants like Gaussian kernel smoothing and Gaussian process rate estimation [112]). However, quantitatively we observed a strong effect of the selected values of the hyperparameters. Surprisingly, this was even true for inhomogeneous-Poisson models when smaller resolutions were used to simulate the data as compared to when analyzing the data.

Furthermore, for carefully simulated *independent* Poisson spike trains of grid cells we found positive noise correlations that depend on the grid phase offset. This was the case, whenever the temporal or spatial scale Δx of analysis was larger than the resolution δx with which the data was generated. That is due to a doubly stochastic component that is induced by such discrepancy: A bin centered at value x_i summarizes an extended range of possible values on some axis x , for example from $x_i - \Delta x$ to $x_i + \Delta x$. When analyzing data using such bins, the information about the exact positions of sample points in that bin is omitted and a constant value f_i is associated in that area. Potential gradients and skewed distributions of possible values within the bin are ignored. However, whenever the underlying function $f(x)$ producing the data of interest acts on smaller scales than the bin width that we chose, each sample is truly associated with its precise value on x and not x_i .

The dramatic effect that this difference can have is illustrated in an example in the Appendix (Fig. E.3): Simulating independent Poisson processes with a spatial resolution of 1cm and analysing such data with a binning of 10cm leads to spurious noise correlations that are dependent on the phase offset between grid cells.

Notably, any smoothing of the firing rate map generalizes the estimate whereas using a firing rate histogram is the most naive approach and is prone to overfitting for small binning - as every spatial bin is associated with one parameter. For a bin-width close to zero the actual spiking is perfectly recovered by the rate model resulting in a deterministic model without uncertainty. Therefore using a rate histogram without smoothing and a reasonable small binning is conservative when estimating variability. In contrast using too large bin-width in comparison to the internal resolution of a spike-generating tuning curve can also increase the measured variability as sketched before. Spike distance metrics and

quantification of noise correlations based on spike jittering are mostly model free and offer a valuable alternative.

Hence, we emphasize that bin width has a strong effect on the outcomes in spike-train analysis and can lead to wrong conclusions. This effect is also true for 1D analysis where a temporal phase offset across trials or neurons is present. Whereas for comparative studies using the same bin-size for any data (as in our analysis shown in Fig. 2.9a) this issue is less of importance, it is crucial to carefully control for its effects when quantifying and interpreting overdispersion or noise correlations.

In fact, this problem is widely ignored in the literature. Many authors use surrogate models to test data against some nullhypothesis in the following way: Data is discretized and a tuning curve is estimated. This is potentially smoothed and normalized to match the average firing rate of the data. From this and based on the trajectory in stimulus space, an expected number of spikes in some time window is computed and used as the expectation values for an inhomogeneous Poisson process. Finally spikes are drawn from the model. By comparing variability and noise correlations to the data conclusions are drawn. In this way, data and simulations are generated on different scales, but compared on the same scale. As already mentioned, however, this approach is ill-suited.

4.3 Outlook

To understand an organisms idea of space, and how it navigates within this idea physically and mentally, it is key to understand how such a complex concept is represented in the dynamics of the networks that are assumed to process spatial information. The term *cognitive map* was established in literature for this idea [14]. Wikipedia defines: “A map is a symbolic depiction emphasizing relationships between elements of some space, such as objects, regions, or themes”¹. Not only can relational information be read out from a map, but it is also essential for constructing a map. It took centuries for cartographers to create reliable world maps, mostly because of the absence of reliable distance and angular estimates between remote locations [152]. When thinking about maps in the context of spatial navigation, I typically think of an aerial view - a classic two-dimensional map on paper like an atlas or maps displayed on GPS systems. On the axes of such maps are the latitude and longitude. However, the actual role of maps is to depict relationships between two reference points. Coordinates cannot be absolute, but must be understood as relations to certain reference points on or outside the map. In the game “Topfschlagen”, mentioned in the introduction of this work, knowledge of absolute spatial location of the “self” of the kid would be irrelevant without the absolute coordinates of all obstacles in the room. And relationships would need to be calculated in order to make such information useful. In essence, the important but missing piece of information for the blindfolded kid is the distance and orientation relative to the pot with the sweets. In contrast, the availability of this information to the observers, provides the entertainment. This small

¹Wikipedia contributors. (2020, April 10). Map. In Wikipedia, The Free Encyclopedia. Retrieved 11:15, July 17, 2019, from <https://en.wikipedia.org/w/index.php?title=Map&oldid=906749629>

example motivates us to ask why brains should bother with absolute coordinates at all. I argue that just as relative quantities like free energy and entropy changes are key to thermodynamics and biophysics, relative spatial information is key to navigation. And just as there is no such thing as absolute entropy, there is no absolute location.

In the context of neural representations of space, relational information about places is reflected in the degree of co-activity of cells that are tuned to specific locations in space. When two cells tend to fire at the same spot, their activity will be correlated. However, cells that tend to signal very different areas of the environment will usually not fire together, as the animal is only in one location at a time. Indeed in this scenario, the brain seems to have a very similar problem to early cartographers: There is a lack of relational information. In the first computational steps of spatial information processing of an organism, there is no information present between remote locations that have never been sensed simultaneously. How would one overcome this gap in information processing and how is this related to the unstable grid representations that we observed?

While trying to understand how an animal's trajectory is reflected in the activity of a population of spatially modulated neurons, I made an important observation: The correlation structure between them is maintained across different situations - both in simulations as well as in experiments. In the linear track experiments described in Chapter 2 for example the firing fields of the grid cells drifted in accordance and the degree of co-activity between two grid cells did not vary. This is in accordance with preserved correlation structures in sleep versus awake conditions [153] and light versus dark conditions [154]. A stationary correlation structure combined with a highly variable tuning has some interesting implications and supports the previous statements: Since the relationships encoded in co-activities seem to be so robust, they could consequently also be the essential information building blocks for the construction of a cognitive map. What are these information that such "a co-activity matrix" of neuronal activity contains? Besides it just reflecting the frequency of simultaneous firing of pairs of neurons it also contains a even more valuable feature: That is information about the topology of the neural network [121]. However, if we limit the analysis to pairs of neurons with spatially modulated activity, the resulting matrix contains information on the neighbourhood relationships of the firing fields of these cells in space. Unfortunately, it only contains indirect information about the relative position of distant firing fields. That is, because a pair of neurons with very distant firing fields has as low co-activity as a pair with less distant but still not overlapping fields. Essentially, the matrix will therefore consist of many zeros, i.e. entries representing pairs of neurons whose fields do not overlap. In such cases we speak of a sparse matrix.

To achieve a two-dimensional map containing all the relevant distance information from a co-activity matrix can be seen as a dimensionality-reduction problem. However, a dimensional reduction based on incomplete distance information is not trivial. In order to decode the trajectory of a rat in space from this matrix, a few tricks are needed. Science has found some solutions here, e.g. concatenating local graphs [155] or semi-definite programming [156]. In general, such methods can also be summarized under the term *matrix completion*. These methods find many applications, for example, in coordinating swarms of drones with only short range sensors, generally in the area of sensory network localiza-

tion, and in recommendation systems of well-known streaming services or web shops. In order to optimize these relatively complex procedures, the company Netflix even organized a contest with a reward of US\$1,000,000. Thus, a variant of this problem became popular as the “Netflix problem”. There is also a reason why so much money was paid out to the winners: It is a so-called NP-hard problem.

An efficient way to construct a full distance matrix from a sparse correlation matrix, mainly containing information about close neighborhoods would be desirable. In our case, the study of the coding of motion in two or three dimensions of space, this distance matrix is a Euclidean distance matrix (EDM). A series of points in a Euclidean space is completely described by the EDM up to affine transformations (except shearing) [157]. Through the associated constraints due to the triangle inequality, the problem also becomes more tractable.

The evolution of mammalian brains had plenty of time to find a solution to such a challenge; this will have also happened because the efficient planning of paths to distant destinations requires a vague idea of the global arrangement of points in space and that corresponds to a distance matrix. Somewhere in the brain there must therefore be systems that not only signal current places in themselves, but also put remembered places in relation to each other. Relationships and references are the decisive information for survival; not absolute coordinates of the “self”, other living beings or things in space. Absolute coordinates do not even exist without the definition of an artificial spatially fixed reference system.

Let’s assume, then, that a strategy of matrix completion has been successfully manifested and optimized in the brains of rodents over millions of years: Will “non-negative matrix factorization” methods (a common method of semi-definite programming based on exact linear algebraic methods to approximate an optimal solution) ultimately be the answer? Probably not. This is because the brain must also work efficiently in moments when the information is noisy, distorted or when circumstances and information change dynamically. It would also save resources to be able to use the same system for abstract distances such as numbers on a ray of numbers, times and other things. So we are talking about a system that, as is often the case in biology, functions flexibly in a robust manner and is based on first principles.

Remember, here the leading question is: How can the brain easily estimate distant distances based on local dependencies? That would be nice to know, because then we could use the same procedure to geometrically display the cognitive map in the animal’s head on our screens. In fact, the previously mentioned methods of algebraic-topological data analysis come quite close [122, 123, 124, 125, 126, 127, 121]. However, these studies focus primarily on new approaches for the investigation of neural networks. How the brain itself might perform that is ignored. Similarly, the role of grid cells, which may be involved in the biological implementation of such a method, is rarely discussed in this context.

The role of grid cells: A hypothesis I would like to suggest another, very pragmatic method, that astonishingly is not mentioned anywhere in the literature and is based on a

very simple principle. We start with a number N of points randomly scattered into a fixed geometry. Each point (place field) is associated with a place cell. The typical distance $\Delta(K)$ of nearest neighbors of points decreases with the number of points K (while density increases). If we know the geometry of the area in which we scatter the points as well as their number n , then we can estimate Δ . This almost banal information can be exploited.

If we do not scatter the N points completely randomly but with mutual repulsion (or also by Poisson disc sampling) this functional connection between N and Δ becomes very reliable. There are many ways to obtain such point distributions, including clustering (see Fig. 4.1D for an example of $\Delta(K)$ obtained via clustering). If the brain tries to categorize its environment into “places”, it may also avoid excessive redundancies. I assume that groups of several place cells that are linked to the identical spatial memory rarely occur in the same network. If that is not true $\Delta(K)$ is worse in predicting next nearest neighbor distances.

The dependency of Δ on the number of points can be utilized as follows: We start out by selecting a set $k = \{i : i \in \{1, \dots, N\}, \#k = K \ll N, \sum_{i \in k, j \in k}^{i \neq j} |C_{i,j}| = 0\}$ of K row indices i with $K \ll N$ in the $N \times N$ correlation matrix C such that $C_{i,j} = 0$ for all $i \in k$, $j \in k$ and $i \neq j$. Then we can set the values $D_{i,j} = \Delta(K)$ in a distance matrix D to be estimated, which gives a first estimate of few previously unknown distances. That is not yet very satisfying, because we have only set a fraction of the entries in D ; therefore we repeat this step sufficiently often. But we still can’t be satisfied, because not all pairs in k will have a distance near Δ to each other. In particular, many have a distance equal to a multiple of Δ . Therefore we repeat the whole procedure in the next step; this time with a smaller K and thus with a larger $\Delta(K)$. While we continue to reduce K , we overwrite existing entries in D by the least common multiples of the existing values and the newly suggested ones again and again. It is advisable to stop at $K = 2$ at the latest, as it is well known that a relationship always requires two partners. This algorithm is still quite inefficient, it does not care about the triangle inequality and is anything but precise. But given the complexity of the problem it is very simple and quite useful as a basis. Simple modifications prevent redundancies, at small K the direct neighbors of the nodes k can also be overwritten and the whole can be structured hierarchically. In particular, the subset k to an iteration step can be selected from the set k of the previous iteration step and so on. It is unclear how precise distance information is represented in mammalian brains and if they are subject to the triangle inequality. The approach heavily relies on the existence and reliability of the function $\Delta(K)$. The sketched algorithm itself is “blind to modality” (how hippocampal researcher Geörgy Buzsáki would say) and does not require the notion of euclidean space.

Somewhat more biological this could probably be implemented relatively simple with a multi-layered neural network with lateral inhibition and layers that become smaller (analogous to the smaller becoming K), so that it implements a hierarchical clustering and thereby a compression. Such a kind of clustering belongs to the natural functions of a neural network [158]. The weights of the network would only have to be learned when the input correlations change. Distances could be read out via the propagation depth after activation of input neurons i and j . Once the network is setup distances could be estimated

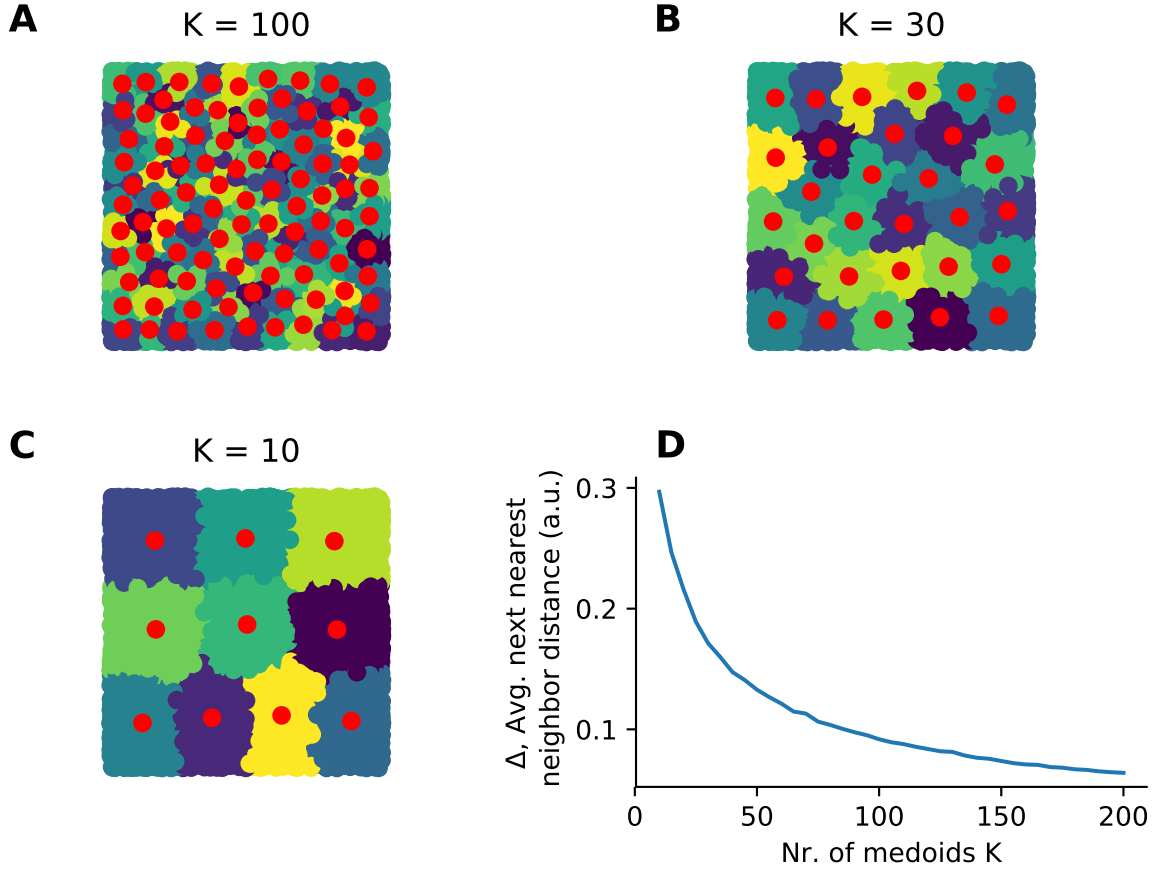


Figure 4.1: An illustration for the effect of different number k of clusters on the average next nearest neighbor distance Δ . Initially $N=5000$ points (place field centers) were scattered randomly within a square area. From that set of points K cluster centroids were computed using the kMeans++ algorithm. A) $K = 100$ cluster centers (red) and respective cluster members (color coded). B) Same as in A) but with $K = 30$. C) Same as in A) and B) but with $K = 10$. D) Δ as a function of the number of clusters K .

on the fly. It should be noted that clustering of a large number of uniformly distributed points in 2D often results in hexagonally arranged clusters.

Following these thoughts, I suggest that the role of grid cells is not to signal the current position of the animal, but that these cells are involved in a clustering based embedding machinery. In that view, each grid cell is just a hub connecting centroids of one clustering.

Specifically, I claim that the subset of grid cells linking a cluster centroid i (for example a place cell) and centroid j (another place cell) gets active whenever relationship information of node i to j is requested. Typically a rodent might relate its own location i to some other location j (for example goals, obstacles or environmental boundaries). Therefore, a grid cell happens to be most active at particular locations in space. In this view the MEC is a static machinery that provides distance information in arbitrary spaces to downstream

systems. This is in agreement with the findings in Chapter 2.

This hypothesis predicts that inactivation of a grid cell module results in impaired distance estimation on the corresponding spatial scale. The most simple hierarchical scenario suggests a ratio between scales of $\sqrt{3} \approx 1.7$ in combination with a change in grid rotation of $\pi/6$. Furthermore, it is predicted that a corresponding subset of grid cells is also active when distances between two remote places are estimated even when the animal is not moving.

The implementation of the described algorithm in a biologically inspired neural network would be an interesting project for future research.

Grid cells reflect rather edges than nodes!

4.4 Conclusion

In summary, during my doctoral thesis I was able to identify a characteristic motif in the spike count distributions of neurons involved in the spatial navigation of rodents. These observations explain a large part of the spike count variability from experiment to experiment and have not been reported so far. It is important that the motif can be traced back to a measurable quantity: Shifts in the positions in space where such cells tend to signal the position of the animals. In addition, these drifts were correlated across neurons. Thus, a source of neuronal noise correlations was identified that supports a newly emerging and fundamentally different perspective on spatial coding in the entorhinal cortex. In the second part, we found depolarizing afterpotentials in the membrane potential of spatially encoding neurons that have not been shown to behave in mice before. Furthermore, we could show that this phenomenon correlates with the cell's burst properties.

With these results, I hope to contribute to the understanding of neural computations utilized by mammals to orient themselves in natural environments.

Appendix A

Additional information on Fenton's variability measure σ_z^2

A.1 The relation of variability σ_z^2 to chi-Square testing

The variance of the z-scores is proportional to the Chi-Square statistic as $\sigma_z^2 = 1/N \cdot \sum_i ((s_i - n_i)^2 / n_i) = 1/N \cdot \chi^2$. Therefore by comparing the value of $N \cdot \sigma_z^2$ to a Chi-Square distribution a p-value can be computed. For Poisson the number of degrees of freedom is equal to the number of estimated spike count expectations minus the number of parameters of the firing rate map estimation minus one. The reduction by one is due to the constraint on the overall firing rate. For example, a z-score variance higher than 1.15 indicates statistically significant deviation from Poisson in the case of a 20 minute recording (divided into 5s time windows, gives 300 categories) and a significance level $\alpha = 5\%$. This is because the 5 percent quantile of a chi-square distribution with 300-3 degrees of freedom is at about $1.15 \cdot 300$. Dependencies between the expectation values in each time bin further decrease the number of degrees of freedom. In that case the percentile is reduced as $df = n_\Omega - n_P - n_D - 1$, where n_Ω denotes the number of possible outcomes, n_P is the number of parameters to be estimated and n_D is the number of dependencies between the parameters. In other word, given any firing rate model leading to the expected counts n_i in time windows i with a known number of degrees of freedom, we can define a significance threshold on σ_z^2 given a spiking distribution like Poisson or some other count distribution.

A.2 The upper bound on variability and its relation to the zero inflated Poisson parameter

The minimal MSE estimator as well as the MLE estimator of the expectation value of the underlying yet unknown distribution of data s is given by its mean. The average squared residual to the mean is minimal and is equal to the variance of s. Therefore the most naive

and model free result for the variance of z-scores as defined above is given by:

$$\text{var}(z) = \text{var}\left(\frac{s - \langle s \rangle}{\sqrt{\langle s \rangle}}\right) = \frac{\text{var}(s)}{\langle s \rangle} = \frac{\text{var}(s)}{\bar{s}} \quad (\text{A.1})$$

which is equal to the fano factor if s are measured across trials.

Given a second series of observables $\{\phi\}$ (i.e. the signal/stimulus like positions, speed etc.) from a unique and finite set Φ of all possible stimulus combinations we can compute the conditioned expectation to one specific stimulus combination Φ_i which is given by the mean of all data s_t where $\phi_t = \Phi_i$:

$$n_{\Phi_i} := \langle s | \Phi_i \rangle = \frac{1}{N_{\phi_t = \Phi_i}} \sum_t s_t \delta_{\phi_t, \Phi_i} \quad (\text{A.2})$$

For each of these stimulus combinations the minimal MSE estimator again is the mean but this time only of the respective subset of the data. The variance $\text{var}(s)$ of s can be divided into a sum of sums over members of the subsets:

$$\text{var}(s) = \frac{1}{N} \sum_t (s_t - \bar{s})^2 = \frac{1}{N} \sum_i^{|\Phi|} \sum_t^N (s_t - \bar{s})^2 \delta_{\phi_t, \Phi_i} \quad (\text{A.3})$$

Note that the right part is not a sum over conditional variances but conditional mean squared errors. Knowing $MSE(x|y) \geq \text{var}(x)$ we can write:

$$\frac{1}{N} \sum_i^{|\Phi|} \sum_t^N (s_t - \bar{s})^2 \delta_{\phi_t, \Phi_i} \geq \frac{1}{N} \sum_i^{|\Phi|} \sum_t^N (s_t - \langle s | \Phi_i \rangle)^2 \delta_{\phi_t, \Phi_i} \quad (\text{A.4})$$

$$\frac{1}{N} \sum_i^{|\Phi|} \sum_t^N (s_t - \bar{s})^2 \delta_{\phi_t, \Phi_i} \geq \frac{1}{N} \sum_i^{|\Phi|} N_{\phi_t = \Phi_i} \text{var}(s | \Phi_i) \quad (\text{A.5})$$

Now we assume s being Gamma distributed within a subset (meaning $\text{variance} = a \cdot \text{mean}$) and replace the MSE (the inner sum on the left, for which we know that it is larger than the variance of the subset) by a linear function of the subsets mean:

$$MSE_i := (a_i + \epsilon_i) \langle s \rangle \geq a_i \langle s | \Phi_i \rangle \quad (\text{A.6})$$

for $\epsilon_i \geq 0$ and $a_i > 0$.

$$\frac{1}{N} \sum_i^{|\Phi|} N_{\phi_t = \Phi_i} ((a_i + \epsilon_i) \langle s \rangle) \geq \frac{1}{N} \sum_i^{|\Phi|} N_{\phi_t = \Phi_i} (a_i \langle s | \Phi_i \rangle) \quad (\text{A.7})$$

Then for the variance-to-mean ratio we get:

$$\text{var}(z) = \frac{\text{var}(s)}{\bar{s}} = \frac{1}{N} \sum_i^{|\Phi|} N_{\phi_t = \Phi_i} \frac{(a_i + \epsilon_i) \langle s \rangle}{\langle s \rangle} \quad (\text{A.8})$$

$$= \frac{1}{N} \sum_i^{|\Phi|} N_{\phi_t = \Phi_i} (a_i + \epsilon_i) = \langle a \rangle + \langle \epsilon \rangle \quad (\text{A.9})$$

and for the subset-wise z-score variance:

$$var(z|\phi) = \frac{1}{N} \sum_i^{|\Phi|} N_{\phi_t=\Phi_i} \frac{a_i \langle s|\Phi_i \rangle}{\langle s|\Phi_i \rangle} = \langle a \rangle \quad (\text{A.10})$$

Therefore

$$var(z) \geq var(z|\phi) \quad (\text{A.11})$$

holds for data where a zero mean implies a zero variance like inh. Poisson count data or more generally Gamma distributed data. In other words for any data that can be grouped such that within the groups the variances are smaller than the mean squared errors to the overall mean a model can be found for the expectation values such that the variance of z-scores with respect to that model drops (i.e. *variability is explained*).

The set of all $\langle s|\Phi_i \rangle \forall i \in \{1, \dots, |\Phi|\}$ estimated via MLE or MMSE may be denoted as *informative model*. Such an informative model allows us to reduce the variance of the z-scores. When assuming there exists an informative model than the inequality (A.11) allows to specify an upper bound on the zero-inflation parameter α without knowing the informative model or the respective mappings $\phi_t \rightarrow \Phi_i$ or the stimulus at all.

The left side of eqn. (A.11) describes the variability of data under homogeneity assumption whereas the right side refers to situations in which a model about the inhomogeneity is available. For example, the measured variability in spike counts of a grid cells will always be lowered when not being ignorant to the presence of spatial modulations of the firing. This is not surprising but comes with a useful implication for estimating the parameter α of a zero inflated model as described in the following paragraph.

Given a informative sequence of expectation values $\{n\} \in \mathbb{R}_+^N$, which we will refer to as informative model, and a series of measured counts $\{s\}$ we can now ask what is the value of α . By informative model we are referring to models where the variability is smaller or equal than the variance-to-mean ratio:

$$\sigma_z^2 \leq var(s)/\bar{s} \quad (\text{A.12})$$

and thereby all models that hold eqn. (A.11). To predict values of $\sigma_z^2(\{n\}, \alpha)$ we first derive an analytical expression for the variance $\sigma_{z,i}^2(n_i, \alpha)$ of z-scores in one bin i (in which we can assume the process to be homogeneous) and demonstrate that the variance σ_z^2 taking all N bins i into account satisfies $\langle \sigma_{z,i}^2(n_i, \alpha) \rangle_{i \in \{0, \dots, N\}} = \sigma_z^2(\bar{n}, \alpha)$. This allows us to define an upper bound for α .

The theoretical variance of the z-score distribution for a ZIP process is

$$\begin{aligned} \sigma_{z,i}^2(n_i, \alpha) &= \sum_s (z(s, n_i) - \langle z(s', n_i) \rangle_{s'})^2 P(z) = \sum_s \left(\frac{s - n_i}{\sqrt{n_i}} - 0 \right)^2 P(s|n_i, \alpha) = \\ &= n_i \left(\alpha + (1 - \alpha) e^{-\frac{n_i}{1-\alpha}} \right) + \sum_{s=1}^{\infty} \frac{(s - n_i)^2}{n_i} \cdot \frac{(1 - \alpha)}{s!} \cdot e^{-\frac{n_i}{1-\alpha}} \cdot \left(\frac{n_i}{1 - \alpha} \right)^s = \\ &= 1 + \frac{n_i \alpha}{1 - \alpha} = \frac{\sigma_S^2}{n_i} \quad (\text{A.13}) \end{aligned}$$

with $\sigma_S^2 = n_i + \frac{\alpha}{1-\alpha} \cdot n_i^2$. For N independent repetitions with different expectations we can compute the variance being the variance of a mixed distribution:

$$\sigma_z^2 = \frac{1}{N} \sum_i^N \sigma_{z,i}^2(n_i, \alpha) = \langle \sigma_{z,i}^2(n_i, \alpha) \rangle_{i \in \{0, \dots, N\}} = 1 + \bar{n} \cdot \frac{\alpha}{1-\alpha} = \sigma_z^2(\bar{n}, \alpha) \quad (\text{A.14})$$

which is equal for the homogeneous case ($n_i = \bar{n}$) and the inhomogeneous case (where n_i can vary) if the underlying process is a ZIP process. Here we used that the average of z-transformed variables is zero by definition. Note that this is not contradicting our definition of informative models as here we computed the mean of the bin-wise variances and did not compute the variance of observed z-scores across the bins where the bin-wise variance is unknown.

We can thus estimate α by solving

$$\hat{\alpha} = \left(\frac{\bar{n}}{\sigma_z^2 - 1} + 1 \right)^{-1} \quad (\text{A.15})$$

with $\sigma_z^2 = \sigma_{z, \text{empirical}}^2$. Without any knowledge about n_i eqn. (A.15) provides an upper limit on $\alpha(\{n\})$ via $\bar{n} = \bar{s}$ for an informative model $\{n\}$. This is true because for the empirical variabilities it holds $\sigma_{z, \text{emp.}}^2(S, \{n\}) < \sigma_{z, \text{emp.}}^2(S, \bar{s}) = \sigma_s^2/\bar{s}$ as we postulated above.

With $\sigma_z^2 = \sigma_z^2(S, \bar{s})$ eqn. (A.15) is equal to

$$\hat{\alpha}(S) = \left(\frac{1}{CV^2 - 1/\bar{s}} + 1 \right)^{-1} \quad (\text{A.16})$$

with $CV = \sigma_S/\bar{s} = \sigma_S/\bar{n} = \sigma_z^2(S, \bar{s})/\sigma_S$.

It is worth to mention that the requirement for informative models can be met even for models with high likelihood that are not strictly informative via restricting the analysis to regions where the ratio between s and n is not too large (usually when $n \geq 1$). In the perspective of studying zero inflation this is natural to do because low expectations are superimposing excess zeros. The inflation probability α could be arbitrarily overestimated when taking data into account where there are no or only rarely spikes expected. In that scenario there is no chance to find out whether zero counts are thrown due to low expectation values and high or low zero inflation.

We conclude that by measuring the variance and the mean of spike counts across time windows or trials an upper limit for α can be estimated without model fitting.

Appendix B

The Hidden Markov model and its relationship to the ZIP model

The HMM is often used to describe multi-state activities with latent variables. For Markovian processes it is assumed that the probability $p(z_t|z_{\{1,\dots,t-1\}})$ of the system being in a state z_t at time-point t solely depends on the state z_{t-1} in the previous moment $t - 1$:

$$p(z_t|z_{\{1,\dots,t-1\}}) = p(z_t|z_{t-1}) \quad (\text{B.1})$$

The probabilities for these momentary transitions are usually described by a state transition matrix $T_{i,j} = p(z_t = i|z_{t-1} = j)$. An HMM entails hidden states, which govern the probability of observing the output. In other words given a series of measurements $\{s\}$ the corresponding series of states $\{z\}$ is unknown as the state z_t is not determined by outcome s_t . Furthermore it is assumed, that an outcome s_t solely depends on the parameters in that moment which are themselves assumed to be independent of the past:

$$p(s_t|\{z_0, z_1, \dots, z_N\}, \Theta_t) = p(s_t|z_t, \Theta_t) \quad (\text{B.2})$$

with parameters Θ_t describing the hidden state (which are not shown in the following equations for better readability).

The emission probability for a sequence of data $\{s\}$ is the product of the probabilities of observations over time summed over all possible state sequences:

$$P(\{s\}) = \sum_{\forall \{z\} \in |Z|^N} \prod_{t=1}^N p(s_t|\{z_0, z_1, \dots, z_N\}) p(z_t|z_{\{1,\dots,t-1\}}) \quad (\text{B.3})$$

with Z being the set of all possible states and N the number of observations. The summation goes over all possible state series $\{z\}$. Together with the Markov assumption and the assumption on the hidden states we can write

$$P(\{s\}) = \sum_{\forall \{z\} \in |Z|^N} \prod_{t=1}^N p(s_t|z_t) p(z_t|z_{t-1}) \quad (\text{B.4})$$

$$= \sum_{\forall \{z\} \in |Z|^N} \prod_{t=1}^N p(s_t|z_t) \prod_{t=1}^N p(z_t|z_{t-1}) \quad (\text{B.5})$$

This formulation is the starting point to the inference algorithms for Hidden Markov Models (Forward/-Backward-Inference, Baum-Welch and Viberti).

For a two state system with no dependency on the previous time step the elements of that matrix are independent of the first index denoting the past state ($T_{00} = T_{10}$ and $T_{01} = T_{11}$). For the ZIP model these entries become $T_{00} = T_{10} = \alpha$ and $T_{01} = T_{11} = 1 - \alpha$. It is worth to mention that the ZIP model is only semi-hidden as for $s_t > 0$ the state at time t is determined to be $z_t = 1$.

As for the the ZIP transition matrix as described before there is no temporal dependency we can iteratively rewrite (B.4):

$$P(\{s\}) = \sum_{\forall \{z\} \in |Z|^N} \prod_{t=1}^N p(s_t|z_t)p(z_t|z_{t-1}) \quad (\text{B.6})$$

$$= p(s_1|z_1=0)p(z_1=0) \sum_{\forall \{z\}_{N-1}} \prod_{t=2}^N p(s_t|z_t)p(z_t) \quad (\text{B.7})$$

$$+ p(s_1|z_1=1)p(z_1=1) \sum_{\forall \{z\}_{N-1}} \prod_{t=2}^N p(s_t|z_t)p(z_t) \quad (\text{B.8})$$

$$= \dots = \prod_{t=1}^N (p(s_t|z_t=0, \Theta_t)p(z_t=0) + \quad (\text{B.9})$$

$$p(s_t|z_t=1, \Theta_t)p(z_t=1)) \quad (\text{B.10})$$

Using $P(z_t=0) = T_{00} = T_{10} = \alpha$ and $P(z_t=1) = T_{01} = T_{11} = 1 - \alpha$, $p(s_t > 0|z_t=0, \Theta_t) = 0$, $p(s_t=0|z_t=0, \Theta_t) = 1$ and $p(s_t|z_t=1, \Theta_t) = \text{Poisson}(s_t|\Theta_t)$ we get:

$$P(\{s\}) = \prod_{t|s_t=0} (\alpha + (1 - \alpha)e^{-\mu_t(\Theta_t)}) \cdot \prod_{t|s_t>0} ((1 - \alpha) \frac{\mu_t(\Theta_t)^{s_t}}{s_t!} e^{-\mu_t(\Theta_t)}) \quad (\text{B.11})$$

which is exactly the probability of a zero inflated Poisson process like described in the following session. Due to the partially observed nature of the described HMM we can directly compute this likelihood without require to perform the EM procedure. Note that there already exists applications of combining HMMs and ZIP in different ways [159, 160, 161]. It is also worth mentioning that there are multiple of possible variations like non-static transition probabilities (that is non-static zero-inflation). The previous paragraph did not make any assumptions on the model $\mu_t(\Theta_t)$, only for Poisson and other count distributions it is restricted to \mathbb{R}^+ . Furthermore we like to mention that in this formulation there exists no temporal structure in the state occurrence. This can easily be changed by allowing $T_{00} \neq \alpha$ and $T_{11} \neq \alpha$, but it shows to be useful to first compute α interpreted as the overall ratio of time in which the zero-state is occupied which we will demonstrate in the following paragraphs.

Continuous zero inflation: From observations we suspect the zero inflation to be temporally extended as we observe long pauses of spiking (and that these pauses seem to be correlated across ensembles of neurons in the region). Neither spatially modulated Poisson nor negative Binomial spiking is likely to generate long lasting pauses in the order of a second without further modifications and being independent on the covariates. As a zero inflation in biological systems is unlikely to have no temporal structure and extended pauses are observed we use the parameters of the ZIP model to restrict a more realistic two-state HMM model. This is an alternative approach to the one used in generalized zero inflated Poisson models (gZIP, see Giles et al. [162]) as here the temporal structure of the zero inflation is not determined by dependencies on covariates. The number of parameters therefore is dramatically less than in gZIP with time varying zero inflation.

For any transition matrix T it is required that

$$\sum_j T_{ij} = 1 \forall i \quad (\text{B.12})$$

hence $T_{00} + T_{01} = T_{10} + T_{11} = 1$. The expected duration d_i of a state i in the two-state system is given by $d_i = 1 + (1 - T_{ii}) \sum_n n \cdot T_{ii}^n = (1 - T_{ii})^{-1}$. Starting from this we can ask how the transition probabilities need to look like when considering desired average state durations d_0 and/or d_1 in the binary scenario and get:

$$T_{ii} = \frac{d_i - 1}{d_i} \quad (\text{B.13})$$

The remaining probabilities $T_{ij|i \neq j}$ can be computed via eqn. (B.12). Thus we get the full transition matrix for continuous zero inflation:

$$T = \begin{bmatrix} \frac{d_0-1}{d_0} & \frac{1}{d_0} \\ \frac{1}{d_1} & \frac{d_1-1}{d_1} \end{bmatrix}$$

How can we now make use of the parameter α which is easy to access from the data? The steady stationary distribution of state j (describing the total average ratio of the system occupying state j) is given as follows:

$$\pi_j = \sum_i \pi_i T_{ij} \quad (\text{B.14})$$

Together with $\sum_{i \in \{z\}} \pi_i = 1$ for our system we get:

$$\pi_1 = \frac{T_{01}}{T_{01} + T_{10}} \quad (\text{B.15})$$

$$\pi_0 = 1 - \pi_1 \quad (\text{B.16})$$

Equations (B.13) and (B.15) serve as a link between the state durations and the zero inflation probability when setting $\pi_0 = \alpha$. We arrive at an equation for d_0 , d_1 and α :

$$d_1 = \frac{d_0}{\alpha} \quad (\text{B.17})$$

with $d_1 \geq 1$, $d_0 \geq 1$ and $0 < \alpha \leq 1$. By measuring the amount of zero-inflation $\pi_0 \rightarrow \alpha$ together with the equations above we can therefore restrict the possible entries of the transition matrix based on any average zero-state duration d_0 :

$$T = \frac{1}{d_0} \begin{bmatrix} d_0 - 1 & 1 \\ \alpha & d_0 - \alpha \end{bmatrix}$$

The calculation based on d_1 instead of d_0 is analogous and not shown here.

In short knowing α as well as the average duration of one of the two states the Markov system is fully described and no fitting of the entries of T is required. In combination with a Poisson emission model for state 1 this two state model can account for the overdispersion we observe in grid cells as well as for the temporal correlations in the spike trains but fails to model spatial jitter and noise correlations across neurons.

Appendix C

Tuning displacement in 2D recordings

Several approaches have been studied to measure the dynamics of grid phase and single grid field locations in 2D recordings in the open field: For example Hardcastle et al. [119] suggested a “spike distance metric” that compares inbound and outbound spiking during firing field crossings. Hägglund et al. [163] use a flow map to study the dynamics of grid fields in response to changing the shape of the arena. Such attempts are fraught with difficulties, however. Grid cells often fire at low rates, even at the center of a firing field. Given relatively low firing rates, the firing fields only become apparent after the animal’s trajectory has passed through each location multiple times, which takes several minutes at the very least. So 2D tracking of the precise position of firing fields across time remains difficult.

Therefore, we decided to build an explicit stochastic model for the spike count distributions in the presence of field jitter or drifting. This model is sketched in Fig. C.1; details of the model can be found in appendix D. The resulting spike counts obey the beta distribution. In the following section, we show how the two-parameter distribution, often treated in the literature as an abstract and purely theoretical distribution, can be interpreted as arising from a concrete, biophysical model. To our knowledge, there are no other examples of such a derivation in the scientific literature.

In short, we assume that the position of a firing field itself is a random variable. We further assume that the shape of the firing field is a non-negative power function bounded at zero and consider the jitter density to be power of the field shape. Both functions, for example, could be negative parabola shifted to positive maxima. That approximates Gaussian bumps. Following these assumptions, the expected number of spikes at a given position is a beta-distributed random variable. Given the expected number, the spike distribution itself is Poisson, so that the overall spike count distribution is a Beta-Poisson mixture (BP). Such a model has three parameters. It can capture both zero inflation and the effect of gain modulations that lead to classical overdispersion in the spike statistics. In particular, this model can produce the frequently observed bimodality of experimental spike count distributions. Note that in all except one of the examples shown here the jitter

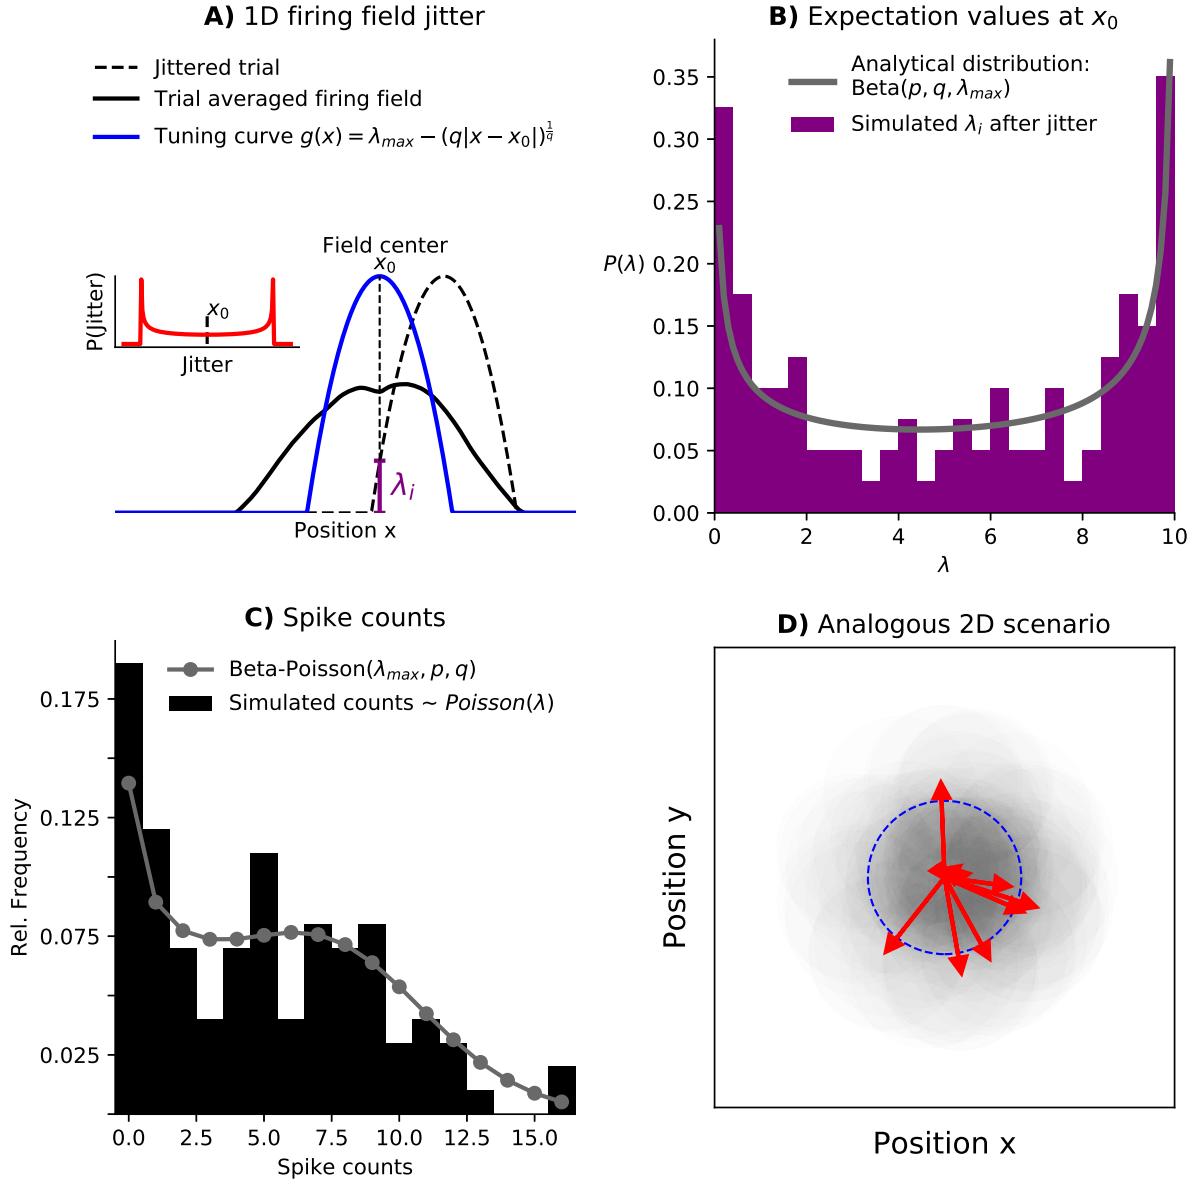


Figure C.1: A jitter model leading to the Beta-Poisson spiking model. **A:** Simulated 1D trial displacement, drawn from the jitter distribution in the inset ($P(jitter) \propto g(x)^{p-1}$). The trial-averaged firing field (solid black line) is broader and has a lower peak rate as compared to the underlying tuning curve without jitter (solid blue line, parameters used: $\lambda_{max} = 10, p = .6, q = .5$). **B:** The resulting distribution of λ at x_0 across $N=100$ trials and the analytical pdf of the Beta distribution. **C:** Simulated spike count distribution from a Poisson process with parameters λ_i where i denote the indices of the trials and the resulting Beta-Poisson distribution as derived in the Appendix D. **D:** Analogous situation in 2D with random shift angles leading to the identical statistics.

density is bimodal. This suggests firing fields that are randomly switching between two typical locations for these cases. However, most time fits resulted in unimodal distributions or rather flat distributions and we observed no stereotypical shape of the jitter.

The BP model connects the spike count statistics from the open field to the field shape and the spatial jitter distribution. The measures estimated for two-dimensional arenas can then be corroborated against the same measures on the linear track, as shown in Fig. C.3; on the linear track, we have direct access to the jitter in field positions from trial to trial.

While fitting multiple models to the spike count data, we observed that the skewness in the BP model’s firing rate distribution is highly correlated with the zero inflation parameters of the ZIP model as well as the ZINB model, both for the 1D and the 2D data ($r = .68$, $p \ll 1e - 10$ and $r = .63$, $p \ll 1e - 10$, see Table C.1). That is, because for a positive skew in the rate distribution a highly zero-inflated regime regime is approached. On the contrary, a negative skew shifts the mode to larger values. Extremely negatively skewed fits suggest that nearly all rates in the field center are equal. In that case resulting spike counts are close to Poisson.

	Δ_{ζ}	Σ_{1D}	Σ_{2D}	α_{1D}	α_{2D}
Δ_{ζ}	1	.42	.35	.27	.34
Σ_{1D}	.42	1	.55	.70	.57
Σ_{2D}	.35	.55	1	.31	.63
α_{1D}	.27	.70	.31	1	.54
α_{2D}	.34	.57	.63	.54	1

Table C.1: Correlations between the skewness Σ as a function of the parameters of the BP model fits, the ZIP model parameters α and the difference Δ_{ζ} of peak amplitudes in spike-triggered firing maps for short and long delays. The numerical values are the Pearson correlation coefficients ($p < .01$ for all pairs). The 1D linear track measures shown in the table are the medians over contexts and running directions. Data is shown for $N = 99$ grid cells for which all the required variables could be measured both in 1D and 2D.

The BP model is preferred over the Poisson model in 83% (likelihood ratio test, $p < .001$, $df = 2$) and to the ZINB model in 61% (by direct comparison of the likelihoods) of the cell recordings. For each cell, the BP model fits in 1D and 2D tended to yield similar parameters, which is consistent with firing field drifts occurring in 2D as well as in 1D. To address this question directly, we analyzed the spike-triggered rate maps [5] of grid cells over fixed time windows. In this approach, the animal’s location at each spike is set as the origin of a standard coordinate system. Spikes that occur within a specific time window *after* the triggering spike are used to build a firing rate map relative to the origin defined by the trigger. These relative firing rate maps are then averaged and collapsed onto a single dimension, which corresponds to a radial distance. If the spatial grid is stable, then the radial spike-triggered firing rate map will have peaks at regularly occurring intervals given by the grid’s period. On the other hand, if a cell’s fields undergo slow field displacement, the peaks in the radial-distance will gradually smear out. So we computed the spike triggered

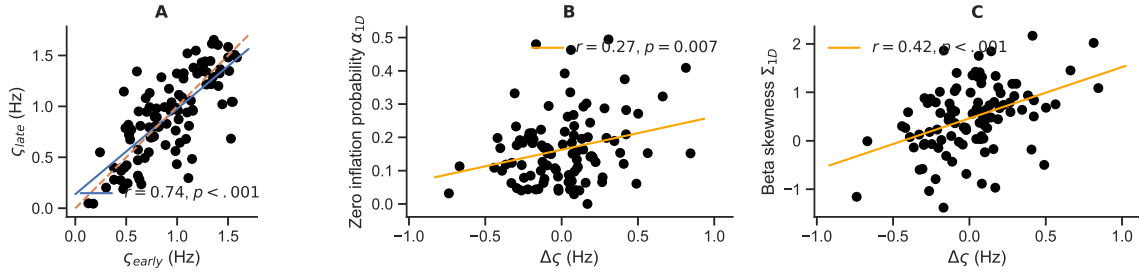


Figure C.2: Scatter plot of measures for the quantification of non-stationary tuning curves. **A:** Amplitudes ζ in the spike-triggered firing rate map at the grid period measured at long delays (up to 15s after trigger) versus short delays (between 60s and 75s after trigger). **B:** Zero count excess (measured via α from ZIP fits on the linear track) against rate change $\Delta\zeta$ at the grid period in the spike-triggered rate map in the open field. **C:** Skewness Σ_{1D} of the mixing distribution computed from BP model fits on the linear track against rate change $\Delta\zeta$. The correlations in **B** and **C** indicate that features of spike count distributions in one context predict drifts of firing fields in a different context.

firing rate maps for successive 15-second intervals after the trigger. If drift occurs, the spatial modulation in the spike-triggered map during later intervals will be more strongly blurred. In particular, we compared the maps for the first 15s after each spike to the maps in the time window from 60s to 75s after each spike.

The amplitude ζ_{early} was quantified as the difference between the first local maximum (grid period) and the average of the minima to the left and to the right of that maximum. The positions of the extrema were detected for the first 15s window. As the peak did not always persist in the later window, the amplitude ζ_{late} for the 60s to 75s window was measured using the positions of the extrema determined for the early window. If ζ_{late} is close to zero, it means that the peak is smeared out, whereas a ζ_{late} that is larger than ζ_{early} means the peak has become more prominent. Therefore, large differences $\Delta\zeta = \zeta_{early} - \zeta_{late}$ indicate a change in the typical relative distances between spike positions on the scale of a minute. While on average $\Delta\zeta$ was close to zero ($mean = .014$, $sem = .028$, Wilcoxon-test: $U = 2438$, $p = .9$, $N = 99$) in fact a significant correlation was found between $\Delta\zeta$ and the skewness of the fitted mixing distributions (see Table C.1). As we pointed out earlier, positive skewness can imply strong jitter. Consequently, jointly drifting firing fields indicate both larger $\Delta\zeta$ as well as positive skewness.

Thus, two very different measures to quantify the non-stationarities of spatial tuning turned out to be correlated. This underscores how grid field displacements yield distinct, but related signatures in the statistics of grid cell spiking and demonstrates the utility of the BP model fits. Such drifts in the spatial representation occur even in cases where trial-averaging yields a clear grid structure to the firing fields. We deduce that such non-stationarities are a major source of trial-to-trial variability in grid-cell spiking in the open field.

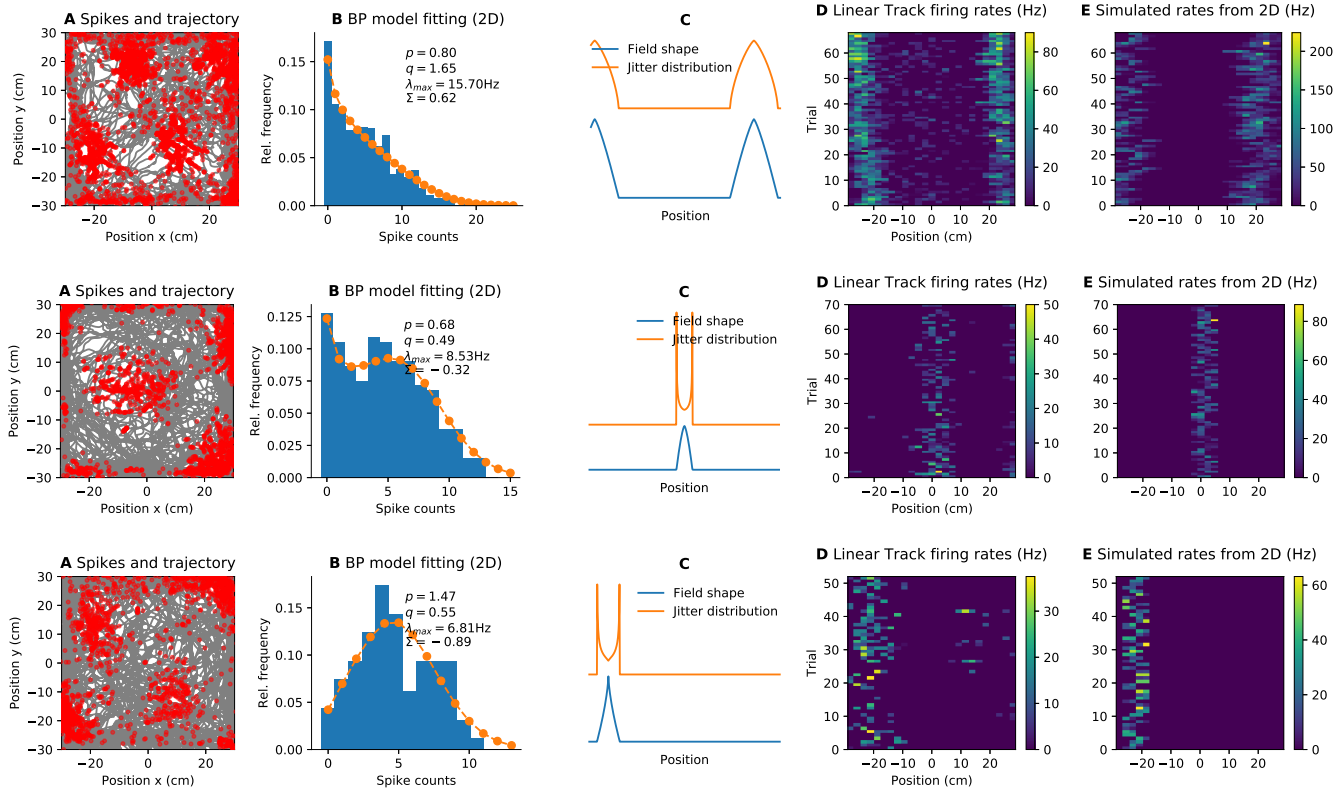


Figure C.3: Each row represents data and model simulation of a different grid cell from the Pèrez-Escobar et al. data [150]. **A:** Spikes (red) and trajectory (grey). **B:** Collected spike count distribution ($n = 5$, $F_{min} = 5\text{Hz}$) and fitted BP distribution with parameters (p , q , λ_{max}) and the skewness Σ of the mixing distribution. **C:** Field shape and jitter distribution sketched as suggested by the fitted parameters (copied to each field center). One displacement was drawn per trial to achieve consistent field drift in the case of multiple fields on the track. The field centers were estimated as the positions of local maxima in the tuning curves of the linear track (threshold for local maximum: $3Hz$). In the last two rows a bimodal field displacement is suggested from the jitter distribution. **D:** Firing rates across trials on left runs in the light condition 1. **E:** Simulated firing rates from the distributions in **C**.

Appendix D

Jitter transformation and the Beta-Poisson model

Let us treat the spatial jitter X as a continuous random variable. We describe the intrinsic field shape (before the position of the field is jittered) using the tuning function $g : \mathbb{R} \rightarrow \mathbb{R}^+$, which represents the expected number of spikes at positions x . Randomly jittering the field's position makes the firing rate at any position a random variable, as well. We can transform from the probability distribution of X to the distribution of $\Lambda = g(X)$ as follows:

$$P(\lambda)d\lambda = P(x)dx \quad (\text{D.1})$$

$$P(\lambda)d\lambda = P(g^{-1}(\lambda))dx \quad (\text{D.2})$$

$$P(\lambda) = P(g^{-1}(\lambda)) \left| \frac{dg^{-1}(\lambda)}{d\lambda} \right| \quad (\text{D.3})$$

Now let us assume that $g(x)$ is a power function of the following form:

$$\lambda = g(x) = \lambda_{max} - (q|x - x_0|)^{\frac{1}{q}}, \quad q > 0, \quad 0 \leq \lambda \leq \lambda_{max} \quad (\text{D.4})$$

For example, with $q = 1/2$ the function $g(x)$ becomes an inverted parabola and, therefore, approximates a Gaussian bump with height λ_{max} to first order. Similarly, it approximates any bell-shaped curves, for example, resulting from diffusion of Brownian particles. Its inverse function is:

$$x = g^{-1}(\lambda) = \frac{1}{q}(\lambda_{max} - \lambda)^q + x_0 \quad (\text{D.5})$$

$$\frac{dg^{-1}(\lambda)}{d\lambda} = (\lambda_{max} - \lambda)^{q-1} \quad (\text{D.6})$$

and for $P(\lambda)$ we get:

$$P(\lambda) = P(x) \cdot (\lambda_{max} - \lambda)^{q-1} \quad (\text{D.7})$$

Assuming that the jitter distribution is similar to the field shape $g(x)$ up to a normalization constant and a power of $(p - 1)$ with $p > 0$, that is $P(x) \propto (\lambda_{max} - (q|x - x_0|)^{\frac{1}{q}})^{p-1}$, we get:

$$P(\lambda) = C_1 \cdot \lambda^{p-1} \cdot (\lambda_{max} - \lambda)^{q-1} \quad (D.8)$$

The normalization constant C_1 is the Beta function $\beta(p, q, \lambda_{max})$, while the resulting distribution is a Beta distribution with support $[0, \lambda_{max}]$. Its expectation value is $\langle \lambda \rangle = p/(p + q) \cdot \lambda_{max}$.

Following our assumption that λ describes an expectation value for spike counts s it comes in handy to use $P(\lambda)$ as a mixing distribution for a doubly stochastic spike count distribution

$$P(k|p, q, \lambda_{max}, \Theta) = \int_0^{\lambda_{max}} P(\lambda|p, q, \lambda_{max}) P(k|\Theta) \cdot d\lambda \quad (D.9)$$

with $\Theta = \{\lambda, \dots\}$ denoting a set of parameters.

For example $P(s|p, q, \lambda_{max})$ is a Beta-Poisson mixture distribution whenever the field shape and the jitter distribution are described by a power function and the spike counts are drawn from a Poisson distribution $P(s|\Theta = \lambda) = \lambda^s \exp(-\lambda)/s!$. Furthermore $P(s|p, q = 1/2, \lambda_{max})$ is a first order approximation of count distributions in situations in which the tuning curve is Gaussian. Following our derivation here k are spike counts in the center of a firing field but generally as any other count distribution that distribution can be used to model counts at any position. For $P(s|\Theta = \lambda)$ being Poisson we can reformulate the probability mass function using the confluent hypergeometric function ${}_1F_1$ which is implemented in common programming languages:

$$P(s|p, q, \lambda_{max}) = \lambda_{max}^s \exp(-\lambda_{max}) \frac{\Gamma(p+q)\Gamma(p+k)}{s!\Gamma(p+q+k)\Gamma(p)} \cdot {}_1F_1(q, p+q+s, \lambda_{max}) \quad (D.10)$$

$$\mu = \langle k \rangle = \frac{p}{p+q} \lambda_{max} \quad (D.11)$$

$$\begin{aligned} \sigma^2 &= \langle s^2 \rangle - \langle s \rangle^2 = \\ &= \frac{pq}{(p+q+1)(p+q)^2} \lambda_{max}^2 + \lambda_{max} \frac{p}{p+q} = \\ &= \frac{q}{p(p+q+1)} \cdot \mu^2 + \mu \end{aligned} \quad (D.12)$$

This distribution approximates Poisson in the limit of $p \rightarrow \infty$ arbitrarily well. It captures the phenomenon of zero inflation for $p < 1$. The statistical law of total variance that any mixture model must increase the variance over the non-mixed model. Hence, a Poisson mixture model must have dispersion index $\sigma^2/\mu \leq 1$. Likelihood ratio tests for comparison to Poisson can be applied. To reduce the number of extra parameters, we fix $q = 1/2$ and $\langle \lambda \rangle = \langle S \rangle$ (the ML estimator of the Poisson parameter) leaving only parameters p and λ_{max} to be fitted. The resulting subset of possible Beta-Poisson mixtures has only two parameters, the same number that the Negative Binomial and Zero Inflated Poisson

distributions have. Therefore, their likelihoods can be compared directly. Alternatively, goodness of fit scores such as the Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC) or deviance can be used for model comparison scenarios with different number of parameters.

Similarly, interpreting λ itself as a probability $\rho \in [\lambda_{min} = 0, \lambda_{max} = 1]$, for a Beta mixture of Negative Binomial distributions with parameters $\Theta = \{\rho, r\}$ we get

$$P(s|p, q, r) = \frac{\Gamma(r+s)}{s! \Gamma(r)} \frac{\beta(p+r, q+s)}{\beta(p, q)} \quad (D.13)$$

$$\mu = \langle s \rangle = \frac{rq}{p-1} \text{ for } p > 1 \text{ otherwise } \infty \quad (D.14)$$

$$\sigma^2 = \langle s^2 \rangle - \langle s \rangle^2 = \mu \cdot \frac{(p+r-1)(p+q-1)}{(p-2)(p-1)} \text{ for } p > 2 \text{ otherwise } \infty \quad (D.15)$$

which is also known as the *generalized Waring distribution* [164] where $r > 0$ is the integer parameter of the Negative Binomial distribution and can be interpreted as a required number of failure crossings until a binomial experiment is stopped. The Negative Binomial distribution is itself a Gamma mixture of Poisson distributions. Again, it is clear that the Beta mixture of Negative Binomial distributions exhibits overdispersion. For large p and q , this distribution approximates the Negative Binomial distribution arbitrarily well. When r is large in addition, a Poisson distribution can be approximated arbitrarily well. Therefore, Likelihood ratio tests offer themselves for comparisons to Poisson and Negative Binomial.

Allowing the field shape and the jitter distributions to have different widths (w_f and w_j) we obtain a generalized version of Beta-mixtures:

$$\lambda = g(x) = \lambda_{max} - [w_f(x - x_0)]^{\frac{1}{q}}, \quad q > 0 \quad (D.16)$$

$$P(x) \propto (\lambda_{max} - [w_j(x - x_0)]^{\frac{1}{q}})^{p-1} \quad (D.17)$$

$$P(s|p, q, \lambda_{max}, \delta x, w_f, w_j) = C_2 \int_0^{\lambda_{max}} (\lambda - \lambda_{min})^{p-1} (\lambda_{max} - \lambda)^{q-1} P(s|\Theta) d\lambda \quad (D.18)$$

with $\lambda_{min} = \lambda_{max}(1 - (\frac{w_j}{w_f})^{\frac{1}{q}})/(\frac{w_j}{w_f})^{\frac{1}{q}}$ and normalization constant C_2 . As the Beta-Poisson mixture shows to be flexible enough and powerful to model grid cell spiking data we do not require further description of its generalized versions.

Appendix E

Miscellaneous

E.1 Pooled return map of grid cells

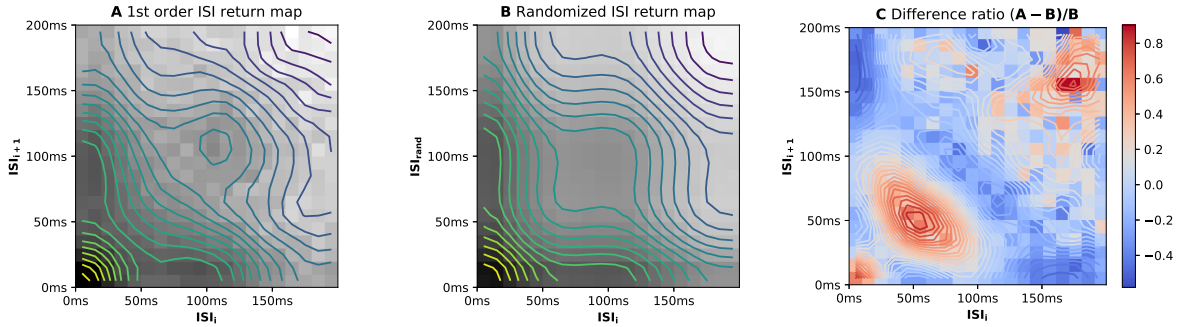


Figure E.1: First order inter spike interval (ISI) return map (**A**) and randomized ISI return map (**B**) of the 25 grid cells from the Stensola data [36] as well their difference divided by the randomized map (**C**). The (ISI_i, ISI_{i+1}) -pairs were collected from all the cells. Then a two-dimensional histogram was computed. The grey values describes a gradient from high density (black) to low density (white). The 20 contour lines are linearly spaced in the logarithmic scale of a smoothed version of the histogram (gaussian kernel smoothing, bandwidth 10ms). The shoulders at around 100ms along the axes as well as on the diagonal indicate theta modulation being present in the spike trains. The randomized return map in **B** was computed from random samples from the full ISI distribution before computing the histogram as in **A**. Using this sampling allows to compare the return map to a renewal process with the same ISI distribution: The difference ratio map (positive values red, negative values blue) in **C** indicates larger fraction of small ISI pairs ($< 20ms$, high gamma range) than expected from a renewal process. Additionally a predominance of pairs of ISIs in an intermediate range from 10Hz to 40Hz is suggested and less switches from 20Hz to high gamma or vice versa than expected from random sampling are observed.

E.2 mecPhysio Viewer - A graphical user interface for hands-on cell classification

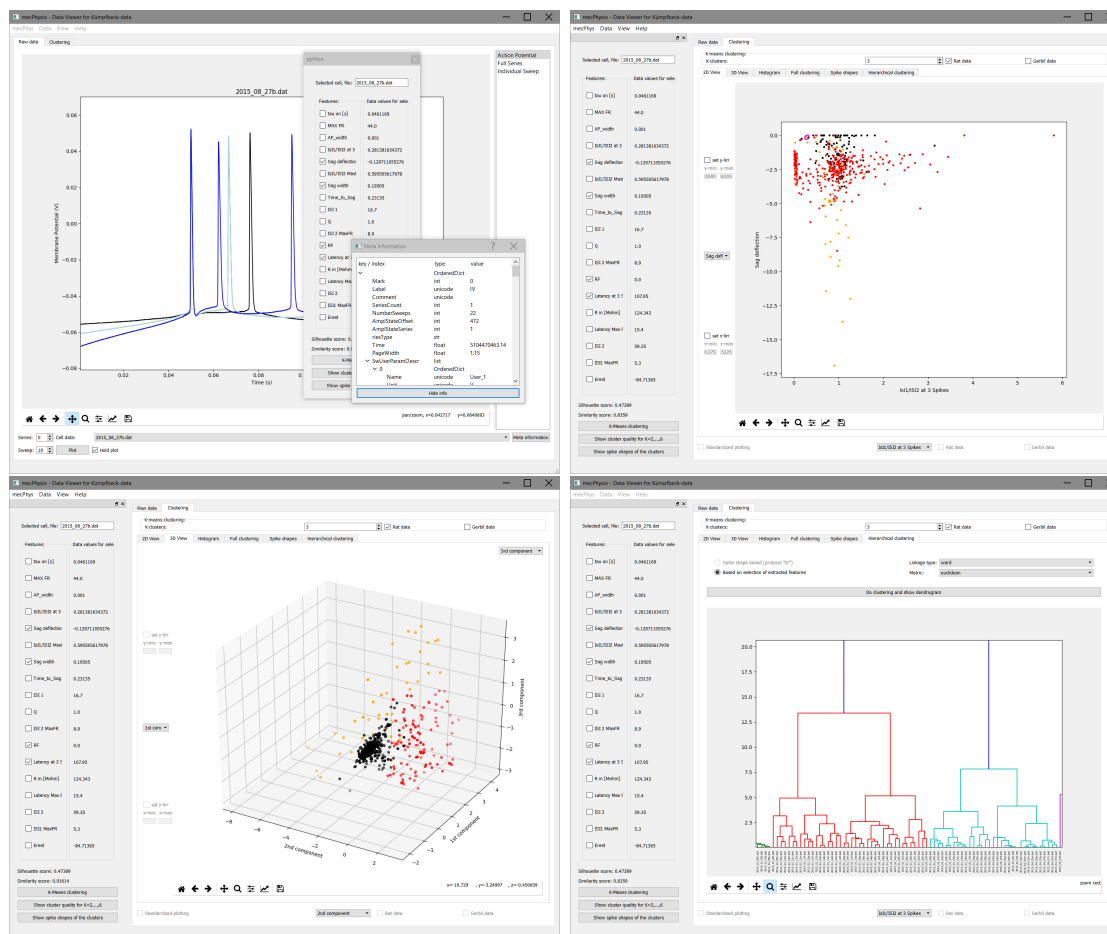


Figure E.2: During the work on a collaboration with Dr. Franziska Kämpfbeck it became evident that tools are missing for analyzing in-vivo patch-clamp recordings without profound programming and machine-learning skills. Therefore, I developed a graphical user interface (GUI) based on *PyQT5* (<https://riverbankcomputing.com/software/pyqt/intro>): *mecPhysio Viewer*. This software allows to visualize and analyze data from patch-clamp recordings of neurons. To this end, the user can import the widely used *heka* files and additional metadata in Excel format (optional). Making use of the Python libraries *scipy*, *pandas* and *sklearn* the *mecPhysio Viewer* enables extensive feature selection, clustering and plotting of voltage traces and electrophysiological features of the neurons in various ways. A complete description of the functionality of the software would go beyond the scope of this work. The tool is still in a work-in-progress state and is planned to be published on an open-source platform in the future.

E.3 Bin size effect

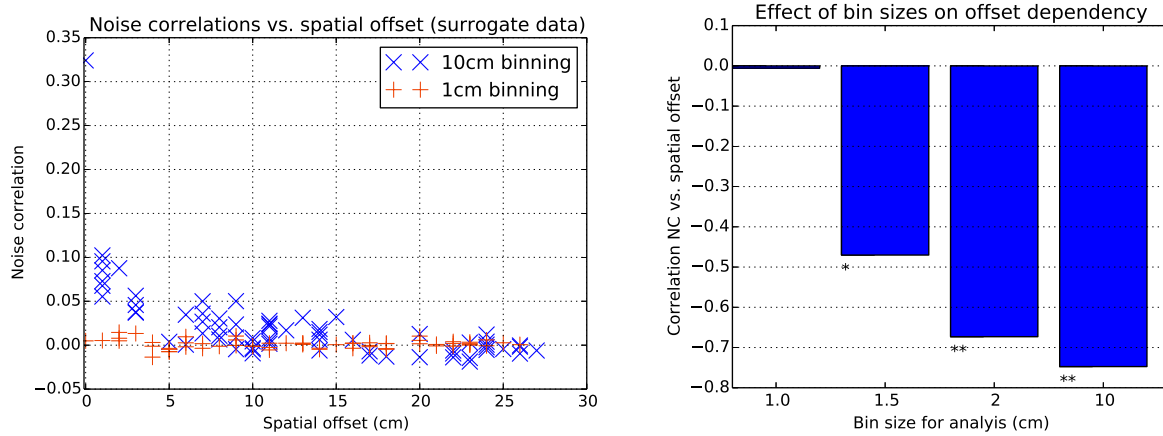


Figure E.3: Bin size affects measures of noise correlations. Left: Data shown for simulated grid cells based on the trajectory of session 10073-17010302 from the Sargolini et al. data [42] and the firing rate map of cell T3C2 (1cm binning). Spikes were generated with inhomogeneous Poisson processes. The firing rate map was shifted by 60 random displacement vectors of the length shown on the x-axis (spatial offset). Blue x: Noise-correlation estimated with 10cm binning. Red +: Same but estimated with 1cm binning. Right: Resulting dependencies (measured as Pearson correlation) of the noise correlations on spatial offsets for different bin size. For the creation of the surrogate cell pairs (60 per binning) a resolution of 1cm was used throughout (*: $p < 1e-3$, **: $p < 1e-10$). In the normal situation of analyzing neural data the internal resolution is not known. Therefore the dependency of noise correlation to tuning offset can easily be overestimated.

E.4 Grid field detection in 2D: An algorithm inspired by physics

Approach extends clustering algorithms that do not rely on a fixed cluster number:


- Exact number of clusters does not need to be set
- Permits uniform space coverage despite spatial undersampling
- No need for manual parameter tuning for individual grid cells

Self-organized grid clustering – the algorithm:


1. Initiation:
 - 1.1 Weight spikes by kernel smoothed spike train
 - 1.2 Space initial cluster nodes uniformly in arena with high density
2. Update node positions:
 - 2.1 Attraction by weighted spikes
 - 2.2 Repulsion by other nodes
3. Merge nearby nodes (e.g., closer than 10 cm)
4. Check for convergence. Possible criterion: Total translation of node points. If not converged, go back to step 2.

Advantages: Tunable method with minimal set of parameters: kernel width and ratio of attraction vs. repulsion. Fast convergence, additional assumptions can be easily added, no prior knowledge about the exact number of nodes required. Probabilistic description through Boltzmann statistics and Bayesian framework.

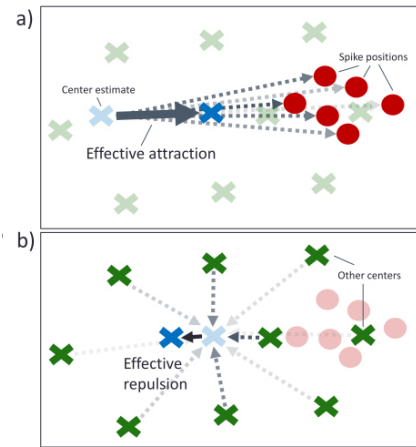
Potential pitfalls: Algorithm might push nodes to locations where there are no data.



Spikes and trajectory: a) Original data b) after initiation steps 1.1 and 1.2




Total translation over iterations



Update of node positions:

- a) Attraction by spikes and
- b) Repulsion by other nodes



Left: Final clusters: Size of red dots (estimated field centers) shows their posterior probability. Blue crosses: Initial positions. Small dots: Spikes, colored by assigned clusters. Grey: Firing-rate estimate from Gaussian Process (7). Grey arrows: Trajectories of cluster centers during the convergence process. **Right: Heat map** (Gaussian Process)

Figure E.4: Probabilistic and self-organized firing field detection based on dwell-time weighted spike positions. Taken from my poster presented at the Bernstein Conference 2015 in Göttingen [165].

Bibliography

- [1] Moser EI, Moser MB, McNaughton BL (2017) Spatial representation in the hippocampal formation: a history. *Nature Publishing Group* 20:1448–1464.
- [2] Tosches MA, et al. (2018) Evolution of pallium, hippocampus, and cortical cell types revealed by single-cell transcriptomics in reptiles. *Science* 360(6391).
- [3] Cheung TH, Cardinal RN (2005) Hippocampal lesions facilitate instrumental learning with delayed reinforcement but induce impulsive choice in rats. *BMC Neuroscience* 6(1):36.
- [4] O’Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* 34(1):171–175.
- [5] Pérez-Escobar JA, Kornienko O, Latuske P, Kohler L, Allen K (2016) Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. *eLife* 5(JULY):1–21.
- [6] Schmidt-Hieber C, Häusser M (2013) Cellular mechanisms of spatial navigation in the medial entorhinal cortex. *Nature neuroscience* 16(3):325–31.
- [7] Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA (1998) A Statistical Paradigm for Neural Spike Train Decoding Applied to Position Prediction from Ensemble Firing Patterns of Rat Hippocampal Place Cells. *J. Neurosci.* 18(18):7411–7425.
- [8] Acharya L, Aghajian ZM, Vuong C, Moore JJ, Mehta MR (2016) Causal Influence of Visual Cues on Hippocampal Directional Selectivity. *Cell* 164(1-2).
- [9] O’Keefe J (1999) Do hippocampal pyramidal cells signal non-spatial as well as spatial information? *Hippocampus* 9(4):352–364.
- [10] O’keefe J, Nadel L (1979) Précis of O’Keefe and Nadel’s The hippocampus as a cognitive map. *Behavioral and Brain Sciences* 2(4):487–494.
- [11] Mou X, Ji D (2016) Social observation enhances cross- environment activation of hippocampal place cell patterns. *eLife* 5(e18022):1–26.

- [12] Omer DB, Maimon SR, Las L, Ulanovsky N (2018) Social place-cells in the bat hippocampus. *Science* 359.
- [13] Danjo T, Toyozumi T, Fujisawa S (2018) Spatial representations of self and other in the hippocampus. *Science* 359(6372).
- [14] Tolman EC (1948) Cognitive maps in rats and men. *Psychological Review* 55(4):189–208.
- [15] Taube JS, Muller RU, Ranck JB (1990) Head-direction cells recorded from the post-subiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 10(2):420–35.
- [16] Fyhn M, Molden S, Witter MP, Moser EI, Moser MB (2004) Spatial representation in the entorhinal cortex. *Science* 305(5688):1258–64.
- [17] Hafting T, Fyhn M, Molden S, Moser MB, Moser EI (2005) Microstructure of a spatial map in the entorhinal cortex. *Nature* 436(7052):801–6.
- [18] Yartsev MM, Witter MP, Ulanovsky N (2011) Grid cells without theta oscillations in the entorhinal cortex of bats. *Nature* 479(7371):103–107.
- [19] Killian NJ, Jutras MJ, Buffalo EA (2012) A map of visual space in the primate entorhinal cortex. *Nature* pp. 2–15.
- [20] Jacobs J, et al. (2013) Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature neuroscience* 16(9):1188–90.
- [21] Doeller CF, Barry C, Burgess N (2010) Evidence for grid cells in a human memory network. *Nature* 463(7281):657–61.
- [22] Boccara CN, et al. (2010) Grid cells in pre- and parasubiculum. *Nature Neuroscience* 13(8):987–994.
- [23] Knierim JJ, Neunuebel JP, Deshmukh SS (2014) Functional correlates of the lateral and medial entorhinal cortex: objects, path integration and local–global reference frames. *Phil Trans R Soc. B* 369:20130369.
- [24] Tang Q, et al. (2014) Pyramidal and Stellate Cell Specificity of Grid and Border Representations in Layer 2 of Medial Entorhinal Cortex. *Neuron* pp. 1–7.
- [25] Sun C, et al. (2015) Distinct speed dependence of entorhinal island and ocean cells, including respective grid cells. *Proceedings of the National Academy of Sciences* 112(30).
- [26] Domnisoru C, Kinkhabwala AA, Tank DW (2013) Membrane potential dynamics of grid cells. *Nature* 495(7440):199–204.

- [27] Fenton AA (2015) Coordinating with the "inner GPS". *Hippocampus* 7(March):1–7.
- [28] Bush D, Barry C, Manson D, Burgess N (2015) Using Grid Cells for Navigation. *Neuron* 87(3):507–520.
- [29] McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB (2006) Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience* 7(8):663–678.
- [30] Rennó-Costa C, Lisman JE, Verschure PF (2010) The Mechanism of Rate Remapping in the Dentate Gyrus. *Neuron* 68(6):1051–1058.
- [31] Gu Y, et al. (2018) A Map-like Micro-Organization of Grid Cells in the Medial Entorhinal Cortex. *Cell* 175(3):736–750.
- [32] Spalla D, Dubreuil A, Rosay S, Monasson R, Treves A (2019) Can Grid Cell Ensembles Represent Multiple Spaces? *Neural Computation* pp. 1–24.
- [33] Zhu HX, Windle A (2001) The geometrical properties of irregular two-dimensional Voronoi tessellations. *Philosophical Magazine A* 81(12):2765 – 2783.
- [34] Møller J (1994) *Lectures on Random Voronoi Tessellations*, Lecture Notes in Statistics. (Springer New York, New York, NY) Vol. 87.
- [35] Chen G, King JA, Lu Y, Cacucci F, Burgess N (2018) Spatial cell firing during virtual navigation of open arenas by head-restrained mice. *eLife* 7.
- [36] Stensola H, et al. (2012) The entorhinal grid map is discretized. *Nature* 492(7427):72–78.
- [37] Brun VH, et al. (2008) Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus* 18(12):1200–1212.
- [38] Rowland DC, Moser MB (2014) From cortical modules to memories. *Current Opinion in Neurobiology* 24:22–27.
- [39] Mathis A, Herz AVM, Stemmler MB (2013) Multi-Scale Codes in the Nervous System: The Problem of Noise Correlations and the Ambiguity of Periodic Scales. *Physical Review E* 88(2):022713.
- [40] Yoon K, et al. (2013) Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature Neuroscience* 16(8):1077–1084.
- [41] Pröll M, Häusler S, Herz AVM (2018) Grid-cell activity on linear tracks indicates purely translational remapping of 2D firing patterns at movement turning points. *The Journal of neuroscience* 38(31):7004–7011.

- [42] Sargolini F, et al. (2006) Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* 312(5774):758–62.
- [43] Barry C, Burgess N (2017) To be a Grid Cell: Shuffling procedures for determining "Gridness". *bioRxiv* p. 230250.
- [44] Weber SN, Sprekeler H (2019) A local measure of symmetry and orientation for individual spikes of grid cells. *PLoS Computational Biology* 15(2):e1006804.
- [45] Souza BC, Pavão R, Belchior H, Tort AB (2018) On Information Metrics for Spatial Coding. *Neuroscience* 375(February):62–73.
- [46] Burgess N (2008) Grid cells and theta as oscillatory interference: Theory and predictions. *Hippocampus* 18(12):1157–1174.
- [47] Burgess N, Barry C, O'Keefe J (2007) An oscillatory interference model of grid cell firing. *Hippocampus* 17(9):801–812.
- [48] Hasselmo ME (2005) The role of hippocampal regions CA3 and CA1 in matching entorhinal input with retrieval of associations between objects and context: theoretical comment on Lee et al. (2005). *Behavioral neuroscience* 119(1):342–345.
- [49] Fuhs MC, Touretzky DS (2006) A spin glass model of path integration in rat medial entorhinal cortex. *The Journal of neuroscience* 26(16):4266–76.
- [50] Burak Y (2006) Do We Understand the Emergent Dynamics of Grid Cell Activity? *Journal of Neuroscience* 26(37):9352–9354.
- [51] Koenig J, Linder AN, Leutgeb JK, Leutgeb S (2011) The spatial periodicity of grid cells is not sustained during reduced theta oscillations. *Science (New York, N.Y.)* 332(6029):592–5.
- [52] Brandon MP, et al. (2011) Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. *Science (New York, N.Y.)* 332(6029):595–9.
- [53] Wills TJ, Cacucci F, Burgess N, O'Keefe J (2010) Development of the Hippocampal Cognitive Map in Prewaning Rats. *Science* 328(5985):1573–1576.
- [54] Hales J, et al. (2014) Medial Entorhinal Cortex Lesions Only Partially Disrupt Hippocampal Place Cells and Hippocampus-Dependent Place Memory. *Cell Reports* 9(3):893–901.
- [55] Schlesiger MI, et al. (2015) The medial entorhinal cortex is necessary for temporal organization of hippocampal neuronal activity. *Nature Neuroscience* 18(8):1123–1132.
- [56] Kropff E, Treves A (2008) The emergence of grid cells: Intelligent design or just adaptation? *Hippocampus* 18(12):1256–1269.

- [57] Dordek Y, Soudry D, Meir R, Derdikman D (2016) Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *eLife* 5.
- [58] Stepanyuk A (2015) Self-organization of grid fields under supervision of place cells in a neuron model with associative plasticity. *Biologically Inspired Cognitive Architectures* 13(July 2015):48–62.
- [59] Monsalve-Mercado MM, Leibold C (2017) Hippocampal Spike-Timing Correlations Lead to Hexagonal Grid Fields. *Physical Review Letters* 038101(1):1–4.
- [60] Lörincz A, Sárkány A (2017) Semi-supervised learning of Cartesian Factors: A top-down model of the entorhinal hippocampal complex. *Frontiers in Psychology* 8(FEB).
- [61] Banino A, et al. (2018) Vector-based navigation using grid-like representations in artificial agents. *Nature* 557:429–433.
- [62] Stemmler M, Mathis A, Herz AVM (2015) Decoding the Population Activity of Grid Cells for Spatial Localization and Goal-Directed Navigation. *bioRxiv* p. 021204.
- [63] Huhn Z, Somogyvári Z, Kiss T, Érdi P (2009) Distance coding strategies based on the entorhinal grid cell system. *Neural Networks* 22(5-6):536–543.
- [64] Kerdels J, Peters G (2015) A New View on Grid Cells Beyond the Cognitive Map Hypothesis in *Artificial General Intelligence. AGI 2015. Lecture Notes in Computer Science*, eds. J. B, B. G, A. P. (Springer, Cham), 9205 edition, pp. 283–292.
- [65] Behrens TE, et al. (2018) What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*.
- [66] Bicanski A, Burgess N (2019) A Computational Model of Visual Recognition Memory via Grid Cells. *Current Biology* 29(6).
- [67] Kriegeskorte N, Storrs KR (2016) Grid Cells for Conceptual Spaces? *Neuron* 92(2):280–284.
- [68] Constantinescu AO, O'Reilly JX, Behrens TEJ (2016) Organizing conceptual knowledge in humans with a gridlike code. *Science* 352(6292).
- [69] Aronov D, Nevers R, Tank DW (2017) Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature* 543(7647):719–722.
- [70] Garvert MM, Dolan RJ, Behrens TE (2017) A map of abstract relational knowledge in the human hippocampal-entorhinal cortex. *eLife* 6:1–20.
- [71] Bao X, et al. (2019) Grid-like Neural Representations Support Olfactory Navigation of a Two-Dimensional Odor Space. *Neuron*.

- [72] Kropff E, Carmichael JE, Moser MB, Moser EI (2015) Speed cells in the medial entorhinal cortex. *Nature* 523(7561):419–424.
- [73] Diehl GW, Hon OJ, Leutgeb S, Leutgeb JK (2017) Grid and Nongrid Cells in Medial Entorhinal Cortex Represent Spatial Location and Environmental Features with Complementary Coding Schemes. *Neuron* 94(1):83–92.
- [74] Boccara CN, Nardin M, Stella F, O’Neill J, Csicsvari J (2019) The entorhinal cognitive map is attracted to goals. *Science* 363(6434):1443–1447.
- [75] Butler WN, Hardcastle K, Giocomo LM (2019) Remembered reward locations restructure entorhinal spatial maps. *Science* 363(6434):1447–1452.
- [76] Hardcastle K, Maheswaranathan N, Ganguli S, Giocomo LM (2017) A Multiplexed, Heterogeneous, and Adaptive Code for Navigation in Medial Entorhinal Cortex. *Neuron* 94(2):375–387.
- [77] Fyhn M, Hafting T, Treves A, Moser MB, Moser EI (2007) Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446(7132):190–4.
- [78] Barry C, Hayman R, Burgess N, Jeffery KJ (2007) Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience* 10(6):682–684.
- [79] Wagenaar D, Pine J, Potter S (2006) An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neuroscience* 7(1):11.
- [80] Izhikevich EM (2000) Neural excitability, spiking and bursting. *International Journal of Bifurcation and Chaos* 10(06):1171–1266.
- [81] Izhikevich EM, Desai NS, Walcott EC, Hoppensteadt FC (2003) Bursts as a unit of neural information: selective communication via resonance. *Trends in neurosciences* 26(3):161–7.
- [82] Lisman JE (1997) Bursts as a unit of neural information: making unreliable synapses reliable. *Trends in Neurosciences* 20(1):38–43.
- [83] Cooper DC (2002) The significance of action potential bursting in the brain reward circuit. *Neurochemistry International* 41(5):333–340.
- [84] Hinman JR, Brandon MP, Climer JR, Chapman GW, Hasselmo ME (2016) Multiple Running Speed Signals in Medial Entorhinal Cortex. *Neuron* 91(3).
- [85] Reifenstein ET, Kempter R, Schreiber S, Stemmler MB, Herz AVM (2012) Grid cells in rat entorhinal cortex encode physical space with independent firing fields and phase precession at the single-trial level. *Proceedings of the National Academy of Sciences of the United States of America* 109(16):6301–6.

- [86] Reifenstein E, Stemmler M, Herz AVM, Kempter R, Schreiber S (2014) Movement dependence and layer specificity of entorhinal phase precession in two-dimensional environments. *PloS one* 9(6):e100638.
- [87] Latuske P, Toader O, Allen K (2015) Interspike Intervals Reveal Functionally Distinct Cell Populations in the Medial Entorhinal Cortex. *The Journal of Neuroscience* 35(31):10963–10976.
- [88] Ólafsdóttir HF, Carpenter F, Barry C (2016) Coordinated grid and place cell replay during rest. *Nature Neuroscience* 19(April):792–794.
- [89] Neill JO, Boccara CN, Stella F, Schoenenberger P, Csicsvari J (2017) Superficial Layers of the Medial Entorhinal Cortex Replay Independent of the Hippocampus. *Science* 188(January):184–188.
- [90] Ismakov R, Barak O, Jeffery K, Derdikman D (2017) Grid Cells Encode Local Positional Information. *Current Biology* 27(15).
- [91] Dunn B, Wennberg D, Huang Z, Roudi Y (2017) Grid cells show field-to-field variability and this explains the aperiodic response of inhibitory interneurons. *arXiv preprint* pp. 1–33.
- [92] Stemmler M, Herz AVM (2017) Spatial Cognition: Grid Cells Harbour Three Complementary Positional Codes. *Current Biology* 27(15):R755–R758.
- [93] Alonso A, Klink R (1993) Differential electroresponsiveness of stellate and pyramidal-like cells of medial entorhinal cortex layer II. *Journal of neurophysiology* 70(1):128–43.
- [94] Alessi C, Raspanti A, Magistretti J (2016) Two distinct types of depolarizing afterpotentials are differentially expressed in stellate and pyramidal-like neurons of entorhinal-cortex layer II. *Hippocampus* 26(3):380–404.
- [95] Canto CB, Witter MP (2012) Cellular properties of principal neurons in the rat entorhinal cortex. II. The medial entorhinal cortex. *Hippocampus* 22(6):1277–1299.
- [96] Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual neuroscience* 13(1):87–100.
- [97] Renart A, Machens CK (2014) Variability in neural activity and behavior. *Current Opinion in Neurobiology* 25:211–220.
- [98] Maynard Smith J (1978) OPTIMIZATION THEORY IN EVOLUTION. *Ann. Rev. Ecol. Syst* 9:31–56.
- [99] Sarkar S (2005) Maynard Smith, optimization, and evolution. *Biology and Philosophy* 20(5):951–966.

- [100] Stein RB, Gossen ER, Jones KE (2005) Neuronal variability: noise or part of the signal? *Nature reviews. Neuroscience* 6(May):389–397.
- [101] Ölveczky BP, Otchy TM, Goldberg JH, Aronov D, Fee MS (2011) Changes in the neural control of a complex motor sequence during learning. *Journal of Neurophysiology* 106(1):386–397.
- [102] Edelman GM (1993) Neural Darwinism: Selection and reentrant signaling in higher brain function. *Neuron* 10(2):115–125.
- [103] Dinstein I, Heeger DJ, Behrmann M (2015) Neural variability: Friend or foe? *Trends in Cognitive Sciences* 19(6):322–328.
- [104] Masquelier T (2013) Neural variability, or lack thereof. *Frontiers in Computational Neuroscience* 7:7.
- [105] Chervyakov AV, Sinitsyn DO, Piradov MA (2016) Variability of Neuronal Responses: Types and Functional Significance in Neuroplasticity and Neural Darwinism. *Frontiers in human neuroscience* 10:603.
- [106] Benda J, Herz AVM (2003) A Universal Model for Spike-Frequency Adaptation. *Neural Computation* 15(11):2523–2564.
- [107] McFadden JA (1962) On the Lengths of Intervals in a Stationary Point Process. *Journal of the Royal Statistical Society: Series B (Methodological)* 24(2):364–382.
- [108] Nawrot MP (2011) Stochastic point processes. *Factorization calculus and geometric probability* pp. 161–199.
- [109] Reich DS, Victor JD, Knight BW, Ozaki T, Kaplan E (1997) Response Variability and Timing Precision of Neuronal Spike Trains In Vivo. *Journal of Neurophysiology* 77(5):2836–2841.
- [110] Ponce-Alvarez A, Kilavik BE, Riehle A (2010) Comparison of local measures of spike time irregularity and relating variability to firing rate in motor cortical neurons. *Journal of Computational Neuroscience* 29(1-2):351–365.
- [111] Fenton AA, et al. (2010) Attention-like modulation of hippocampus place cell discharge. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 30(13):4613–25.
- [112] Rad KR, Paninski L (2010) Efficient, adaptive estimation of two-dimensional firing rate surfaces via Gaussian process methods. *Network (Bristol, England)* 21(3-4):142–68.
- [113] Shimokawa T, Shinomoto S (2009) Estimating Instantaneous Irregularity of Neuronal Firing. *Neural Computation* 21(7):1931–1951.

- [114] Fenton AA, Muller RU (1998) Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences of the United States of America* 95(6):3182–7.
- [115] Amarasingham A, Geman S, Harrison MT (2015) Ambiguity and nonidentifiability in the statistical analysis of neural codes. *Proceedings of the National Academy of Sciences* 112(20):6455–6460.
- [116] Lee I, Knierim JJ (2007) The relationship between the field-shifting phenomenon and representational coherence of place cells in CA1 and CA3 in a cue-altered environment. *Learning & memory (Cold Spring Harbor, N.Y.)* 14(11):807–815.
- [117] Bahar AS, Shirvalkar PR, Shapiro ML, Kastor GJ (2011) Memory-Guided Learning: CA1 and CA3 Neuronal Ensembles Differentially Encode the Commonalities and Differences between Situations. *The Journal of neuroscience* 31(34):12270–12281.
- [118] Roth ED, Yu X, Rao G, Knierim JJ (2012) Functional differences in the backward shifts of ca1 and ca3 place fields in novel and familiar environments. *PLoS ONE* 7(4):1–10.
- [119] Hardcastle K, Ganguli S, Giocomo L (2015) Environmental Boundaries as an Error Correction Mechanism for Grid Cells. *Neuron* 86(3):827–839.
- [120] Marozzi E, Ginzberg LL, Alenda a, Jeffery KJ (2015) Purely Translational Realignment in Grid Cell Firing Patterns Following Nonmetric Context Change. *Cerebral Cortex* pp. 1–9.
- [121] Babichev A, Dabaghian YA (2018) Topological Schemas of Memory Spaces. *Frontiers in Computational Neuroscience* 12:27.
- [122] Dabaghian Y, Mémoli F, Frank L, Carlsson G (2012) A Topological Paradigm for Hippocampal Spatial Map Formation Using Persistent Homology. *PLoS Computational Biology* 8(8).
- [123] Curto C (2015) What can topology tell us about the neural code? *arXiv preprint* pp. 1–16.
- [124] Babichev A, Ji D, Mémoli F, Dabaghian YA (2016) A Topological Model of the Hippocampal Cell Assembly Network. *Frontiers in Computational Neuroscience* 10(June):1–12.
- [125] Arai M, Brandt V, Dabaghian Y (2014) The Effects of Theta Precession on Spatial Learning and Simplicial Complex Dynamics in a Topological Model of the Hippocampal Spatial Map. *PLoS Computational Biology* 10(6).

- [126] Giusti C, Pastalkova E, Curto C, Itskov V (2015) Clique topology reveals intrinsic geometric structure in neural correlations arXiv : 1502 . 06172v1 [q-bio . NC] 22 Feb 2015. *ArXiv* pp. 1–29.
- [127] Giusti C, Ghrist R, Bassett DS (2016) Two’s company, three (or more) is a simplex: Algebraic-topological tools for understanding higher-order structure in neural data. *Journal of Computational Neuroscience* 41(1):1–14.
- [128] Low RJ, Lewallen S, Aronov D, Nevers R, Tank DW (2018) Probing variability in a cognitive map using manifold inference from neural dynamics. *bioRxiv* p. 418939.
- [129] Chaudhuri R, Gercek B, Pandey B, Peyrache A, Fiete I (2019) The population dynamics of a canonical cognitive circuit. *bioRxiv* p. 516021.
- [130] Colgin LL (2016) Rhythms of the hippocampal network. *Nature Reviews Neuroscience* 17(4).
- [131] Simonnet J, Brecht M (2019) Burst Firing and Spatial Coding in Subicular Principal Cells. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 39(19):3651–3662.
- [132] Brandon MP, Linder AN, Leutgeb JK, Leutgeb S (2011) Reduction of Theta Rhythm Dissociates Grid Cell Spatial Periodicity from Directional Tuning. *Science* 332(6029):595–599.
- [133] Colgin LL, et al. (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* 462(7271):353–7.
- [134] Lisman J (2005) The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus* 15(7):913–922.
- [135] Igarashi KM, Lu L, Colgin LL, Moser MB, Moser EI (2014) Coordination of entorhinal-hippocampal ensemble activity during associative learning. *Nature* 510(7503):143–147.
- [136] Nakazono T, et al. (2017) Impaired In Vivo Gamma Oscillations in the Medial Entorhinal Cortex of Knock-in Alzheimer Model. *Frontiers in Systems Neuroscience* 11:48.
- [137] Shay CF, Ferrante M, Chapman Iv GW, Hasselmo ME (2016) Rebound spiking in layer II medial entorhinal cortex stellate cells: Possible mechanism of grid cell function. *Neurobiology of Learning and Memory* 129:83–98.
- [138] Chrobak JJ, Buzsáki G (1998) Gamma oscillations in the entorhinal cortex of the freely behaving rat. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 18(1):388–98.

- [139] Wong RK, Prince DA (1981) Afterpotential generation in hippocampal pyramidal cells. *Journal of neurophysiology* 45(1):86–97.
- [140] Tsubokawa H, Ross WN (1997) Muscarinic modulation of spike backpropagation in the apical dendrites of hippocampal CA1 pyramidal neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 17(15):5782–91.
- [141] Buzsáki G, Kandel A (1998) Somadendritic Backpropagation of Action Potentials in Cortical Pyramidal Cells of the Awake Rat. *Journal of Neurophysiology* 79(3):1587–1591.
- [142] Couey JJ, et al. (2013) Recurrent inhibitory circuitry as a mechanism for grid formation. *Nature Neuroscience* 16(3):318–324.
- [143] Fuchs EC, et al. (2016) Local and Distant Input Controlling Excitation in Layer II of the Medial Entorhinal Cortex. *Neuron* 89(1):194–208.
- [144] Winterer J, et al. (2017) Excitatory Microcircuits within Superficial Layers of the Medial Entorhinal Cortex. *Cell Reports* 19(6):1110–1116.
- [145] Weber AI, Pillow JW (2017) Capturing the Dynamical Repertoire of Single Neurons with Generalized Linear Models. *arXiv preprint*.
- [146] Consul P, Famoye F (1992) Generalized poisson regression model. *Communications in Statistics - Theory and Methods* 21(1):89–109.
- [147] Gao Y, Buesing L, Shenoy KV, Cunningham JP (2015) High-dimensional neural spike train analysis with generalized count linear dynamical systems. *Advances in Neural Information Processing Systems* pp. 1–9.
- [148] Fyhn M, Molden S, Witter MP, Moser EI, Moser MB (2004) Spatial representation in the entorhinal cortex. *Science (New York, N.Y.)* 305(5688):1258–64.
- [149] Nagele J, Stemmler M, Herz AV (2016) Structure of the noise in the entorhinal cortex. *Bernstein Conference 2016. Program and Abstracts*.
- [150] Pérez-Escobar J, Kornienko O, Latuske P, Kohler L, Allen K (2016) Data from: Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex, <https://datadryad.org/resource/doi:10.5061/dryad.c261c>.
- [151] Graboski JJ, Resnik E, Sirota A (2014) Behavioral state dependence of hippocampal place cell expression in the exploring rat. *Bernstein Conference 2014. Program and Abstracts*.
- [152] Grosjean G (2013) *Geschichte der Kartographie*.

- [153] Gardner RJ, Lu L, Wernle T, Moser MB, Moser EI (2019) Correlation structure of grid cells is preserved during sleep. *Nature Neuroscience* 22(4):598–608.
- [154] Chen G, Manson D, Cacucci F, Wills TJ (2016) Absence of Visual Input Results in the Disruption of Grid Cell Firing in the Mouse. *Current biology* 26(17):1–8.
- [155] Singer A (2008) A Remark on Global Positioning from Local Distances. *Proceedings of the National Academy of Sciences* 105(28):9507–9511.
- [156] Candes EJ, Plan Y (2010) Matrix completion with noise. *Proceedings of the IEEE* 98(6):925–936.
- [157] Dokmanic I, Parhizkar R, Ranieri J, Vetterli M (2015) Euclidean Distance Matrices: Essential theory, algorithms, and applications. *IEEE Signal Processing Magazine* 32(6):12–30.
- [158] Du KL (2010) Clustering: A neural network approach. *Neural Networks* 23(1):89–107.
- [159] Wang P (2001) Markov zero-inflated Poisson regression models for a time series of counts with excess zeros. *Journal of Applied Statistics* 28(5):623–632.
- [160] DeSantis SM, Bandyopadhyay D (2011) Hidden Markov models for zero-inflated Poisson counts with an application to substance use. *Statistics in Medicine* 30(14):1678–1694.
- [161] Olteanu M, Ridgway J (2012) Hidden Markov models for time series of counts with excess zeros. *Proceedings of ESANN 2012*.
- [162] Giles D (2010) Notes on the Zero-Inflated Poisson Regression Model. *Department of Economics, University of Victoria*.
- [163] Häggglund M, Mørreaunet M, Moser MB, Moser EI (2019) Grid-Cell Distortion along Geometric Borders. *Current Biology*.
- [164] Johnson NL, Kemp AW, Kotz S (2005) *Univariate Discrete Distributions*.
- [165] Nagele J, Özmen HB, Csordas D, Stemmler M, Herz AVM (2015) A closer look at grid cells : Discharge statistics and firing-field shapes. *Bernstein Conference 2015. Program and abstracts*.

Acknowledgements

I would like to thank my advisor Prof. Dr. Andreas Herz who supported me and my research by providing continuous encouragement, subtle but decisive guidance and criticism, as well as the freedom to talk to scientists at many conferences and other occasions. The work at his group for Computational Neuroscience at the Bernstein Center in Ludwig Maximilians Universität München is characterized by an open-minded, relaxed, and oftentimes cheerful atmosphere, for which I would like to thank also my current and former colleagues Dinu Patirniche, Dr. Alireza Chenani, Dr. Mauro Monsalve, Dr. Alvaro Tejero, Dr. Alexander Mathis, Dóra Csordás, Franziska Kümpfbeck, Michaela Pröll, Caroline Fischer, Florian Eberhardt, Dr. Mehrdad Salmasi, Dr. Dustin Fetterhoff, Dr. Martin Stemmler, Dr. Stefan Häusler, Dr. Kay Thurley, and many more. In some of them I have found close friends. For his scientific advice and lively discussions during my TAC meetings, in the seminars or somewhere on the floor I would also like to thank Prof. Dr. Anton Sirota.

Furthermore, I would like to thank the Graduate School for Systemic Neuroscience in Munich and all my old friends including Joachim Fröstl, Sergej Breiter, Daniel Brzoska, Patrick Zimmermann and Mariel-Carol Ziegler, for their continuous support and willingness to leave me in my scientific thoughts when needed and pull me out when required. Many thanks to Priyanka Vinnakota for the final proof reading.

Lastly, I would like to thank my family. Their lifelong and unconditional support, the scientific and philosophical discussions with my father (the wisest man on earth), the always supporting words of my mother (the kindest person on earth), as well as the rest of my beloved family are the foundation that allowed me to pursue my research in the first place.

(...not to forget the heartbreaking smiles of my nephews!)