A distributed neural code in ensembles of dentate gyrus granule cells

Fabio Stefanini[1,2,3*], Mazen A. Kheirbek[7,8,9,10,*], Lyudmila Kushnir[4,*], Joshua H. Jennings[13], Garret D. Stuber[14,15,16], René Hen[5,6,7,+], and Stefano Fusi[1,2,3,v]

[1]Center for Theoretical Neuroscience, College of Physicians and Surgeons, Columbia University
[2]Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University
[3]Kavli Institute for Brain Sciences, Columbia University
[4]GNT - LNC, Department d'ètudes cognitives, École Normale Supérieure, INSERM, PSL Research University, 75005, Paris, France
[5]Departments of Neuroscience & Psychiatry, Columbia University, New York, United States of America
[6]Division of Integrative Neuroscience, Department of Psychiatry, New York State Psychiatric Institute, New York, NY, USA
[7]Department of Psychiatry, University of California, San Francisco, San Francisco, CA, USA
[8]Neuroscience Graduate Program, University of California, San Francisco, San Francisco, CA, USA
[9]Weill Institute for Neurosciences, University of California, San Francisco, San Francisco, CA, USA
[10]Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA, USA
[11]Center for Integrative Neuroscience, University of California, San Francisco, San Francisco, CA, USA
[12]Department of Bioengineering, Stanford University, Stanford, CA, 94305, USA
[13]Department of Psychiatry, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA
[14]Neuroscience Center, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA
[15]Department of Cell Biology and Physiology, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA
[*]these authors contributed equally to this work
[+]corresponding author, rh95@cumc.columbia.edu
[v]corresponding author, sf2237@columbia.edu

ABSTRACT

Traditionally, neurons have been viewed as specialized for single functions or a few highly related functions. However, at the higher levels of processing, neural specialization is less obvious, rather a mix of disparate, seemingly unrelated, sensory, cognitive and behavioural quantities drive neural activity. Task relevant variables can be decoded by reading out populations of neurons as the neural code is highly distributed. Here we show that this seems to be the case also in the dentate gyrus subregion of the hippocampus. Using calcium imaging we simultaneously recorded the activity of hundreds of cells from the dentate gyrus of mice while they freely explored an arena. Their instantaneous position, direction of motion and speed could be accurately decoded from the neural activity. Ranking neurons by their contribution to the decoding accuracy revealed that the response properties of individual neurons were often not predictive of their importance for encoding position. Furthermore, we could decode position from neurons that were important for encoding the direction of motion and vice versa, showing that these quantities are encoded by largely overlapping populations. Our analysis indicates that classical methods of analysis based on single cell response properties might be insufficient to characterize the neural code and that the lack of observation of easily interpretable place cells in one brain area is not necessarily the indication that position is not efficiently encoded in that area.

Introduction

The hippocampus is believed to be involved in navigation and spatial memory due to the observation of neurons encoding the position of the animal. The responses of some of these cells are easily interpretable as the cells tend to fire only when the animal is at one or more locations in an environment (place cells). However, it is becoming clear that in many brain areas, which include the hippocampus and entorhinal cortex, the neural responses are very diverse[1–4] and highly variable in time[5,6]. Place cells might respond at single or multiple locations, in an orderly (grid cells) or disorderly way and multiple passes through the same location typically elicit different responses[6]. Part of the diversity can be explained by assuming that each neuron responds non-linearly to multiple variables (mixed selectivity)[1]. Each neuron might respond to different non-linear combinations of these variables, which would explain the diversity. The variability might be because some of these variables are not being monitored in the experiment, and hence contribute to what might appear as noise. A neural code based on mixed selectivity is highly distributed because some variables can be decoded only by reading out the activity of a population of neurons. It was recently shown that the mixed selectivity component of the neuronal responses is important in complex cognitive tasks[1,3] because it is a signature of the high dimensionality of the neural representations. These recent studies naturally pose the question of how position is encoded within population activity of the hippocampus. To answer this question, we used calcium imaging to record the activity of a large population of neurons in the dentate gyrus (DG), a region of the hippocampus in which the neural responses are highly sparse and diverse[7]. We show that the position of a mouse freely exploring an environment can be decoded with high accuracy from the activity of a few tens of granule cells (GCs), despite the sparseness and diversity of the code. Position can be decoded also from cells that would not be considered spatially tuned according to the standard statistical tests based on mutual information. Using machine learning techniques, we ranked neurons by their contribution to position encoding. We found that trial averaged, single neuron properties (e.g., the place fields) are insufficient to predict a neuron's contribution to position encoding. Some single field neurons with high spatial information were unreliable and were ranked low. On the other hand, some neurons that did not meet the criteria for being classified as place cells were important for the position decoder. We further found that DG neurons encoded other variables, as the direction and speed of movement, which could also be decoded. These neurons were not distinct from the neurons that encode position, i.e., the majority of neurons encoded multiple variables. Our results suggest that the information encoded at the population level is far richer than at the single cell level.

Results

We studied the neural code in the DG of freely moving mice. We used miniaturized head-mounted microscopes to perform calcium imaging of DG granule cells (GC). To image activity patterns in DG GCs we injected a virus encoding the calcium indicator GCaMP6m under control of the CaMKII promoter into the dorsal DG (AAVdj-CaMKII-GCaMP6m-WPRE-SV40, Stanford) and implanted a gradient index (GRIN) lens for chronic imaging (Fig. 1). Four weeks after surgery, we imaged cellular activity in the DG while mice foraged for sucrose pellets in an open field arena. We then used a recently developed algorithm for reliably extracting the GCaMP signals from the raw videos, CNMF-e[9](Fig. 1d, e). This algorithm separates local background signals due to changes in fluorescence in the neuropil from the signals due to calcium concentration changes in individual cells and was particularly effective in identifying signal sources in our granule cells data without introducing spurious distortions or correlations among cells due to artifacts. We recorded a total of 1109 cells across 3 animals, among which 352 (32%) were significantly tuned to position (see Methods and Supplementary Fig. 2).

The first step of our analysis was to assess whether the position of the animal is encoded in the recorded

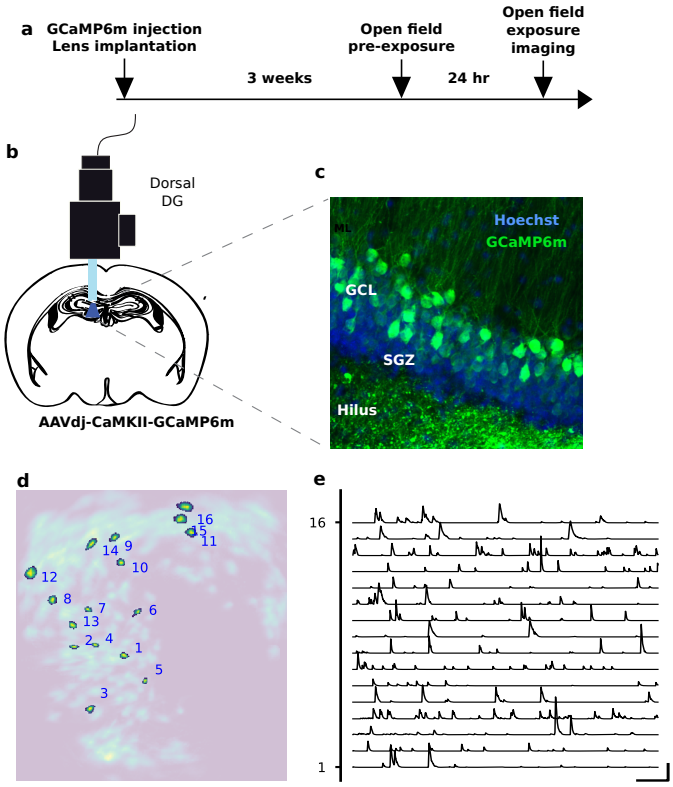

Figure 1. **a**) Experiment protocol. Mice were anesthetized with isoflurane and placed in a stereotactic apparatus. They were then injected in the dorsal DG with a virus encoding GCaMP6m. Mice were then implanted with a GRIN lens and three weeks after surgery they were checked for GCaMP expression with a miniaturized microscope (Inscopix, Palo Alto, CA) and procedures previously described[8]. A baseplate was attached to the skull at the optimal imaging plane and one week later mice were imaged during foraging in an open field task. **b-c**) Recording site. Mice were implanted with a GRIN lens and a baseplate was attached to the skull at the optimal imaging place. The imaging plane was later assessed through histology. (GCL, granule cell layer; SGZ, subgranular zone). **d-e**) Automated signal extraction using CNMF-e[9]. The algorithm identifies the spatial (left) and temporal (right) components of the signal sources, i.e., putative GCs. It uses a generative model of Calcium traces and non-negative matrix factorization to separate actual signal sources from the background due to diffused neuropil fluorescence. The extracted spatial components are displayed on the left, a few representative ones are highlighted, and their temporal profile is shown on panel e. The traces are obtained by convolving the inferred Calcium events with a temporal profile resembling the Calcium indicator dynamics (see Supplementary Fig. 1). Scale bars are 10 s and 10 sd.

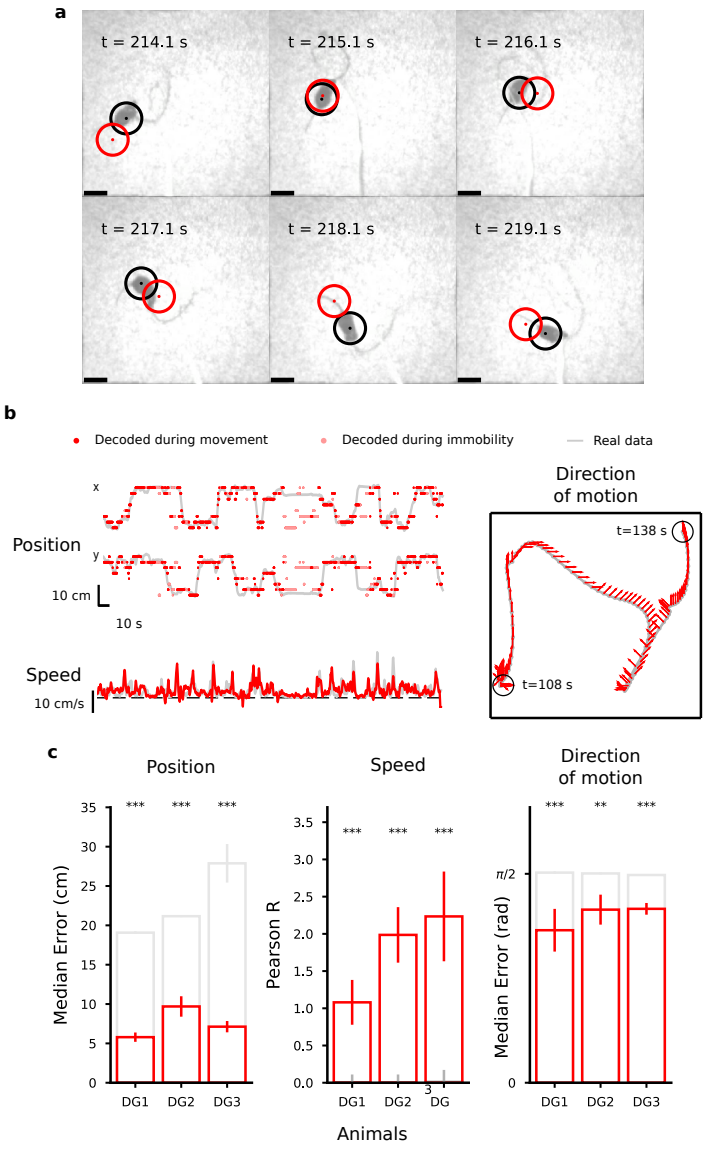

Figure 2. Decoding position, speed and direction of motion. **a**) Selected frames of a video showing the arena and the animal from above. The black circle represents the mouse's actual position and red circle is the decoded position, obtained with a Bayesian decoder that reads the activity of 317 DG GCs. Neural activity has been pre-processed to identify putative calcium events as explained in the Methods. See Supplementary Material for the full video. **b**) Examples of decoding position, speed and direction of motion. Grey lines correspond to the real values of position and speed variables in the top left and lower left panels respectively while the red dots correspond to their decoded values. The time bins marked in light red in for position and direction of movement correspond to moments of immobility that have been excluded from the training data. The grey line in the right panel corresponds to the position of the mouse and the red arrows correspond to the decoded direction of motion in a 30 s time window. **c**) Decoding accuracy for three animals. The decoding error for position and head direction is computed as the median distance between the decoded value in each time bin and the actual value of the decoded variable in the test data. For the direction of motion, the smallest angle between the decoded and the actual value is considered. For velocity, estimated with linear regression (see Methods) we report the Person's R correlation factor. The red vertical bars corresponding to the mean over the 10-fold cross-validation (error bars correspond to s.e.m.). Grey: chance error obtained by decoding from shuffled data in a way that preserves the correlation structure in the data (see Methods and Supplementary Fig. 3).

neural activity during mobility. We therefore removed all the time bins in which the animal was slower than 2 cm/s for a period longer than 1 s. To decode position, we discretized the x and y coordinates of the animal by dividing the 50 cm square of the arena into 64 regions (8 by 8 grid). We then trained a battery of linear classifiers for each pair of discrete locations. Each session was divided into 10 one-minute long intervals, 9 of which were used to train the classifiers and the remaining ones to test them (10-fold cross validation). We used a majority rule[10] to combine the outputs of the linear classifiers as an instantaneous estimate of the animal's location, using the centre of the selected location as the decoded position.

The median decoding error was comparable to the animal size, revealing for the first time that instantaneous position can be decoded from DG GC population activity (Fig. 2). Surprisingly, the decoding accuracy from DG GCs was higher than observed in calcium imaging experiments on CA1 using similar decoding techniques[5] indicating that position is strongly encoded in the DG. Different decoding strategies, such as using linear decoders and Bayesian decoders, as well as decoding from raw Calcium traces or convolved events, produced similar results (see Supplementary Fig. 3 and 4). The decoding error was found to weakly correlate to the speed of movement (see Supplementary Fig. 5) and to be limited mainly by the total number of datapoints available for training the decoder (see Supplementary Fig. 6).

We could also decode the direction and the speed of motion of the animal. Speed was correlated with the overall level of activity and we could decode it using linear regression (Fig. 2b). This is in contrast with recent studies on calcium imaging in DG with head-fixed preparations where speed was found not to correlate with activity[11]. To decode the direction of motion we divided the full range of possible directions into 8 angular bins and labelled time bins according to the instantaneous discrete direction of motion of the mouse (see Supplementary Fig. 7). To our knowledge, this is the first time that decoding of position, direction and speed of motion from populations of DG cells has been reported (Fig. 2c).

To better characterize the neural code, we tried to determine what features of the response properties of individual neurons are important for encoding the variables we could decode. It is important to realize that the neural response properties of an individual cell could be dissociated from its contribution to the accuracy of a decoder that reads out a population of neurons. For example, there could be neurons that are only weakly selective to position, and individually would not pass a statistical selectivity test. However, when combined with other neurons, they can still contribute to position encoding. This is what typically happens when the number of neurons is relatively large, and the number of repetitions for each value of the decoded variable (e.g. the position of the animal) is small. There is an extreme situation in which the decoder might assign the same weight to a neuron which is selective to position and a neuron that has no spatial tuning at all. Although it is unlikely to be encountered in real data, this situation shows nicely how the responses of individual neurons can be completely dissociated from their importance for the decoder. This situation is illustrated in Fig. 3. A simulated animal visits two locations of the arena multiple times. The activity of two hypothetical neurons is represented in the activity space (Fig. 3b), with the X and Y axes representing the activity of the first and the second neuron respectively. At each pass through each location the two neurons have different activity due to the other variables that might also be encoded, e.g., the direction of movement. Each point in the activity plot represents the activity of the neurons in a single pass. The responses of neuron 2 to the two different locations have the same distribution (Fig. 3b,c). A cell with such response properties is untuned to space (a non-place cell) and therefore it is typically considered unimportant for encoding position. However, a linear decoder reading out the population activity to decode the position of the animal can in principle make use of the untuned neuron because of the correlations between the two. While the activity of neuron 1 is only partially predictive of the animal's location (the distributions partially overlap), by reading out neuron 2 together with neuron 1 it is possible to decode position with no errors using a linear decoder. In such situation, the linear decoder would assign equal weights to the two neurons, as shown in Fig. 3c. However, in the real data cells that are tuned

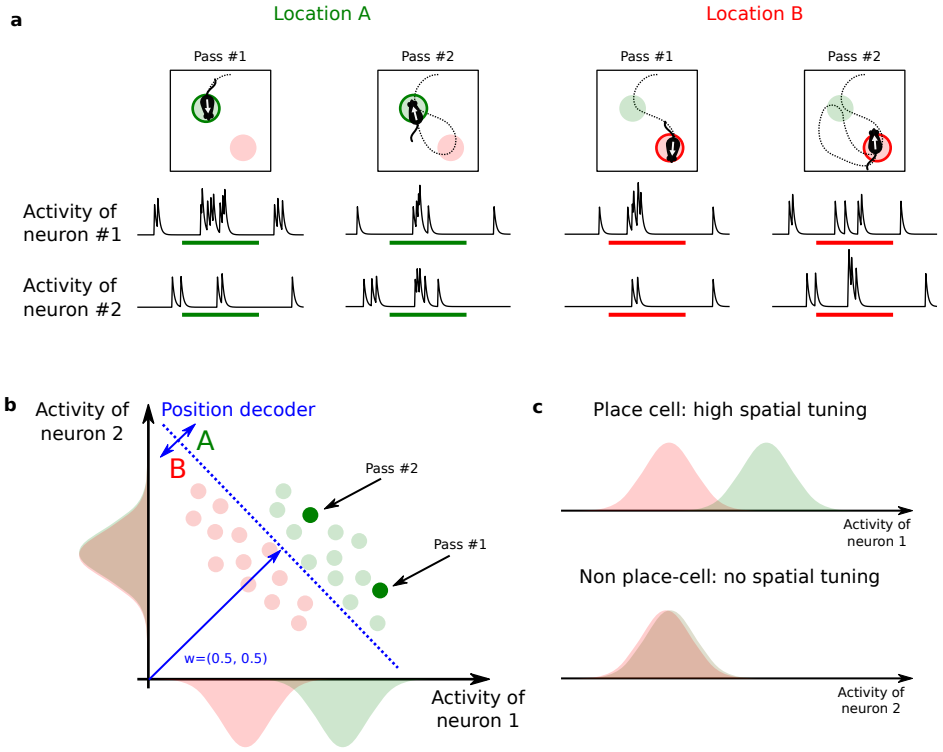

Figure 3. The contribution of untuned cells for encoding position. We show an extreme situation in which one simulated neuron has the same activity distribution when the animal is in two different locations of the arena. Hence the neuron is not selective to position. Nevertheless, for a decoder this neuron can be as important as other selective neurons due to its contribution to the population coding. **a**) Activity of two simulated neurons as a function of time. Top: The simulated animal visits the same discrete location twice (location A in green, location B in red). Bottom: Simulated traces around the time of passage through each location. Different responses for the two neurons are elicited by different experiences, for example due to the different direction of motion. **b**) Example of how place cells and non place-cells can be equally important for encoding the position of the animal. In the scatter plot, the x-axis represents the average activity of the first neuron during one pass and the y-axis is the activity of the second neuron. Each point in the space represents an average population response in a single pass. These responses are in general highly variable and are scattered around their mean values. Despite this variability in the single neuron responses, the two distributions are well separated, making it possible for a linear decoder to read out the animal's position with high accuracy, as we reported. The decoder's weight vector has two equal components corresponding to the importance of the two neurons in encoding position. This example shows that both neurons are important for encoding position despite their very different tuning properties. The distributions of the activity for each individual neuron are plotted along the axes. **c**) When the two neurons are considered in isolation, the situation is different. The tuning curves for neuron 1 only partially overlap, hence neuron 1 is predictive of the animal location to some extent, like the typical tuning of a place cell. The tuning curves of neuron 2 are instead identical and therefore this neuron is not selective to this pair of locations (non place-cell). When considered in isolation, this neuron does not contribute to encoding the position of the animal.

might be like the untuned cell in Fig. 3, or they could simply be so weakly tuned that they would not pass a statistical test. Most likely, untuned cells would have some weak selectivity and they would be at least partially correlated to other cells. In any case, the decoder can use them to improve its accuracy. Analogously a downstream neuron can in principle harness the activity of untuned neurons to readout the animal's position.

In our analysis we took the perspective of such a readout neuron and analysed the weights of our decoder to determine the importance of input neurons in a population for encoding position. We first trained the decoder on each pair of locations and then combined the resulting weights to obtain a single importance index (ω) for each cell (see Methods). Similar methods are used to assess the importance of individual features in a feature space[12, 13] and have been recently used to identify important synapses in learning models[14]. We therefore ranked the neurons according to this importance index and estimated the decoding accuracy for populations of 50 neurons (Fig. 4a) to assess the validity of our approach. The 50 neurons with the largest importance index indeed performed significantly better than the worst 50 neurons, though position could be decoded above chance level even from the worst neurons. The accuracy decreases progressively between the performance for the best and for the worst neurons, validating the method for ranking the neurons on the basis of the importance index. We also controlled that the ranking was stable within the session (see Supplementary Fig. 8).

We then examined the response properties of individual neurons (Fig. 4b). Not too surprisingly, one important feature of an individual neuron is its average activity, which is strongly correlated with the importance index and hence to the overall ability to encode position (Fig. 4c). However, inspection of the firing fields of Fig 4b indicated there were no other obvious properties that predicted whether a neuron is important or not. Moreover, there seem to be spatially tuned neurons (which are all fields not enclosed in dashed red boxes) that rank low and neurons that are not tuned (fields in dashed red boxes) that rank high. The neurons have been considered spatially tuned and called place cells if the spatial information contained in their activity is statistically significant (see Methods for details). The difference between the

Figure 4 ***(following page).*** Ranking neurons according to their contribution to the decoding accuracy for position. **a**) Validation of the importance index. In this figure we show the median error for various selections of 20 cells ranked by their importance index as obtained using the decoder's weight (a, position decoding; e, heading direction decoding). Each point in the plot is aligned to the rank of the first cell in the selection (for example, the first dot corresponds to the selection of the first 50 cells from index 1 to index 50; the shaded region represents the standard error for the 10-fold cross-validation). Grey: chance level and standard error. As expected, the median error for the population of the 50 top ranked (best) cells is much smaller than the median error for the worst last (worst) 50 ones. **b**) Spatial tuning maps for groups of 18 cells ordered by importance index. We ranked the cells using the importance index for position (see Methods). The three groups of best, mid and worst cells are highlighted with the color bands in a for reference. The maps are normalized to the peak rate in each map. Dashed red borders indicate cells that don't pass the criteria for place-cells using a commonly used statistical test for tuning (see Methods). Even among the most important cells there appear some non place-cells (and vice versa). Similarly, some place cells appear in the group of cells with medium and low importance. **c**) Scatter plots and distributions of (top) cell activity and importance for position decoding and (bottom) observed spatial information and importance. Each dot corresponds to one cell in one animal. Pearson's correlation factor ρ between the plotted quantities are reported (***$p<0.001$). A low correlation, although significant is observed between the analysed quantities, suggesting that single cell statistics only partially capture the information available at the population level. **d**) Decoding error for place cells (PCs) and non-place cells (nPCs) for the analysed animals. The position of the animal can be decoded from the activity of either groups of cells with a performance significantly higher than chance (Mann-Whitney U test, ***$p<0.001$). The groups have been normalized to the minimum number of available cells in each group.

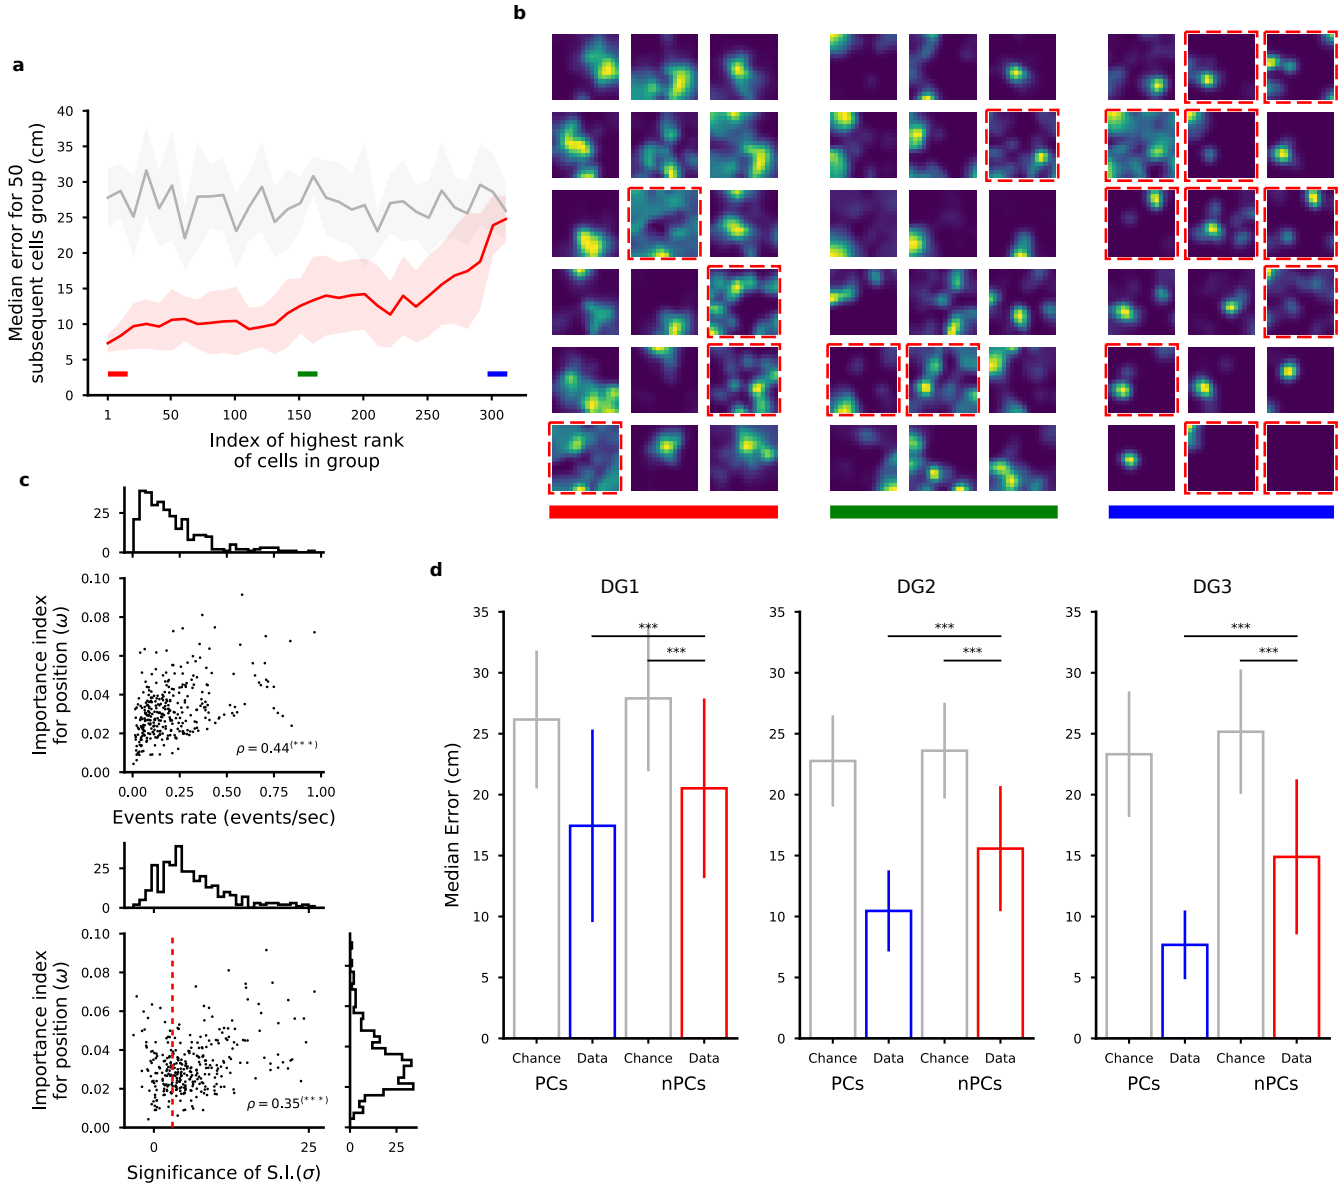

spatial information for the recorded activity and the spatial information obtained for shuffled data, properly normalized, is what we defined as significance of spatial information (SSI). It is indeed a measure used to assess whether a cell is a place cell or not relative to a null distribution16–18. Although there is not a one to one correspondence between SSI and importance index, the two quantities are highly correlated (Fig. 4d) indicating that some individual response properties are at least partially informative about the importance of a cell in encoding position (see Supplementary Fig. 9 and 10). To compute the SSI one has to compute the spatial information and subtract a baseline obtained by shuffling the activity. The spatial information without the baseline subtraction, which is sometimes used as a measure of the tuning of the cells[15], is actually negatively correlated with the importance index (see Supplementary Fig. 9). This is a reflection of the sampling bias problem that affects cells with low activity[16]. Hence, the spatial information without the baseline subtraction is not very informative and should not be considered alone as a measure of the tuning of cells.

From Fig. 4d it is also clear that there are non-place cells that have a large importance index. Positional information can be decoded from these cells alone (Fig. 4e-f). This indicates that non-place cells are not completely untuned, as discussed above. However, because of noise and limited data, they do not pass the statistical test that we adopted to characterize place cells. Among place cells, we also identified single-field and multi-field cells and we found no significant difference of the importance index between the two categories (see Supplementary Fig. 11).

We performed a similar analysis for the direction of movement. In Fig. 5 we show that we can rank the cells according to their contribution to decoding (Fig. 5a) and the important cells are highly heterogeneous in their fields (Fig. 5b). In the case of the direction of movement, we find that a cell's activity correlates with the importance index but not the mutual information between activity and direction of movement (Fig 5c, d).

In conclusion, the importance of a neuron is determined both by its individual response properties and its contribution at the population level. Single neuron response properties are only partially predictive of the importance of a cell for decoding positional information. These results favour the hypothesis whereby populations rather than single neurons are the functional units at work for encoding position, direction of movement and presumably other variables. Since we could decode at least two variables from the neural activities, we were wondering whether there is some form of specialization in which segregated groups of neurons encode only one variable. In Fig. 6 we report the importance index for the direction of movement versus the importance index for position. The plot shows a positive correlation between the two quantities, suggesting that neurons that are important for encoding one variable are also important for encoding the other. This is partially explained by the fact that for both position and direction of movement the most active cells tend to be the most important ones. However, when we regressed out the components explained by the activity, we still found a positive correlation between the importance indexes of the two variables (Fig. 6b). We were wondering whether the positive correlations could be explained by a correlation between the two encoded variables, the direction of movement and position. Supplementary Fig. 12 shows that this not the case. We then focused on those cells that had a high importance for one variable and not for the other. These could indeed be the specialized cells we were looking for. However, we could decode position from the most important cells for encoding direction of motion and vice versa, showing that even the most important cells for one variable carry information about the other variable (Fig. 6c). We conclude that DG neurons have mixed selectivity to the variables we decoded, in line with recent studies both in CA1[17] and in EC[4] as well as in the cortex[1]. Moreover, our results show that part of the elevated variability in the responses of the neurons can be explained by the observation that most of the neurons encode multiple variables.

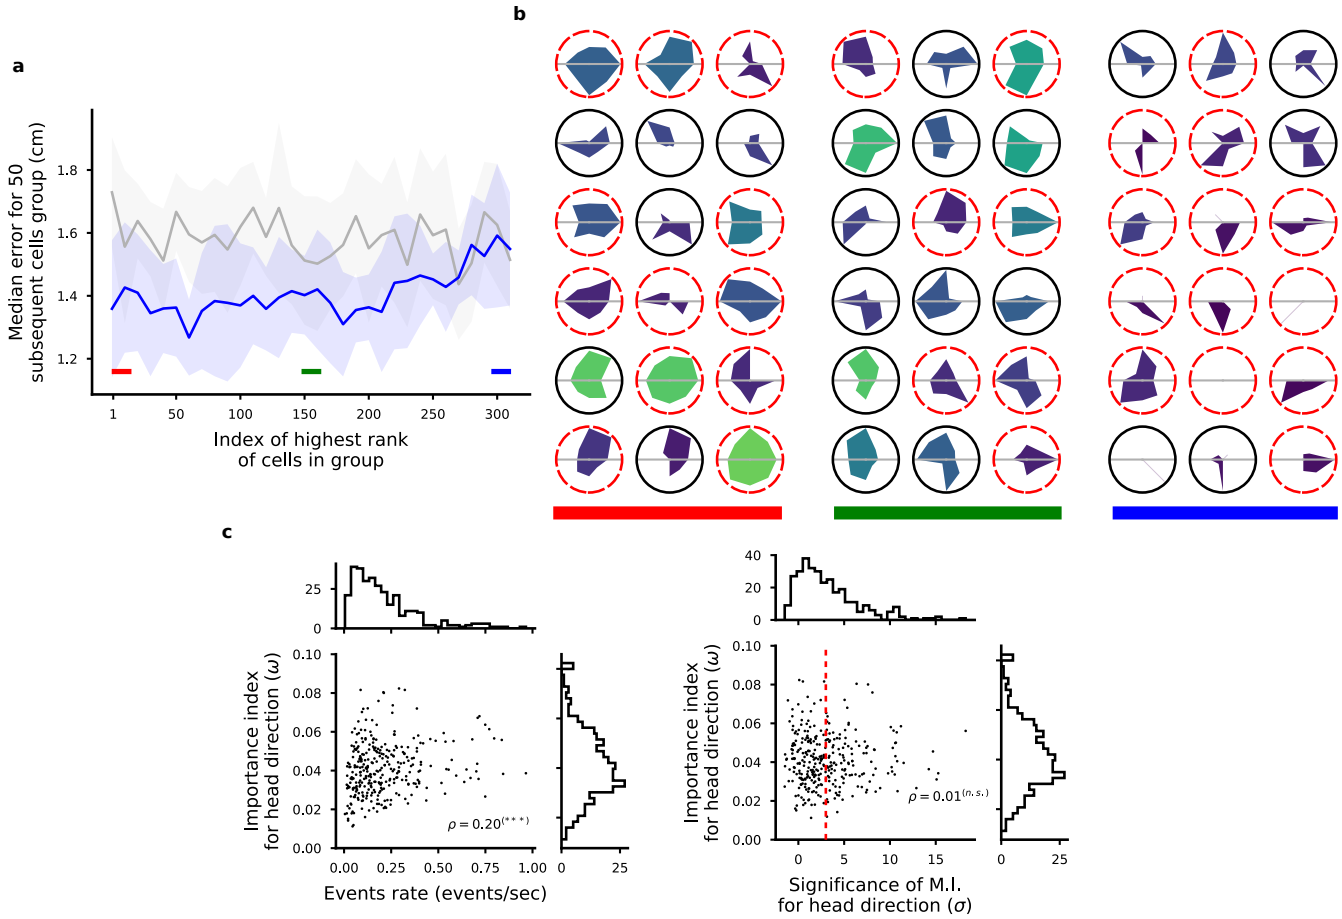

Figure 5. Ranking neurons according to their contribution to the decoding accuracy for head direction. **a**) Validation of the importance index ad in Fig. 4a but we ranked the cells according to the importance index for decoding direction of motion (see Methods). **b**) Tuning maps as in Fig. 4b. Here we show the tuning for direction of motion of single cells as polar tuning maps for groups of 18 cells ordered by importance index. The area colour represents the overall activity of the cell throughout the trial. Dashed red borders indicate cells that don't pass the criteria for significant direction tuning using a commonly used statistical test (see Methods). Even among the most important cells there appear some untuned cells (and vice versa). Similarly, some tuned cells appear in the group of cells with medium and low importance. **c**) Scatter plots of (left) cell activity and importance for position decoding and (right) observed movement direction information and importance for position decoding. Each dot corresponds to one cell in one animal. Pearson correlation factor ρ between the plotted quantities are reported (***p<0.001). A low correlation, although significant, is observed between the analysed quantities, suggesting that single cell statistics only partially capture the information available at the population level.

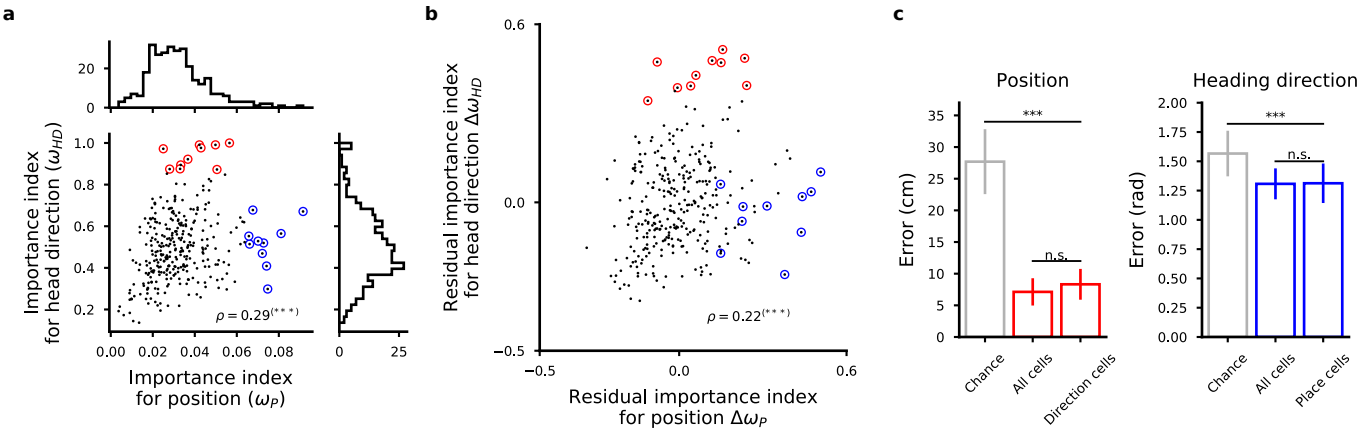

Figure 6. The representations for space and direction of motion are distributed in the DG cells. **a**) Scatter plots of importance index for position and direction of motion. Each dot corresponds to one cell for which we computed the importance index for the variables we decoded. As in a distributed code, the correlation between these quantities is small. We highlight the 10 most important cells for decoding position (blue circles) and the 10 most important cells for direction of movement (red). **b**) Same as in a but from the importance index of each cell we removed the component explained by activity. **c**) Even the most important cells for encoding one variable carry information about the other variable. We show the decoding performance of position (left) and direction of motion (right) using the most important cells for direction (left) and position (right).

Discussion

Neurons in the DG have very sparse and diverse response properties[7,11]. Here we showed that it is possible to decode from the population activity the position the speed and the direction of motion of the animal, despite the seemingly disorganized neural code. Neurons respond to mixtures of the decoded variables as observed in other high cognitive brain areas[1,3]. The information about these variables is highly distributed across neurons to the point that the responses of individual neurons are only weakly predictive of their contribution to the neural code. It is therefore crucial to consider the neurons of the DG as part of an ensemble to assess their importance for processing and transferring information about a particular variable.

One implication of such distributed neural code is that it can be misleading to characterize the function of a brain area based only on the statistics of individual neuron properties. In the specific case of position encoding, for instance, it is not possible to conclude to what extent the position of the animal is encoded in a brain area, only by analysing the tuning of individual cells to space. Indeed, populations of cells whose activities don't pass a selectivity criterion for space encoding, for example through an information theoretical approach, may still encode position via the ensemble activity patterns, as we showed by decoding position and direction of motion from untuned cells.

The population coding rescues the ability of the DG to encode position despite the sparsity of its activity and the variability of its representations. Here we show that indeed even few tens of DG granule cells can encode position with high precision, like CA3 pyramidal neurons. Furthermore, the decoding was accurate even when model training and model test periods were separated by 20 minutes, indicating that at the population level the representations were stable, despite the elevated variability of individual cells (see Supplementary Fig. 13).

All these results indicated that the neural code in the DG is highly distributed and that it is important to analyse it using a population approach. The analysis of the response properties of individual neurons is certainly informative but it is not sufficient to characterize the neural code of a brain area.

Methods

Calcium imaging. Mice were prepared for in vivo calcium imaging as previously described[8]. Briefly, mice were anesthetized with isoflurane and placed in a stereotactic apparatus. Mice were injected in the dorsal DG (1.95AP, 1.4ML, 2.2, 2.1, 2.0, 1.9 DV, ∼90nl per site) with a virus encoding GCaMP6m (AAVdj-CaMKII-GCaMP6m, ∼4X1012vg/ml, Stanford Vector Core). Mice were then implanted with a GRIN lens at -2.0AP, -1.4ML, -1.95 DV. Three weeks after surgery, mice were checked for GCaMP expression with a miniaturized microscope (Inscopix, Palo Alto, CA) and procedures previously described[8]. Anesthetized mice were checked for GCaMP+ neurons and a baseplate was attached to the skull at the optimal imaging plane. One week later, mice were imaged during foraging in an open field task Mice were habituated to the room and enclosure (30min), then 24 hours later mice were imaged as they foraged for sucrose pellets in an open field enclosure (50cm^2). On the day of imaging mice were briefly anesthetized (<5mins) in order to attach the miniscope to the baseplate and were allowed to recover from anaesthesia for one hour before beginning imaging. Imaging frames were recorded with nVista acquisition software (Inscopix, Palo Alto, CA), and time-synced behaviour was acquired EthoVision XT 10. Calcium imaging videos were acquired at 15 frames per second with 66.56 ms exposure.

Behaviour data pre-processing. The behaviour was recorded using a webcam (Logitech) mounted on the ceiling about 3 feet above the arena. The instantaneous position of the animal was then extrapolated from the video using custom code written in Python using the Scikit-image library (version 0.13.0). We first applied a 9 points piecewise affine transformation to correct for barrel camera distortions. We then applied a smoothing filter with a Gaussian profile to reduce the effect of pixel intensity noise due to low lighting and low image resolution and applied a threshold to the gray-scale converted image to get a few contiguous regions of pixels as candidate animal tracking. We then used a method based on the determinant of the Hessian to identify blobs in the pre-processed images and verified that the largest blob was consistently found to be corresponding to the animal silhouette. Hence, we used the centre of the largest blob as the tracked position of the mouse. We further temporally aligned the position data to the imaging data using linear interpolation and smoothed them with a 7 frames time window. Lastly, we identified the time bins in which the speed of the animal was lower than 2 cm/s for more than 1 s and discarded them from the analysis, unless specified.

Signal extraction and spike deconvolution. All calcium movies were initially processed in Mosaic (Inscopix, Palo Alto, CA) for spatial binning and motion correction and subsequently analysed using a recently developed software algorithm written in Matlab (Mathworks) called CNMF-e22. Briefly, the algorithm separates the large, low-frequency fluctuating background components from the signal produced by of multiple sources in the data, allowing the accurate source extraction of cellular signals. It involves a constrained non-negative matrix factorization problem optimized for endoscopic data whereby calcium temporal dynamics and the shape of spatial footprints are used as constraints. It includes 3 main steps which are iterated: obtain a first estimate of spatial and temporal components of single neurons without direct estimation of the background; estimate the background given the estimated neurons' spatiotemporal activity; update the spatial and temporal components of all neurons while fixing the estimated background fluctuations. In each of these steps, manual intervention guided by visual inspection based on temporal profile and spatial footprint shape allowed to further improve the quality of the signal extraction. The result of this process consists of a list of deconvolved calcium events for each cell with associated timestamp and magnitude and the convolved trace obtained with the estimated calcium decay profile. For our analysis, we used an exponential profile to integrate the calcium events in time with a time scale that optimized the position decoding performance, although we empirically observed that results didn't qualitatively change

across a wide range of reasonable integration time values (see Supplementary Material, Fig. 1).

Place fields and heading direction tuning. Place fields for each extracted source were constructed in a manner similar to established method applied to electrophysiology data[7]. We used the calcium events of each cell as its putative spiking activity. We then summed the total number of events that occurred in a given location, divided by the amount of time the animal spent in the location and smoothed using a Gaussian kernel centered on each bin. The rate in each location x was estimated as

$$(x) = \frac{\sum_{i=1}^{n} g\left(\frac{s_i - x}{h}\right)}{\int_0^T g\left(\frac{y(t)-x}{h}\right) dt}$$

where g is a gaussian smoothing kernel, $h = 5$ sets the spatial scale for smoothing, n is the number of events, s_i is the location of the i-th event, $y(t)$ the location of the animal at time t and $[0, T)$ the period of the recording. In this and all subsequent analysis we removed the time bins in which the animal had a speed of less than 2 cm/s for more than 1 s, unless specified otherwise. Similarly, for heading direction tuning, we first discretized the directions of motion into 8 angular bins of 45 degrees each and then computed the mean event rate for each cell in each of the 8 bins.

Spatial information statistics. To quantify the statistical significance of the rate maps we measured their specificity in terms of the information content of cell activity[15]. We used a 16x16 square grid and computed the amount of Shannon information that a single event conveys about the animal's location. The spatial information content of cell discharge was calculated as a mutual information score between event occurrence per cell and animal position or equivalently using the formula:

$$SI = \sum_{i=1}^{N} p_i \frac{r_i}{r} \log_2 \frac{r_i}{r}$$

where i is the spatial bin number, p is the probability for occupancy of bin i, r_i is the mean event rate at bin i and r is the overall mean event rate. We applied the same formula to the direction of motion after discretizing the full angle to 8 bins of 45 degrees. For both measures, we corrected for the sampling bias problem in information measures[16] using shuffled distributions of event occurrences. We assigned randomly chosen time bins to each calcium event of a cell and computed the resulting information content. We repeated this operation for 1000 similar shuffled data for each cell independently. We then compared the original information content to the resulting distribution and considered a cell as a place cell or a head direction cell if the original value exceeded 3 sigmas from the shuffled distribution.

Decoding position. For all the datasets, unless otherwise specified, we used 10-fold cross validation to validate the performance of the decoders. We divided the trial in 10 temporally contiguous periods of equal size in terms of number of datapoints. We then trained the decoders using the data from 9 of them and tested on the remaining data. To decode the position of the animal, we first divided the arena into 8x8 equally sized, squared locations. We then assigned at each time bin the label of the discrete location in which the animal was found. For each pair of locations, we trained a Support Vector Machine (SVM) classifier[18] with a linear kernel to classify the cell activities into either of the two assigned locations using all the identified cells unless specified otherwise. We used only the data corresponding to the two assigned locations and to correct for unbalanced data due to inhomogeneous exploration of the arena we balanced the classes with weights inversely proportional to the class frequencies[19]. The output of the classifiers was then combined to identify the location with the largest number of votes as the most likely location[10]. The decoding error reported corresponds to the median physical distance between the centre of this location and

the actual position of the mouse in each time bin of the test set, unless otherwise specified. To assess the statistical significance of the decoder, we computed chance distributions of decoding error using shuffled distributions of calcium event occurrences. Briefly, for each shuffling, we assigned a random time bin to each calcium event for each cell independently while maintaining the overall density of calcium events across all cells, i.e., by choosing only time bins in which there where calcium events in the original data and keeping the same number and magnitude of the events in each time bin. This method destroys spatial information as well as temporal correlations but I keeps the overall activity across cells. We trained one decoder on each shuffled distribution and pooled all the errors obtained. We finally assessed the statistical significance of the decoding error for the 10-fold cross-validation of the original data by comparing them to the distribution of errors obtained from the shuffled data using the non-parametric Mann-Whitney U test, from which we obtained a p-value of significance. We applied several other strategies for assessing the significance of the decoder to verify that the results didn't depend on the particular strategy adopted (see Supplementary Fig. 3). We also report a comparison between different decoding strategies, including probabilistic approaches such as Näive Bayes, in the Supplementary Material (see Supplementary Fig. 4).

Decoding of the direction of motion. We decoded the instantaneous direction of motion of the animal, expressed in terms of angle with respect to the physical reference of the arena, using the movement data and the calcium traces. We computed the direction of motion using the position in two successive time bins and assigned it to the later time bin. We discretized the values into 8 angles and then applied similar decoding strategies as for position decoding, i.e., we used a battery of linear-kernel SVM decoders to distinguish between pairs of angles after balancing the dataset through class weighting. We report the median error in radiant on the left-out data of the 10-fold cross validation. We applied the methods describe above for position decoding for assessing the statistical significance of the results.

Decoding of speed. To decode the speed of movement of the animal we first computed the speed of motion using two consecutive positions and assigned the computed speed to the later time bin among the two. To decode the instantaneous speed of motion we used Lasso[20], a linear regression analysis method that minimizes the sum of squared errors while selecting a subset of the input cells to improve decoding accuracy and interpretability of the results. For assessing speed decoding chance level, we applied the same strategy as for the other variables except we shuffled all the events of each cells in time independently and convolved them with the same temporal profile we used for the original data.

Importance index. The importance index was introduced to quantify the contribution of each cell in a population to the decoding of a given quantity. We applied a modified version of a traditional method for feature selection in machine learning. In our analysis, a feature of the input space consists of one DG cell. Feature selection is performed using the weights of the decoder after fitting model to the data. In our case, since we employed multiple decoders, one for each pair of physical location in the arena, we introduced a method to combine the weights assigned to the cells by each decoder. We defined the importance index of cell i as:

$$\omega_i = \sum_j \frac{|w_{ij}|}{\sum_k |w_{kj}|}$$

where w_{ij} is the weight of the j-th decoder assigned to the i-th cell.

Code and data availability. The data analysis has been performed using custom code written in Python (version 2.7.12) and routines from the Scipy (ver. 0.19.0), Numpy (ver. 1.11.3) and the Scikit-learn (0.19.1)[19] packages. The source extraction has been performed using Matlab (Mathworks, R2016a) and an early version of the CNMF-e[9] package. The source code and the data are made available upon request.

Acknowledgments

This article is dedicated to Howard Eichenbaum. We thank Loren Frank, André Fenton, Surya Ganguli, Liam Paninski, Jessica Jimenez, William Bialek, Mattia Rigotti and Marcus Benna for insightful discussions and comments on an earlier version of this work and Pengcheng Zhou for his help with the software. FS was supported by the Kavli Foundation. SF was supported by NSF's NeuroNex program award DBI-1707398, the Gatsby Charitable Foundation, the Simons Foundation, the Schwartz foundation, the Kavli foundation. MK was supported by NIMH R01 MH108623, R01 MH111754, and IMHRO/One Mind Rising Star Award.

References

1. Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nat.* **497**, 1–6 (2013). URL http://www.ncbi.nlm.nih.gov/pubmed/23685452http://dx.doi.org/10.1038/nature12160{%}5Cnpapers3://publication/doi/10.1038/nature12160. DOI 10.1038/nature12160.

2. Eichenbaum, H. Barlow versus Hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition? *Neurosci. Lett.* (2017). URL http://dx.doi.org/10.1016/j.neulet.2017.04.006. DOI 10.1016/j.neulet.2017.04.006.

3. Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix: High dimensionality for higher cognition. *Curr. Opin. Neurobiol.* **37**, 66–74 (2016). URL http://dx.doi.org/10.1016/j.conb.2016.01.010. DOI 10.1016/j.conb.2016.01.010.

4. Hardcastle, K., Maheswaranathan, N., Ganguli, S. & Giocomo, L. M. A Multiplexed, Heterogeneous, and Adaptive Code for Navigation in Medial Entorhinal Cortex. *Neuron* **94**, 375–387.e7 (2017). URL http://dx.doi.org/10.1016/j.neuron.2017.03.025. DOI 10.1016/j.neuron.2017.03.025.

5. Ziv, Y. *et al.* Long-term dynamics of ca1 hippocampal place codes. *Nat. neuroscience* **16**, 264 (2013).

6. Fenton, A. A. & Muller, R. U. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Neurobiol. Commun. by Jan Bures Czech Acad. Sci.* **95**, 3182–3187 (1998).

7. Leutgeb, J. K., Leutgeb, S., Moser, M.-B. & Moser, E. I. Pattern separation in the dentate gyrus and ca3 of the hippocampus. *science* **315**, 961–966 (2007).

8. Resendez, S. L. *et al.* Visualization of cortical, subcortical and deep brain neural circuit dynamics during naturalistic mammalian behavior with head-mounted microscopes and chronically implanted lenses. *Nat. protocols* **11**, 566 (2016).

9. Zhou, P. *et al.* Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *eLife* **7**, e28728 (2018).

10. Bishop, C. M. *Pattern Recognition and machine learning* (Springer New York, 2006).

11. Danielson, N. B. *et al.* Distinct Contribution of Adult-Born Hippocampal Granule Cells to Context Encoding. *Neuron* **90**, 101–112 (2016). DOI 10.1016/j.neuron.2016.02.019.

12. Haufe, S. *et al.* On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage* **87**, 96–110 (2014). URL http://dx.doi.org/10.1016/j.neuroimage.2013.10.067. DOI 10.1016/j.neuroimage.2013.10.067.

13. Mladenić, D., Brank, J., Grobelnik, M. & Natasa Milic-Frayling, I. Feature Selection using Linear Classifier Weights: Interaction with Classification Models. *Proc. 27th annual international ACM SIGIR conference on Res. development information retrieval* 234–241 (2004). DOI 10.1145/1008992.1009034.

14. Zenke, F., Poole, B. & Ganguli, S. Continual learning through synaptic intelligence. In Precup, D. & Teh, Y. W. (eds.) *Proceedings of the 34th International Conference on Machine Learning*, vol. 70 of *Proceedings of Machine Learning Research*, 3987–3995 (PMLR, International Convention Centre, Sydney, Australia, 2017). URL http://proceedings.mlr.press/v70/zenke17a.html.

15. Skaggs, W. E., McNaughton, B. L. & Gothard, K. M. An information-theoretic approach to deciphering the hippocampal code. In *Advances in neural information processing systems*, 1030–1037 (1993).

16. Panzeri, S., Senatore, R., Montemurro, M. A. & Petersen, R. S. Correcting for the Sampling Bias Problem in Spike Train Information Measures. *J Neurophysiol* **98**, 1064–1072 (2007). DOI 10.1152/jn.00559.2007.

17. Meshulam, L., Gauthier, J. L., Brody, C. D., Tank, D. W. & Bialek, W. Collective behavior of place and non-place neurons in the hippocampal network. *Neuron* **96**, 1178–1191 (2017).

18. Cortes, C. & Vapnik, V. Support-vector networks. *Mach. learning* **20**, 273–297 (1995).

19. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2012). URL http://dl.acm.org/citation.cfm?id=2078195{%}5Cnhttp://arxiv.org/abs/1201.0490. DOI 10.1007/s13398-014-0173-7.2. 1201.0490.

20. Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *J. Royal Stat. Soc. Ser. B (Methodological)* **57**, 267–288 (1996).