SIMNETS: a computationally efficient and scalable framework for identifying sub-networks of functionally similar neurons

Jacqueline B. Hynes [1, 2], **David M. Brandman**,[2, 3], **Jonas B. Zimmerman**[4], **John P. Donoghue**[1, 2, 4, 5], **Carlos E. Vargas-Irwin** [*, 1, 2]

1 Department of Neuroscience, Brown University, Providence, RI, USA.
2 Robert J. and Nancy D. Carney Institute for Brain Science, Brown University, Providence, RI, USA
3 Department of Surgery (Neurosurgery), Dalhousie University, Halifax, NS, Canada.
4 WYSS Institute, Chemin des Mines 9, CH-1202, Geneve, Switzerland.
5. Department of Engineering, Brown University, Providence, RI, USA.

∗ Correspondence to: carlos_vargas_irwin@brown.edu

Keywords: functional sub-networks, neuron clustering, spike train distances, neural ensembles, large-scale recordings, dimensionality reduction,

Abstract

Recent technological advances have now made it possible to simultaneously record the activity of thousands of individual neurons in the cortex of awake behaving animals. However, the comparatively slower development of analytical tools capable of handling the scale and complexity of large-scale recordings is a growing problem for the field of neuroscience. Here, we present the Similarity Networks (SIMNETS) algorithm: a computationally efficient and scalable method for identifying and visualizing sub-networks of functionally similar neurons within larger, simultaneously recorded ensembles. While traditional approaches tend to group neurons according to the statistical similarities of their inter-neuron spike train outputs, our approach begins by mathematically capturing the intrinsic relationship between the spike train outputs of each neuron across experimental conditions, before any comparisons are made between neurons. This strategy estimates the intrinsic geometry of each neuron's output space and allows us to capture the information processing properties of each neuron in a common format that is easily compared between neurons. Dimensionality reduction tools are then used to map high-dimensional neuron similarity vectors into a low-dimensional space where functional groupings are identified using clustering and statistical techniques. SIMNETS makes minimal assumptions about single neuron encoding properties; is efficient enough to run on consumer-grade hardware (100 neurons < 4s run-time); and has a computational complexity that scales near-linearly with neuron number. These properties make SIMNETS well-suited for examining large networks of neurons engaged in complex behaviors. We validate the ability of our approach to detect functional groupings in a population of synthetic neurons with known ground-truth, as well as three publicly available datasets of ensemble recordings from primate primary visual and motor cortex and the rat hippocampal CA1 region.

Introduction

The neural computations underlying complex sensory, cognitive, and motor information processing are thought to emerge from the interactions of vast networks of functionally interrelated neurons. Within these networks, smaller sub-networks of neurons ('sub-nets') engaged in similar information processing tasks have been proposed to embody the computational units that support specific functions including perceptual integration, memory storage/retrieval, and dexterous motor control[1–3]. Identifying functional sub-nets would greatly simplify the process of tracking information flow in cortical circuits, modeling population-level neural dynamics, and ultimately understanding the general principles of neural computation[4–7]. While it is now possible to record ever larger neural populations, detecting functional groupings of neurons and characterizing their computational operations has proven notoriously difficult because of the scale of data processing involved and the lack of accepted mathematical tools to partition large networks into smaller functional components[8].

One of the critical challenges lies with the selection of an appropriate quantitative definition of 'functional similarity' across neurons. Motivated by theories of Hebbian cell assemblies, several neuron clustering approaches have relied on using synchrony or firing rate covariations to detect functional associations between neurons[9–14]. One widely discussed hypothesis proposes that synchronously-active neurons might serve as an independent coding dimension to facilitate perceptual or cognitive integration of the information encoded in the firing rates of individual neurons[15–17]. Although several studies have observed synchronous and correlated activity between neurons in multiple brain regions, discrepant reports regarding the functional and statistical significance of the features of neural activity have led to some doubts regarding the usefulness of this approach for detecting functionally relevant ensemble motifs[18,19]. While focusing on spike-rate and spike-time covariations is intuitive, cross-correlation methods do not scale well for large datasets and firing rate covariation measures could limit the complexity of the functional relationships that can potentially be detected between neurons. More specifically, these methods prioritize grouping neurons according to the similarity of their spiking statistics – as opposed to their information processing properties. Although the underlying premise of these methods is that similar inter-neuron spike patterns imply similar information processing properties[9,11,14,20], a growing body of work suggests that this may be an oversimplified view[21].

Recent work has shown that individual neurons in higher-level brain areas[22,23], motor areas[24–26], and even primary sensory areas[27], can exhibit highly heterogeneous and complex responses dynamics, both across conditions and neurons[21]. Historically, these features of single neuron spiking activity were interpreted as biological noise; however, these studies suggest that trial-to-trial variability and temporal complexity are important features of the information coding operations taking place across the network. This work also suggest that

measures of trial-averaged spike rate or spike time covariations may neglect important aspects of a neuron's activity that can reveal a functional association to other neurons within in sub-network. Here, we propose a more general strategy for identifying groups of neurons with similar information processing properties.

SIMNETS: a novel mathematical framework to identify functional neuronal sub-ensembles

The SIMNETS algorithm is designed to map a set of spike trains generated by a single simultaneously recorded neuron into an abstract metric space in such a way that the structure of the space captures the neuron's information processing characteristics (see below). Rather than comparing the spike train outputs between pairs of neurons directly, we compare the intrinsic structure[27,2] of the neurons' output state-spaces, i.e., the self-similarity pattern of each neuron's spike trains across conditions[8].

We envision each neuron as performing an unknown operation on a set of high dimensional inputs (potentially 1000's of synaptic inputs carrying sensory and/or time-varying internal signals) (Fig 1*a*). To be useful as a computational element, each neuron should have a relatively consistent internal mapping between inputs and outputs, allowing for stochasticity in spike generation and the potential to change the mapping over time through learning. When examined within the context of the observed population, a neuron may be insensitive to certain changes taking place across the network, such that the neuron's spike train pattern appears similar across many different global network patterns. On the other hand, other global network patterns may elicit dramatic changes in the neuron's spike outputs that will differ depending on the strength of the synaptic inputs across time. Thus, the observed changes in the spike train patterns of a single neuron across different network states will highlight the differences between some global network states and generalize over others. The key insight is that it is possible to represent the information processing characteristics (i.e., the input-output relationship) of a neuron by examining the similarities and dissimilarities between its spike train outputs across different points in time. Note that this approach makes it possible to compare the operations performed by simultaneously recorded neurons on a trial-by-trial basis, without requiring explicit knowledge of the type of function they may be computing.

Our neuron sub-net identification strategy is built upon the premise that 'computationally equivalent' neurons will generalize and differentiate across the same subset of network states (i.e., trials). A spike train similarity analysis framework[30,31] can be used to project the set of activity patterns generated by each neuron into an abstract metric space, such that near-by points correspond to similar output states (i.e., similar spike trains) and far-away points correspond to dissimilar output states (i.e., dissimilar spike trains) (Fig. 1*b*). Mathematically, we can represent the relationships between the set of spike trains originating from a given neuron across a set of events of interest (e.g., stimulus presentations), using a pairwise distance matrix, where each entry represents the similarity between a pair of spike trains (Fig. 1*c*). We refer this type of matrix as a Spike Train Similarity (SSIM) matrix (see Online Methods for more details). The spike train distance matrix can be thought of as a high-dimensional representation of the relationship between the neuron's outputs across experimental conditions[30] (see Supplementary Fig. 1-2 for further demonstrations). This general strategy captures the intrinsic geometry of each neuron's output space and allows us

to capture the information processing properties of each neuron in a format that is easily compared between neurons (Fig. 1*d-f*).

Geometric models of similarity data have a long history of application in the field of psychology, where they have been used to model the perceptual relationships between sensory stimuli, i.e. perceptual metric-space[32]; however, it is only more recently that this approach has found application in the field of Neuroscience, where it has been successfully used to model the relational structure of neuronal ensemble activity patterns[33–35] and fMRI activity patterns[36].

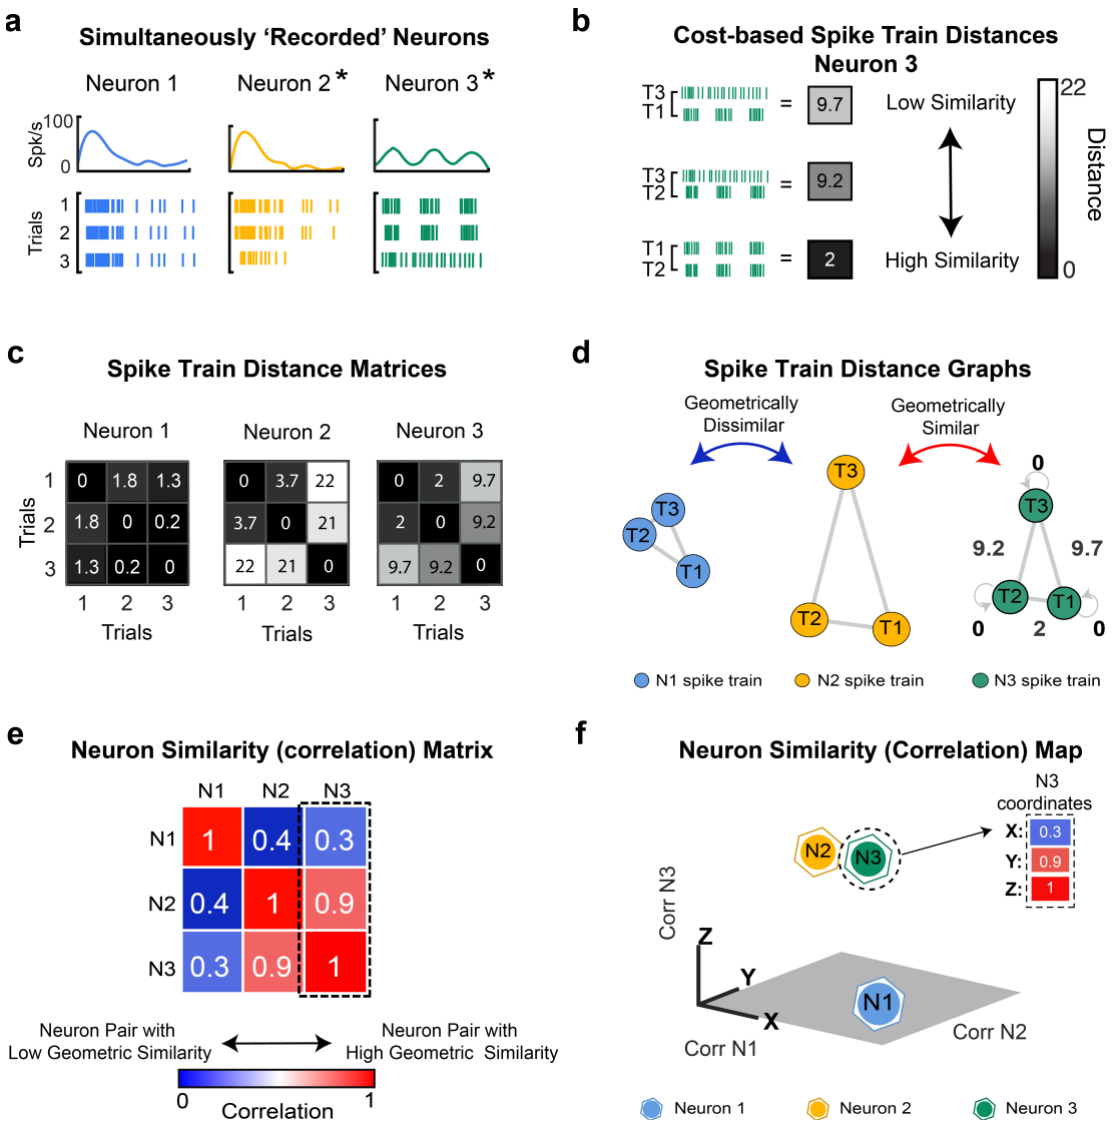

Figure 1| Assessing functional relationships between neurons by comparing state spaces generated from within-neuron single trial spike train distance measures

a. Simulated spike train data for three 'simultaneously recorded' synthetic neurons (N) during repeated presentations of a stimulus: N2 was designed to exhibit similar trial-averaged firing rates to N1 and similar trial-to-trial response variability to N3. *Top row*: trial-averaged firing rate functions for each neuron. *Bottom row*: spike train outputs of each neuron in response to three repetitions of the same stimulus. Note*, both N2 and N3 exhibit a spike pattern modulation on trial 3 to an 'unknown' experimental variable, which is expressed as a reduction in spike rate (N2) and a change in the temporal pattern of the output (N3). **b.** Victor-Purpura (VP) spike train metric applied to N3 spike trains: the metric assigns an 'edit cost' to spike train pairs in accordance to the similarity of their spike patterns. Note that several spike train similarity metrics have been

proposed in the literature. **c.** The relationship between a set of spike trains originating from a common neuron is summarized as a pairwise distance matrix; each entry corresponds to the edit-distance between two spike trains. We refer to these data structures as spike train similarity (SSIM) matrices **d.** Visualization of SSIM matrix using a graph representation (**c**); a separate graph is plotted for each matrix. Each point represents a spike train from a given neuron and the line connecting pair of points corresponds to the distance values assigned to the spike trains. Note that the sizes of the spike train graphs are different across neurons, but that shape of N2 and N3 graphs are similar, i.e., covariation in their side lengths. (**e**) Pairs of single neuron SSIM matrices are compared using Pearson's correlation and represented as a Neuron Similarity (NS) matrix. Note that the diagonal represents the comparison of each neuron to itself, and so always equal to one. Each NS matrix column (broken line) represents the correlations between one neuron versus all others. Each of these column vectors can be used to map a neuron into an *N*-dimensional neuron space. **f**. Scatter plot showing the data contained in the NS matrix (**e)**: each point corresponds to a single neuron with coordinates defined by a *N* dimensional vector of correlation values, i.e., the NS matrix column. Note that N2 and N3 are located next to each other in the map because of the high correlation between their SSIM matrices, which is ultimately dependent upon the geometric similarity of their spike train output spaces.

Our proposed algorithm calculates the pairwise spike train distances between all *S* spike trains generated by a given neuron (Fig. 2; *step 1*), and then compares each neuron's *S*x*S* SSIM matrix to that of every other neuron (Fig. 2; *step 2*). This can be accomplished by comparing their SSIM matrices using standard correlation statistics, such as Pearson's correlation. The resultin*g* pairwise correlation measures are represented as a single *N*x*N* Neuron Similarity (NS) matrix, where each column of the matrix can be viewed as a vector that represents the functional similarity of a given neuron to all others *N*-1 neurons in the population (Fig. 2; *step 2*). Standard dimensionality reduction techniques, e.g. multidimensional scaling, t-distributed stochastic neighborhood embedding[37] can be used to project these neuron similarity vectors into a low-dimensional Neuron Similarity (NS) map such that neurons are positioned according to their information processing properties (Fig. 2; *step 3-4*). Applying dimensionality reduction makes the data easier to visualize and facilitates statistical analysis. Overall, this representation reduces the problem of identifying functional sub-nets to one of detecting clusters of neurons within the NS map. This step can be accomplished using standard clustering algorithms (e.g. k-means) and validated using a shuffle-based statistical test that relies on shuffling the rows/columns of the SSIM matrices to avoid false cluster discovery (Fig. 2; *inset*; see Methods and supplementary Fig. 4 for details). We call this new strategy for identifying sub-nets of neurons with similar informational properties SIMNETS. Note that there is a wide choice of (1) similarity metrics for spike trains, (2) dimensionality reduction algorithms, and (3) clustering algorithms that can be employed within the proposed analytical framework.

Unlike methods based on measures of synchronous spiking activity[9,38,39], SIMNETS identifies neurons with similar functional properties, even if functionally interrelated neurons exhibit diverse firing statistics (i.e., different encoding schemes). Critically, unlike other pairwise methods [9,11,38], SIMNETS is well suited for studying datasets with large numbers of neurons and relatively small numbers of experimental trials. The computational cost of the generating the high-dimensional neuron embedding (Fig. 2, *step 2*) grows nearly linearly with the number of neurons (but quadratically with the number of trials). Further, SIMNETS can be implemented without *a priori* knowledge of neural tuning functions or trial labels, making it particularly useful for the analysis of complex, naturalistic behaviors.

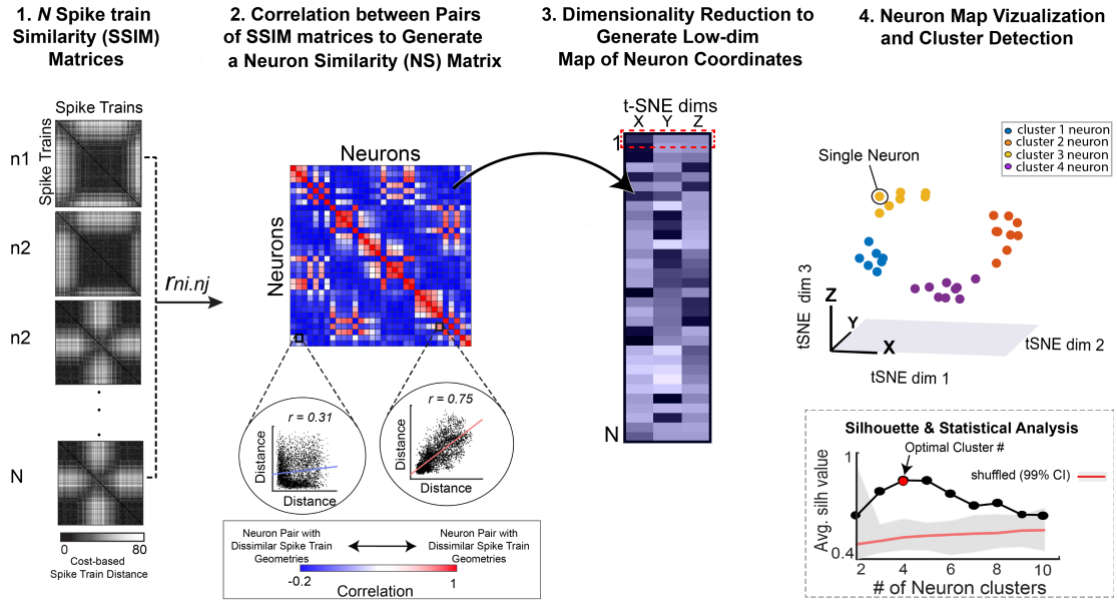

Figure 2|The SIMNETS algorithm can be used to identify functionally related groups of neurons.

The basic steps of the algorithm (demonstrated using synthetic data) are as follows: Step 1. Generate *N* Single-neuron Spike Train Similarity (SSIM) matrices**:** extract *S* spike train events of equal duration from *N* simultaneously recorded neurons and calculate the pairwise spike train distances between all spike trains using a predetermined metric, e.g., Victor-Purpura. Step 2. Neuron Similarity (NS) Matrix: pairwise correlations between single neuron SSIM matrices are used to generate an *N*x*N* NS matrix. *Left Circular Inset*: example neuron pair with a low SSIM correlation value indicating dissimilar spike train geometries. *Right Circular Inset*: neuron pair with a high SSIM correlation value indicating similar spike train geometries. Step 3: NS Matrix Dimensionality reduction: project the *N*x*N* correlation matrix down into a smaller number of *d*x*N* dimensions using standard techniques, such as t-SNE. Each row represents the coordinates of a single neuron (red broken line) in a low-dimensional space. The t-SNE parameter, *perplexity*, controls the number of effect nearest neighbors factored into the mapping of each neuron from the high- to low-dimensional space (see Methods for details). Step 4. Visualization of Neuron Similarity (NS) map and Cluster Detection: each colored point represents a single neuron with *x*, *y*, *z* coordinates that correspond to a single row of the low-dimensional data structure shown in step 3. Clustering: neuron clusters (denoted using color) are detected in a 10-*d* space using agglomerative clustering methods (e.g., *k*-means clustering). *Inset*: Optimal cluster number is estimated as the maximum average silhouette value for a range of test cluster number, where a high silhouette value indicates good cluster separation. The statistical significance of the estimated optimal cluster number is determined from a null-distribution of silhouette values (gray band), which was obtained by repeating steps 2 - 3 on shuffled SSIM matrices (step 1) over multiple iterations (see Methods for more details). Neurons assigned to same cluster will have similar spike train geometries.

Results

Here, we apply SIMNETS to four different datasets to validate its ability to detect sub-nets of functional similar neurons. We first apply the algorithm to a dataset of synthetic neurons with known ground-truth functional ensembles. Next, we apply the algorithm to three datasets of multi-electrode, extra-cellular single-unit (i.e., neuron) recordings from non-human primate primary visual cortex (V1)[40] and primary motor cortex (M1) [41], as well as the rat CA1 hippocampal region[42] (see Methods for more details on datasets). For the V1 and M1 datasets, the SIMNETS neuron functional maps are validated against the estimated computational properties of the neurons calculated using parametric tuning models. We use the rat CA1 dataset to demonstrate how SIMNETS can be used for exploring the functional properties of neurons when the tuning functions of the neurons that are not easily quantifiable or are unknown. Performance of the Synthetic, V1, M1, and CA1 datasets are compared

against various alternative neuron similarity measures that representative of traditional methods that use compare the spike patterns between neurons directly[14,39,43–45].

Synthetic Neuron Population — clustering functionally similar neurons exhibiting distinct firing patterns

We applied SIMNETS to simulated spike train data from a population of 180 synthetic neurons comprised of 3 functionally distinct 'ensembles' (E_1, E_2, E_3). Each ensemble (n = 60, neurons) was designed to represent a sub-group of computationally equivalent neurons that exhibited heterogenous firing patterns. Sub-groups of neurons within ensemble E_i responded to a common 'preferred' test condition through either a change in spike rates (n = 20, neurons), a change in the precise timing of their spikes (n = 20, neurons), or a change in both spike rates and spike timing (n = 20, neurons) (Fig. 3*a*) (see Methods for more details). We simulated 30, one-second spike trains for each of the 180 neurons, which included 10 repetitions of each stimulus (S = 30, spike trains per neuron).

SIMNETS was applied to the resulting NxS spike trains using three different temporal accuracy settings for the Victor-Purpura spike train metric: 5ms (q = 200), 100ms (q = 10), and pure rate code (q = 0). The temporal accuracy setting dictates the temporal accuracy at which two spikes should occur in order to be considered as occurring 'at the same time'. As expected, with a setting of q = 0, the neurons operating with a rate-based encoding scheme ('rate-code' neurons) and rate/temporal-based coding scheme ('mixed-code' neurons) are grouped into three functionally distinct clusters in the NS map, while the functionally dissimilar 'temporal-code' neurons form a single cluster at the center of the map (Fig. 3c). As the value of q increases, the algorithm becomes sensitive to differences in spike timing in addition to the total number of spikes. Using these settings, SIMNETS correctly groups all neurons into three distinct clusters that reflect the ground-truth functional ensemble assignments (Fig. 3*d*). At very high q values, the NR map shows sub-groupings within each of the main clusters that reflect the coding properties of the neurons (Fig. 3*e*); however, the optimal number of clusters remain in agreement with the ground-truth functional ensemble assignments. By specifying a higher partition value for the k-means clustering step of the algorithm (e.g., k = 9), it is possible to confirm that the sub-groupings within the detected clusters are defined by the coding properties of the neurons (data not shown). For a demonstration of the interaction between the cluster number and the SIMNETS hyper parameters, *perplexity* and q, see Supplementary Fig. 5.

SIMNETS is designed to identify neurons in the same sub-net as being more similar to each other than to neurons in a different sub-net. In order to quantify this trend, we compared the distribution of similarity estimates (entries in the NS matrix) *within* and *between* the artificially generated ensembles (Fig. 3*f*). Within-ensemble similarity was significantly higher than between-ensemble values in all cases (Mann-Whitney $p<0.001$). For q values > 0, there was no overlap between the two distributions, indicating complete separation of the functional classes. Our results demonstrate that the SIMNETS algorithm can accurately separate neurons according to their computational properties, even if they employ different coding schemes to represent information.

In order to demonstrate the potential pitfalls of traditional approaches that directly compare

spike trains *between* neurons on a trial-by-trial basis[14] we applied two additional methods that rely on between-neuron spike train comparison to the synthetic data: the 'Direct Comparison (DC) method and the Firing Rate Covariance (FRC) method[44].

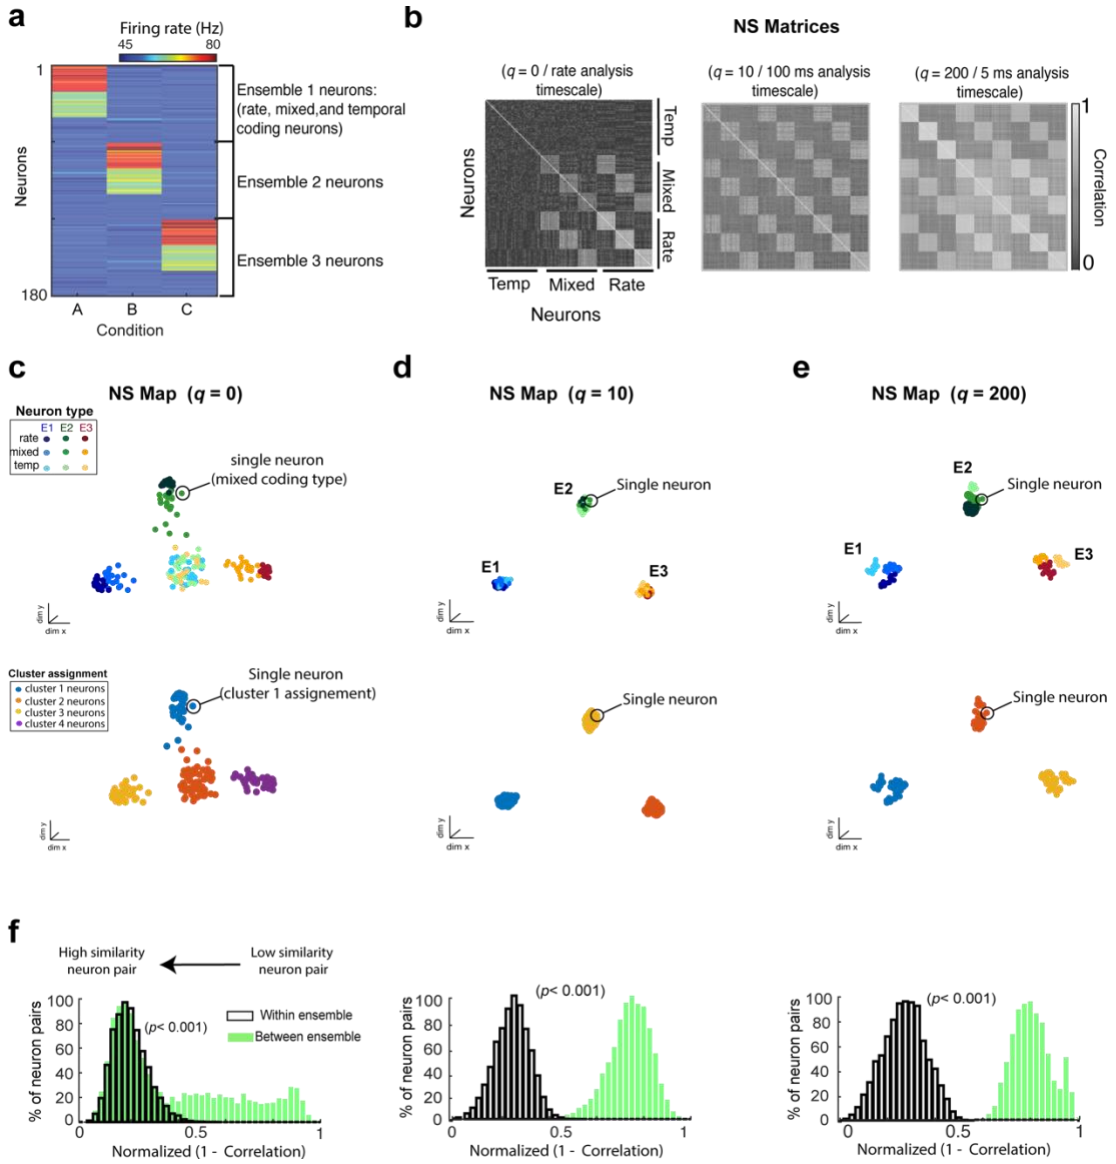

Figure 3| SIMNETS detects ensembles of functionally similar neurons in population of synthetic neurons with mixed coding schemes.

a.Trial-averaged firing rates for a population of 180 synthetic neurons as a function of three different stimulus condition. Neurons are ordered along the y-axis according to ensemble membership and by encoding strategies. Each of the three ensembles (E1, E2, and E3) are made up three neuron sub-groups defined by their different encoding formats, i.e., rate-based coding scheme, temporal-based coding scheme, or a mixed (temporal/rate-based) coding scheme. **b**. SIMNETS NS matrices for three different analysis temporal sensitivity values: rate (q = 0), 100 ms (q = *10)*, 5 ms (q = *200)*. Color-bar represent a range of correlation values where higher correlation values (light pixels) indicate a high similarity neuron pair and low correlation values indicate a low similarity neuron pair. **c** - **e.** Low-dimensional NS maps for each of the three temporal sensitivity values: q = 0 (**c**) q = 10 (**d**), q = 200 (**e**). Each dot represents a single neuron. *Top row:* colors indicate the neurons ensemble membership (E1, red/orange; E2, green; and E3, blue color-hue range) and whether it utilizes a rate, mixed, or temporal coding scheme (dark, mid, or light hue, respectively). *Top row:* colors indicate the neurons ensemble membership (E1, red/orange; E2, green; and E3, blue color-hue range) and whether it utilizes a rate, mixed, or temporal coding scheme (dark, mid, or light hue, respectively). *Bottom row:* same map as top row but colors indicate the neuron's *k*-means cluster assignments. **f**. Histograms of normalized 'Within-' and 'Between-ensemble' correlations from SIMNETS

NS matrices (**b**). Non–overlapping Within/Between distributions (i.e., *middle and right column*) corresponds to good separation between each of the three ground-truth ensembles in the NS Maps (*ranksum, p<0.001).*

The DC method also uses Victor-Purpura spike train metrics, but compares spike trains from different neurons directly, i.e., without generating SSIM matrices as an intermediate step. In this case, each entry of the resulting *N*x*N* matrix is the sum of the spike train distances between matching trials across a neuron pair. A neuron pair that generates similar spike train outputs on matching trials will have a low sum of spike train distances, whereas a neuron pair that generates dissimilar spike train outputs on matching trials will have a high sum of spike train distance value (Fig. 4*a*, *columns* 1-3; see Methods for details).

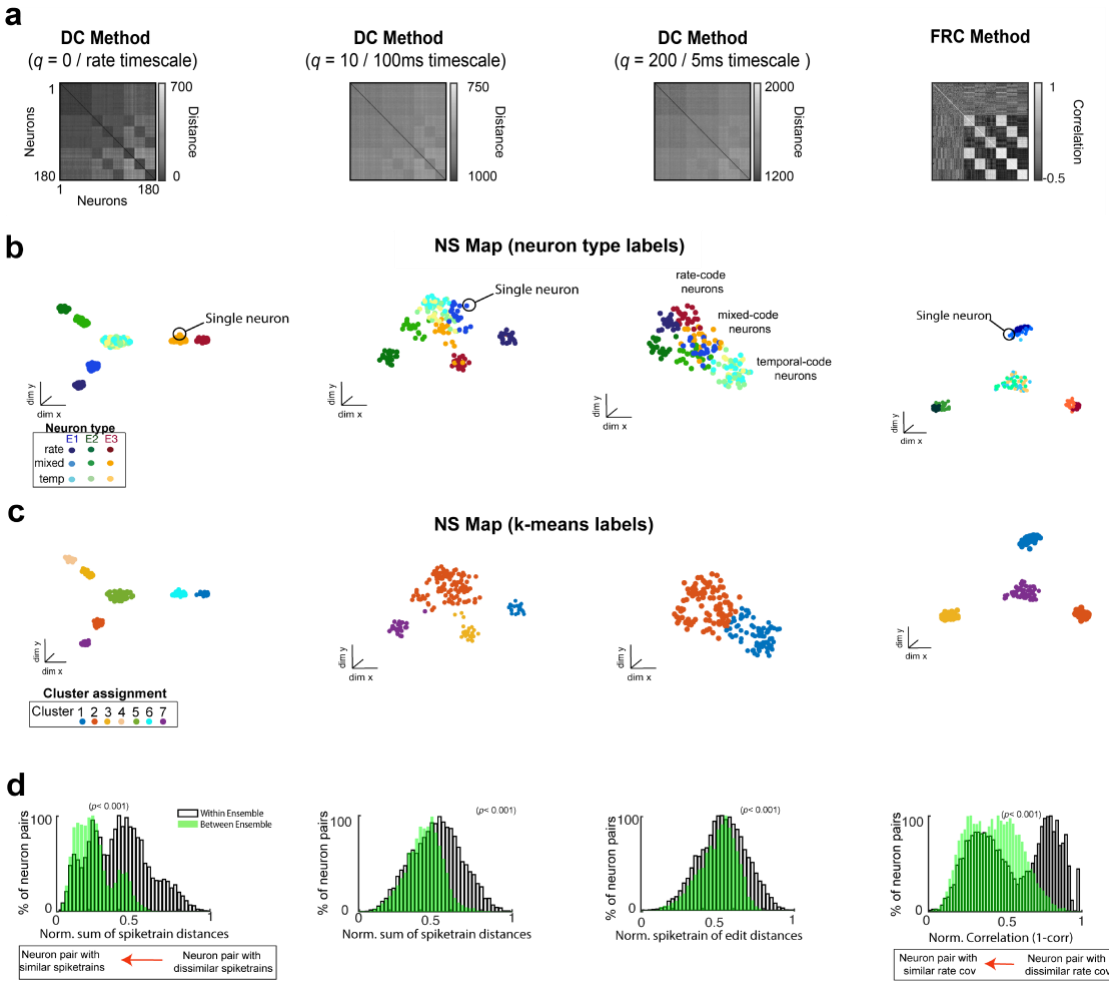

Figure 4| Direct Comparison (DC) method and Firing Rate Covariance (FRC) method fail to detect functional ensembles in population of synthetic neurons.

a. DC *NxN* distance matrices shown for three different analysis temporal sensitivity values (*columns 1-3*): rate, 100 ms, 5 ms, and the FR Covariance *NxN* correlation matrix (*column 4*). DC matrix color-bars represent a range of summed spike train distance values for a neuron pair, where smaller distance values (dark pixels) indicates a neuron pair with similar trial-by-trial spike train outputs. The FRC matrix color-bar represent the range of firing rate correlation values for the neuron pairs, where higher correlation values (light pixels) indicates a neuron pair with a strong relationship between their trial-by-trial firing rate outputs. **b.** DC and FRC NS maps with neuron-type color labels: low-dimensional representation of the *NxN* matrices (shown in **a)**. Each point represents a single neuron and the different colors indicate each neuron's ensemble membership (E1, E2, or E3) and coding scheme (color notation same as Fig. 3). **c.** DC and FRC NS maps with neuron *k*-means cluster assignment labels. K-means clusters were detected using the silhouette analysis and shuffle statistical test (step 4 of SIMNETS algorithm). **d.** Histogram of normalized 'Within-' and 'Between- ensemble' values from DC distance matrices (*Columns 1-3*) and FRC correlation matrix (*Column 4*) (**a**) demonstrates a poor cluster separation for neurons belonging to different ground-truth ensembles.

For the FRC method, each NxN correlation matrix entry represents the between-neuron single trial firing rate covariation for a given neuron pair (see Methods; Kiani et al (2015)). Neuron pairs that exhibit a strong relationship between their firing rates across individual trials will have high correlation values (Fig. 4*a*, *column* 4). Dimensionality reduction (Fig. 4*b*) and cluster detection (Fig. 4*c*) was performed on the DC and FRC matrices in accordance with steps 3 - 4 of the SIMNETS algorithm (i.e. using t-SNE and k-means). Overall, the DC and FRC method failed to cluster functionally similar neurons into the three ground-truth functional ensembles (Fig. 4*c*). The distributions of similarity estimates for neurons within and between ensembles (Fig. 4*d*) displayed broad overlaps, reflecting the poor separation between functional ensembles. However, the FRC method was successful in clustering functionally equivalent rate-coding and mixed-coding neurons into their correct functional ensembles (Fig. 4*d*). Our results demonstrate that grouping neurons based on the similarity of their spike train outputs or firing rate covariations does not necessarily reflect their informational content (and presumed computational properties).

V1 Neuron Population – clustering real neurons with known tuning functions

We next analyzed a previously described dataset of 112 Macaque V1 neurons simultaneously recorded using a 96-channel electrode array during the presentation of drifting sinusoidal gratings [40,46] (Fig. 5). We extracted 1 second of spiking data from the first 30 repetitions of each stimulus (S = 360), starting 0.28 seconds after stimulus onset (Fig. 5*a*). Each neuron's receptive field orientation ('preferred' orientation) was estimated by finding the orientation that maximizes a Gaussian function fitted to the stimulus-dependent firing rates (Fig. 5*b*) (see Methods for more details).

We examined the NS map produced using SIMNETS in order to determine if it accurately captured the functional relationships between neurons (Fig. 5*d*-*e*). A circular-linear correlation ($r$cl) analysis shows a significant positive relationship between preferred orientation and neuron location in the map (Pearson, $r$cl = 0.89; p = 0.001), confirming that neurons with similar computational properties are organized in near-by regions of the NS map (see Methods for more details). Applying the k-means algorithm to the NS map revealed an optimal number of $\hat{k}$ = 3 neuron clusters ($\hat{h}$= 0.74, max average silhouette value), indicating that the neurons may be organized into three separate sub-nets (Fig. 5*e, right*). The statistical significance of the number of estimated optimal clusters was determined using the shuffle-based statistical test. The shuffle-test involves generating a null-distribution of silhouette values by shuffling each of the N SSIM matrices, calculating a new NS matrix, and the associated silhouette value. This procedure is repeated over multiple iterations until a distribution of silhouette values is generated. The estimated number of neuron clusters is considered statistically significant if the original silhouette value falls outside the 99% confidence interval of the null-distribution of silhouette values (see Methods and Supplementary Fig. 4 for more details). We examined the computational properties of each of the detected clusters by calculating ensemble tuning functions that take into account the average activity of all neurons within each identified cluster. Our analysis revealed that sub-ensembles displayed significant tuning with peaks evenly distributed at Δ60° intervals. (Fig.

5*h*). Tuning strength and direction-of-motion tuning preferences did not appear to contribute to the cluster organization (data not shown).

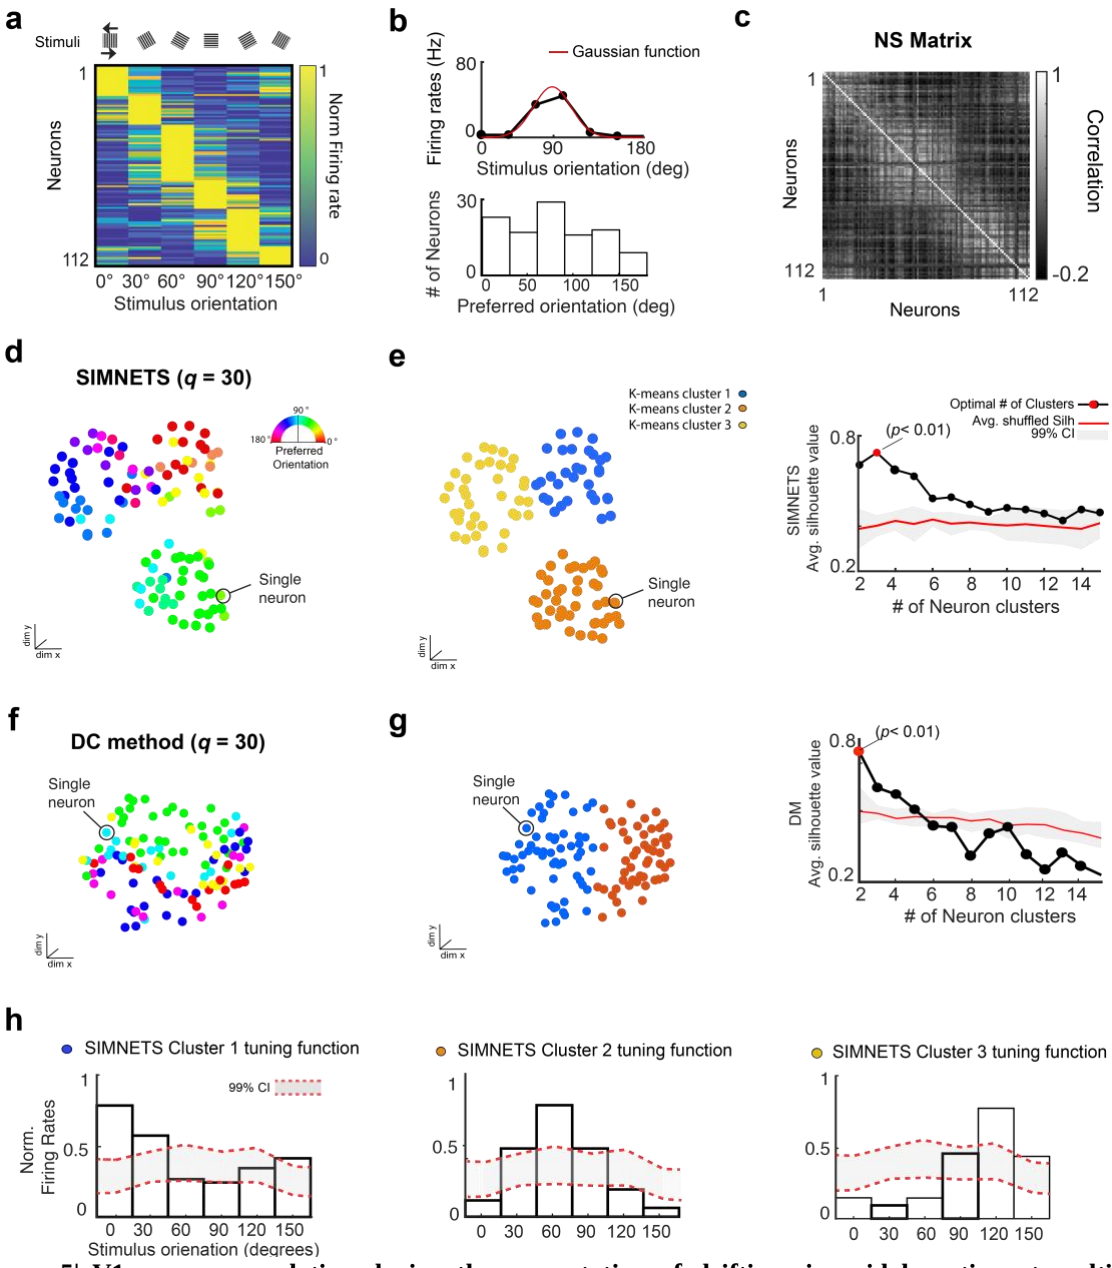

Figure 5| V1 neuron population during the presentation of drifting sinusoidal grating at multiple orientation

a. Normalized trial-averaged firing rates of a population of V1 neurons (N = 112, neurons) during the presentation of 12 different drifting grating stimuli (T = 360, trials). Stimuli were presented for 1.28 s at 6 different orientations (0, 60, 90, 120, 150 degrees) and 2 drift directions (rightward and leftward drift, orthogonal to orientation). **b**. Distribution of calculated preferred grating orientations for all neurons. **c**. SIMNETS NS matrix with neurons ordered according to preferred grating orientation. Color-bar represent a range of correlation values where higher correlation values (light pixels) correspond to greater similarities between neurons. **d.** SIMNETS NS map with neurons labeled according to preferred orientations. **e.** SIMNETS NS map with neurons labeled according to k-means cluster assignments. *Right* (g): Average silhouette value for SIMNETS NS map as a function of the number of clusters (red filled circle denotes optimal cluster separation), compared to expected chance distribution obtained from shuffled data. (bootstrap, p < 0.01). **f-g.** DC NS map with neurons labeled according to preferred directions (**f**) and k-means cluster assignments (**g**). Right (g): Average silhouette value for DC NS map as a function of the number of clusters (red filled circle denotes optimal cluster separation), compared to expected chance distribution obtained from shuffled data. **h.** Ensemble orientation tuning functions, computing by taking the normalized mean firing rates, for each of the SIMNETS neuron clusters (**e**) with significant peaks around 0, 60, 120 degrees

for clusters 1,2, and 3, respectively. Red line and gray band correspond to the mean and 99% confidence intervals of the null-distribution calculated using the shuffle-based statistical test (see Main text).

The SIMNETS algorithm results were compared against the DC method to demonstrate how a more traditional approach fails to organize the neurons according to their functional properties (Fig. 5*f-g*). Although neuron clusters were detected using DC, we observed a weak and non-significant r_{cl} correlation between neuron location and preferred orientation (Pearson, $r_{cl} = 0.01$; $p = 0.56$), indicating that the two detected DC clusters ($\hat{k} = 2$; $\hat{h} = 0.82$) were unlikely to exhibit a tuning preference for any particular orientation. Again, we also compared SIMNETS' performance to the FRC method (Supplementary Fig. 7*a*) and modified cross-correlation (CCH) analysis (Supplementary Fig. 8*a-b*). The performance of the FRC method was comparable to SIMNETS in its ability to organize the neurons according to their estimated computational properties (FRC, r*cl* = 0.86; SIMNETS, *r*cl = 0.89), whereas the cross-correlation based method failed to capture the neuron's estimated functional properties (CCH, r*cl* = 0.16).

M1 Neuron Population – clustering real neurons with known tuning functions

We next applied the SIMNETS algorithm to a dataset of 103 M1 neurons recorded using a 96-channel electrode array in a macaque performing a planar 8-direction instructed-delay reaching task (see Methods). Each neuron's preferred reach direction was estimated by fitting a von Mises distribution[47] to the firing rates as a function of direction (Fig. 6*b*). This dataset and task has previously been described [48,49] (see Methods for more details).

We extracted 1-second spike train events ($S = 114$) from each neuron during all trials where the monkey successfully reached the cued target, starting 0.1 seconds before movement onset. As with the V1 data, the layout of the neurons in the SIMNETS NS map accurately reflected the estimated tuning properties (Fig. 6*d* - *e*). A circular-linear correlation analysis found a significant positive relationship between preferred direction and mapped location (Pearson, *rcl* = 0.92; $p = 0.001$). SIMNETS revealed a statistically significant optimal number of $\hat{k} = 3$ clusters ($\hat{h} = 0.71$), indicative of three functional sub-ensembles. Each cluster displayed ensemble-level tuning with significant peaks at 45°, 180°, and 315°. Additionally, our results show that neurons are not distributed along a uniform continuum within the NS map, but instead form statistically separable clusters in space. These results are in agreement with previous findings, supporting the hypothesis that the biomechanical constraints of the limb are reflected in an uneven distribution of preferred directions among motor cortical neurons [51].

As before, the DC method failed to organize the neurons according to their preferred directions (Fig. 6*f-g*), resulting in weak, non-significant relationship between the neurons' preferred direction and location in the NS map (Pearson, *rcl* = 0.18; $p = 0.56$). Again, we also compared SIMNETS performance to the FRC method and the modified cross-correlation analysis (Supplementary Fig. 7*b*, 8*c*-d). We found that FRC method performed similarly to SIMNETS in organizing the neurons according to their estimated computational properties (FRC, r*cl* = 0.93; SIMNETS, *r*cl = 0.92) and the CCH method failed to organize the majority of the neurons according to their estimated functional properties (CCH, r*cl* = 0.40; Supplementary Fig. 8*c*-d).

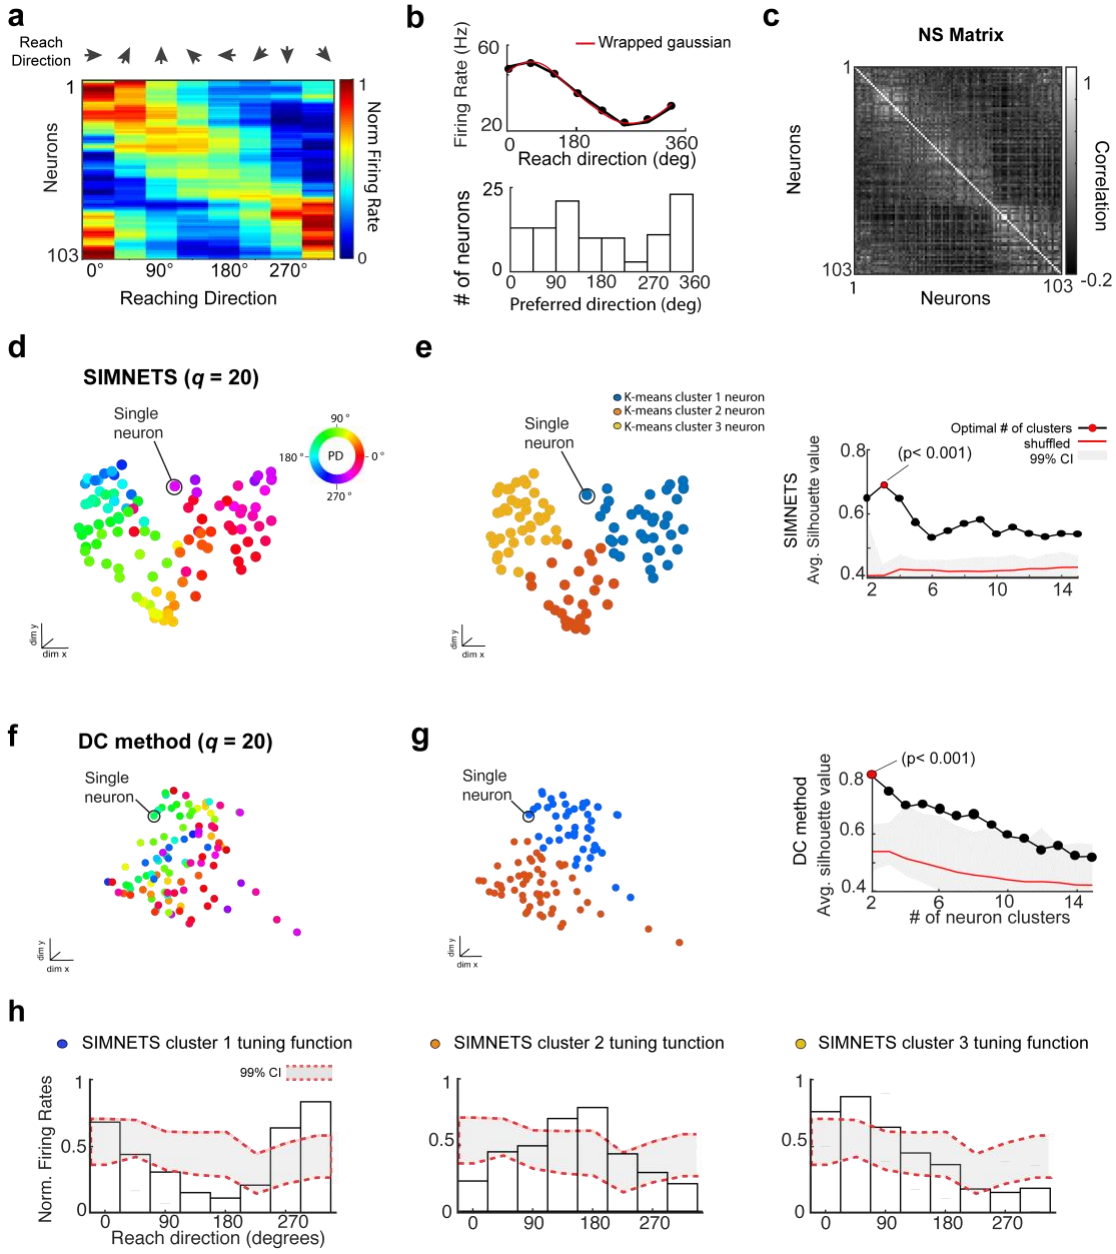

Figure 6| M1 neuron population during center-out reaching task

a. Trial-averaged, normalized firing rates for each neuron ($N = 103$) as a function of reach direction for a planar 8-directional reaching task. **b**. Histogram of estimated preferred reach direction for the population. **c**. SIMNETS NS matrix. Color-bar represent a range of correlation values, where higher correlation values (light pixels) correspond to greater similarities between neurons. **d - e**. SIMNETS NS map with neurons labeled according to the estimated preferred reach directions (**d**) or k-means cluster assignments (**e**). *Right* (**e**): Average silhouette value for SIMNETS NS map as a function of the number of clusters (red filled circle denotes optimal cluster separation), compared to expected chance distribution obtained from shuffled data (bootstrap, $p < 0.01$). **f - g**. DC NS map with neurons labeled according to preferred reach directions (**f**) or k-means cluster assignments (**g**). Right (**g**): Average silhouette value for DC NS map as a function of the number of (bootstrap, $p < 0.01$). **h.** Ensemble tuning functions for each of the SIMNETS neuron clusters (**e**) with significant peaks around 30, 180, 315 degrees for clusters 1, 2, and 3, respectively. Red line and gray band correspond to the mean and 99% confidence intervals of the bootstrapped null-distribution (see Methods).

Hippocampal Dataset – clustering neurons with complex or unknown tuning properties

We applied SIMNETS to a publicly available[42] dataset of N = 80 rat CA1 hippocampal neurons recorded using Multi-site silicon probes while the rat performed left/right-alternation navigation task in a 'figure-8' maze[52,53] (Fig. 7*a*). The rat performed 17 correct trials (T = 17, trials) taking on average 4.3 seconds to reach the reward location at either end of the arms. The input to the SIMNETS algorithm was obtained by dividing the linearized trajectories [42] of the rat's path along the track into six equal segments and extracting 0.75 s spike train events, starting from the time that the rat entered each segment (*see* Methods for more details). This resulted in S = 102 spike train events from each of the N = 80 neurons. In order to validate SIMNETS performance against the other tested algorithms (i.e., DC, FRC, and CCH), we characterized the CA1 neurons as having place cell-like activity (n = 60 PCs) or lacking place cell-like activity (n = 20) based on their spatial firing properties (Fig. 7*b* and Supplementary Fig. 6*c* for example neurons) and took note of the number and location of the significant firing fields (see Methods for more details on neuron characterization and exclusion criteria).

Out of the four algorithms examined, SIMNETS was the only method capable of separating the non-PC and PC groups into statistically significant sub-groupings in the recorded population (Fig. 7*c*-*g*). When we evaluated the distributions of distances *between* the non-PCs and PCs groups and *within* the non-PCs for all four NS maps, we observed a significant difference between the two distance distributions for the SIMNETS (Fig. 7*c*) and DC maps (Fig. 7*e*) *(ranksum*, $p<0.001$) but not the FRC (Fig. 7*f*) and CCH maps (Fig. 7*g*), indicating that SIMNETS and DC method organized the two neuron groups into non-overlapping regions of the NS maps (see Supplementary Table 1 for summaries of distance distributions for all four methods). Interestingly, we observed a strong positive relationship between neuron location and average firing rate in the DC neuron map (Spearman's correlation, r_s = 0.97, $p<0.001$) but only a weak non-significant correlation for that in SIMNETS map (Spearman's correlation, r_s = 0.03, p = 0.74), indicating that although both the DC and SIMNETS method organized the non-PCs into a specific sub-region of their respective maps, the DM method appeared to organize the neurons along a continuum of firing rates statistics. As further evidence to this point, a silhouette analysis and shuffle-based significance test confirmed that the DC NS map did not contain any statistically significant neuron clusters (*shuffle-based statistical test*, $\hat{h}$ = 0.55; see Methods for details). In contrast, a similar silhouette analysis and shuffle-based statistical test determined that SIMNETS organized the neurons into $\hat{k}$ = 6 statistically separate clusters (shuffle-significance test, $p<0.01$). Supplementary Table 1 provides a summary for the distance, silhouette, and clustering analysis results for all four methods.

An examination of the physiological properties of the SIMNETS clusters revealed that one of the putative sub-ensembles was composed almost entirely (95%) of non-PCs, while the other five clusters were either entirely or almost entirely (>92%) composed of PCs. An inspection of the single neuron and SIMNETS cluster (i.e., ensemble) firing rate maps indicates that each of the 'PC clusters' were comprised of neurons with overlapping or partially overlapping place-fields (Fig. 7*h;* Supplementary Fig. 6*b*). PCs with single and multiple peaks in their spatial firing maps were found within the same clusters if they shared a common peak firing field (for example, see Supplementary Fig. 6*c*, *clusters 2 and 4*). Interestingly, despite being primarily made up of neurons that lack place-dependent signals,

the cluster-6 spatial firing map exhibited a single significant firing field (Fig. 7*i*, *cluster 6;* see Method for more details). This phenomenon could reflect the contribution of the single PC in the cluster or reflect an emergent feature of the neuron cluster, e.g., ensemble coding [54]. In order to get a better understanding of the computational properties of this cluster in relation to the other detected clusters, we examined the ensemble activity patterns of each detected cluster using a spike train relational analysis framework [33] (*see* Methods or Vargas-Irwin et al (2015) for details).

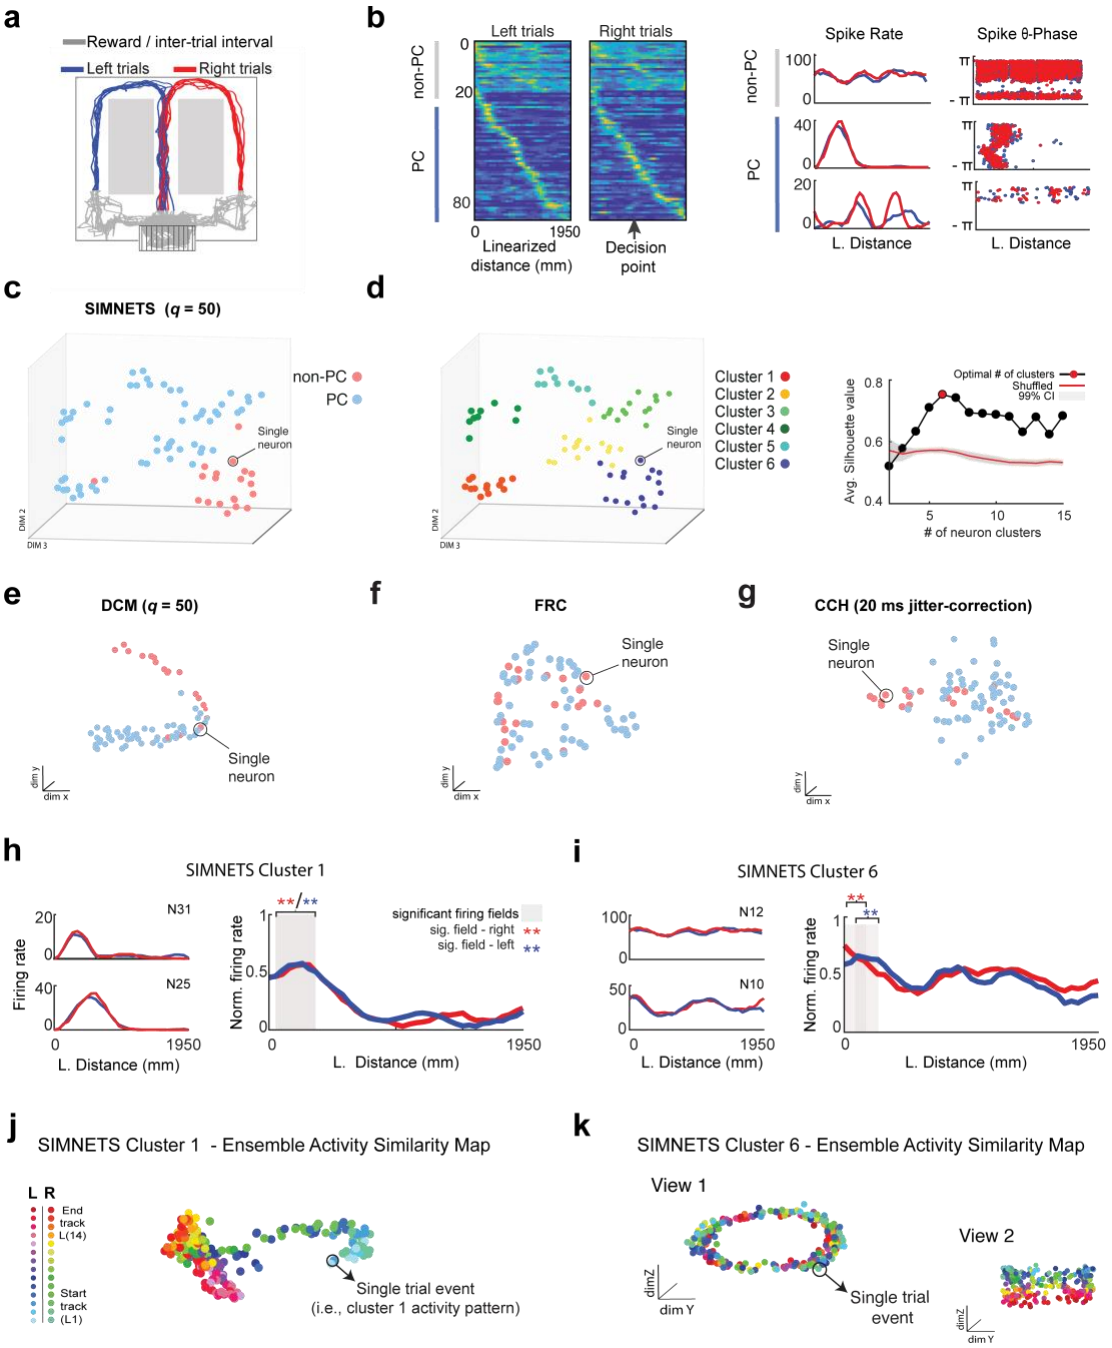

Figure 7|Hippocampal neurons during left-right alternation task.

a. Maze showing rats position in a 'figure-8' maze during left-right alternation task. Red and blue lines show the rat's location during correct right and left trials (T = 18, trials), respectively (similar color convention used throughout b, e-f). Gray lines show the rats location during reward and inter-trial interval period. **b.** *Left*: Normalized firing rates are shown for each neuron (N = 80, neurons) as a function of linearized distance on track (50 mm bins, 150mm Gaussian smoothing)

during left and right trials. Neurons were ordered according to the latency of their peak response along the track and according to their characterization as a non-place cell (n = 20, non-PC) or a place cell (n = 60, PC). Arrow indicates the left/right decision point in maze. *Right*: Spike rate maps and spike theta-phase spatial maps for one example non-PC (*top row*) and two example PCs (bottom *rows*). **c-d.** SIMNETS NS map with neurons labeled as non-PCs (gray points) and PCs labels (blue points) (**c**) and k-means clustering labels (**d,** *left*) and Silhouette plot (**d**, *right*). An optimal number of k = 6 neuron clusters were detected in the NS map. **e.** Ensemble spatial firing rate maps shown for three example SIMNETS clusters (top row), and the spatial firing map for two representative neurons from each example cluster (*middle, bottoms rows*). Gray bands indicate the significant firing fields (rank-sum, $p < 0.01$, in-field vs. out-of-field comparison) for the left and right trials (blue and red asterisks)**. f**. *Top row*: Ensemble Activity Relational map showing the similarity between ensemble activity patterns as a function of the rat's location on the linearized track during left and right trials (L1 - L14). Each point represents an ensemble activity pattern on a single trial and the distance between points represents the similarity of the ensemble spiking patterns. Different colors correspond to the different locations on the left and right track. *Bottom row*: distance plots showing the mean distance between points corresponding to each event type and all other points (e.g., distance between L1 points and all other points). Gray bands represent the 99% confidence interval of a null-distribution (bootstrap analysis).

An ensemble spike train similarity analysis generates an Ensemble Activity Similarity map similar to those presented in Fig. 1*c* (and Supplementary Figure 1-2) but encompasses the activity of multiple neurons rather than just a single neuron. Each point in the ensemble Activity Similarity maps (Fig. 7*j*,*k*) corresponds to the activity pattern generated across all neurons in a particular SIMNETS neuron cluster on *a single trial*. As expected from the previous analysis, the topology of the Ensemble Similarity map for the 'PC neuron clusters' (Fig. 7*j*) captured the place-dependent modulation of the sub-net firing patterns as the rat traversed the a specific region of the maze. Interestingly, the Ensemble Activity Similarity map for the neuron cluster composed primarily of non-PCs (Fig. 7*k*), had a 'torus-shaped' topology that exhibited variance along the z-dimension in relation to rat's position along the track (Fig. 7*k*, *view 2*), and a periodic variance in the *x*- and y-dimension (Fig. 7*k*, view 1) according to an unknown variable, or variables. This suggests that the activity of the non-PC sub-net displays dynamics that may reflect a non-spatial task-variable or, potentially, the intrinsic dynamics of the circuit. The toroid structure of the Ensemble Activity Similarity map suggests that the unknown variable is likely to be periodic in nature. Although this phenomenon warrants further investigation, it is outside of the scope of the present work. These particular results highlight the advantages of applying SIMNETS to neural recordings where the tuning properties of the neurons are not readily apparent or known *a priori*.

Computational efficiency and analysis run-time

The SIMNETS algorithm processed each of the synthetic, M1, and Hippocampal datasets in fewer than 5 seconds using a standard desktop computer (Intel Xeon® Processor, 24 GB of RAM ; see Methods for more details). Because of a larger trial number, the V1 dataset was processed in a relatively slower time of around 20 seconds. In general, SIMNETS' run-time for a dataset of 100 neurons, with 100 spike trains per neuron, takes approximately ~4 s (see Supplementary Fig. 9*a*). Importantly, the computational complexity of the algorithm scales near-linearly with neuron number and quadratically with the number of spike trains, meaning that datasets of up to 1000 neurons can be analyzed in a reasonable amount of time on consumer-grade hardware (~ 4 minutes). By comparison, calculating the pairwise cross-correlation [45,46] for 100 and 1000 neurons would take approximately 6 minutes and 6 hours, respectively, using the same hardware (supplementary Fig. 9*b*).

Discussion

Summary of Findings

Advances in multi-electrode recording technology have now made it possible to record or image the activity of thousands upon thousands of individual neurons simultaneously [7,55–57]. By contrast, the development of analytical tools capable of parsing the complexity of large-scale neural activity patterns have lagged behind advances in recording technology. Here, we introduce a computationally efficient and scalable method for characterizing the information processing properties of individual, simultaneously recorded neurons, allowing us to extract and visualize sub-networks of functionally similar neurons within large neuronal ensembles.

The critical component that differentiates the SIMNETS framework from the previous methods is our novel application of spike train metrics for capturing the relationship between the output states of each individual neurons in a universal format that is easily compared across neurons. We emphasize that, unlike other related algorithms, SIMNETS does not directly compare spiking responses between neurons; instead, our approach compares intrinsic geometry of the output spaces of each neuron (represented by SSIM matrices), which reflects information-processing properties in a universal format (i.e. regardless of the format of the encoded information).

Our analysis of simulated data with known ground truth demonstrates that SIMNETS can organize the neurons into functionally related sub-nets, even when computationally equivalent neurons utilize very different encoding schemes (e.g., rate, temporal, or mixed encoding schemes). Our analysis of V1 and M1 recordings shows that SIMNETS neuron functional maps could recapitulate the computational properties of the neurons without the need for imposing stimulus or movement driven parametric tuning models *a priori.* Our results with the hippocampal CA1 data suggest that it will be possible to use this approach to identify and separate neurons with complex or unknown tuning functions into distinct groups according to the similarities of their information processing properties. Our results also suggest that SIMNETS may be able to detect functional sub-ensembles hypothesized to support ensemble place-coding[58] or complex feature conjunction[59]. Although it was beyond the scope of this report to demonstrate functional significance of the detected putative functional sub-ensembles, our results strongly suggest that sub-nets detected using SIMNETS are statistically and physiologically meaningful. We also demonstrated that SIMNETS can identify groups of neurons with similar functional properties that other comparable methods cannot detect using standard statistical approaches. Finally, our particular choice of datasets allowed us to demonstrate that this method generalizes well to neural recordings from a variety of brain regions (including sensory, motor, and hippocampal areas) and across multiple species (including rat and non-human primate).

Collectively, our results demonstrate that estimates of correlation between SSIM matrices provide a simple, yet powerful approach for quantifying the functional similarities between neurons. Our analysis strategy shifts the emphasis from detecting coincident or correlated activation to comparing the intrinsic relational structure of single-trial firing patterns. This

critical difference allows our method to detect neurons with similar computational properties even if the neurons exhibit diverse firing patterns. The combination of a short processing time (< 5 second per 100 neurons) and a computational complexity that scales near-linearly with the size of the neuron population makes SIMNETS an extremely efficient and, thus, appealing tool for exploring very large-scale neuron populations.

Comparison to existing methods

The concept of a low dimensional embedding that captures the functional relationship between neurons was introduced in the seminal papers by Gerstein & Aertsen[45,60], where they describe 'Gravitational Clustering' (GC): a neuron clustering and visualization tool for identifying groups of neurons with synchronous spiking patterns. GC is based on an analogy of the physics of the gravitational forces governing the dynamics and interactions of macroscopic particles. It treats the *N* neurons as *N* particles moving within an *N*-dimensional space, where charges that influence the attractive and repulsive interactions between particles are dictated by the temporal dynamics of pairwise synchronous spiking activity between neurons. The end result is a visualization of particle clusters (and their trajectories) that represent dynamically evolving assemblies of synchronously-active neurons. Recent formulations of the GC algorithm have improved visualization and sensitivity, but retain the same basic strategy[11,39]. More recent work by Kiani et al. (2015) involves the application of the dimensionality reduction techniques (i.e., MDS) to pairwise measures of between-neuron spike rate covariations in order to detect 'natural' functional modules in populations of neurons. Although, the goal of this method is similar to that of SIMNETS – to the extent that it makes use of single trial information to group neurons in an unsupervised manner – this approach only takes into account the covariation information carried in the absolute firing rates of the neurons' responses across trials, rather than the information carried in the intrinsic geometry of the neurons' spike train patterns across trials. For neurons that operate using predominantly rate-based codes, this approach generates results roughly equivalent to SIMNETS, as shown in our analysis of rate based synthetic data (Fig. 3) as well as MI and V1 ensembles (supplementary Fig. 7b,c). However, our analysis of the synthetic neurons using temporal codes (Fig. 3, supplementary Fig 7a), as well rat CA1 neurons (supplementary Fig 7d) demonstrates that SIMNETS can detect a broader range of information processing motifs when task information is carried in the temporal dynamics of spike train outputs [26].

Several previous studies have used inter-neuron spike train metrics to identify putative functional sub-nets[61,62]. As with the GC method, these studies have operated under the general assumption that the detection of similar sequences of spike patterns across neurons is indicative of a potential functional link[62,63]. By contrast, our approach is capable of detecting neurons with similar informational content even if they employ different encoding schemes (i.e. rate vs. precise spike timing). SIMNETS clustered the simulated neurons according to their ground-truth functional ensembles and, when we adjusted the tunable temporal accuracy parameter, was capable of further sub-dividing the neurons according to their distinctive firing characteristic, such as their encoding timescale. This feature of SIMNETS could be particularly useful for determining if neurons operating on different timescales – such as, different interneuron subtypes – are involved in different information processing operations[64,65]. Additionally, this feature of SIMNETS enabled us to cluster known functional groups in the CA1 dataset (i.e. PCs vs. non-PCs) that other algorithms missed.

Capturing interactions between pairs of neurons can be accomplished in a more general way using information theory, e.g., mutual information[66], or by leveraging asymmetries in the predictive power of variables at different lags to generate 'directed' estimates of functional connectivity, e.g., Granger Causality[67], Transfer Entropy[68], or the Directed Transfer Function[69,70]. These strategies require estimates of joint probability distributions across the activity patterns of pairs of neurons. The number of possible activity patterns is very large, so this type of calculation can be challenging, even when relatively large amounts of data are available. Unlike SIMNETS, the computational cost of this approach grows quadratically with the number of neurons examined, making it impractical to process datasets with hundreds or thousands of neurons without specialized hardware.

Limitations of SIMNETS

Several important limitations of SIMNETS are worth noting. First, estimates of similarity using spike train metrics require that the time windows of interest be of equal length, making it difficult to compare neural responses with different time courses. This particular weakness is common to all trial-averaging models commonly used in the literature that we are aware of. Second, although the SIMNETS framework does not require *a priori* assumptions about the variables potentially encoded by neural activity, experimental design and data selection will still have a direct effect on the results obtained. For example, a set of neurons identified as a functional sub-net could separate into smaller groups with different computational properties when additional task conditions are added to the analysis. Thus, the functional properties identified using SIMNETS are only valid within the context of the data examined and may not necessarily extrapolate to different experimental conditions. Third, it is likely that neuronal sub-nets are constantly re-arranged depending on ethological demands. A neuron could potentially be functionally interacting with one group of neurons for one computation (or moment in time) and then another group of neurons for another computation (or another moment in time). The current version of the SIMNETS algorithm was not designed to distinguish between such rapidly changing sub-network memberships. However, it is possible to apply the SIMNETS algorithm multiple times over different epochs in order to determine if different sub-nets are present across different conditions. Our future work will focus on examining the temporal evolution of sub-net clustering using this approach. Fourth, SIMNETS may fail to recover certain functional characteristic similarities among neurons. For example, the neuronal mechanism underlying receptive field overlap (i.e., common synaptic input) is thought to be a major source of synchrony in the cortex[71]. As such, receptive field overlap may be best recovered with the use of cross-correlation based methods. However, using SIMNETS to narrow down the sets of neurons evaluated for synchrony could potentially reduce the considerable computational load incurred by correlation-based analyses. Fifth, the computations required to generate a SSIM matrix for a neuron scales quadratically with the number of trials. Despite this computational cost, SIMNETS scales better than methods that require computing power-sets and would thereby grow exponentially with the number of neurons.

The role for SIMNETS to mitigate experimenter bias

A considerable amount of research in systems neuroscience has focused on identifying new classes of neurons based on their information processing properties. The standard approach

for many of these experiments involves recording single unit activity while a certain experimental variable of interest is manipulated (for example, providing different stimuli, or eliciting different movements, etc.). Standard statistical tests (ANOVA, etc.) are then used to determine if each neuron displays significant changes in firing rate across the experimental conditions. The percentage of significant neurons is usually reported as a functionally distinct 'class' of neurons sensitive to the variable of interest. It is common to exclude neurons that do not reach statistical significance or cannot be fit using a predetermined model from further analysis. This approach is prone to both selection and confirmation bias, and ultimately produces 'classes' of neurons identified based on arbitrary statistical thresholds imposed on what are likely continuous distributions of properties[72,73]. SIMNETS provides a way to determine if neurons are functionally organized across a continuum or are organized into statistically separable neural sub-nets, there-by mitigating the experimenter's bias inherent in parametric neural discovery methods. In addition to providing a principled way to determine if a consistent organization of information processing modules can be found across sessions and subjects, we believe that the ability to intuitively visualize relationships within networks of neurons will provide a unique perspective leading to new data-driven hypotheses and experimental refinement.

Acknowledgments:
This work was supported by NINDS-Javits (NS25074), The Israeli Brain Prize, NIH Director's New Innovator award, and the Killam Trust Award Foundation. We thank Stuart Geman and Adam Kohn for their feedback on the method, as well as M. Nevor and J Murphy for their assistance with animal care and instrumentation design. We also thank the Buzsaki and Kohn Lab for allowing us (and the field) to work with their data via the CRCNS data repository.

Bibliography

1. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13,** 51–62 (2011).
2. Zagha, E., Ge, X. & McCormick, D. A. Competing Neural Ensembles in Motor Cortex Gate Goal-Directed Motor Output. *Neuron* **88,** 565–577 (2015).
3. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321,** 1322–7 (2008).
4. Harris, K. D. Opinion: Neural signatures of cell assembly organization. *Nat. Rev. Neurosci.* (2005). doi:10.1038/nrn1669
5. Brown, E. N., Kass, R. E. & Mitra, P. P. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat. Neurosci.* (2004). doi:10.1038/nn1228
6. Briggman, K. L., Abarbanel, H. DI & Kristan, W. B. From crawling to cognition: analyzing the dynamical interactions among populations of neurons. *Current Opinion in Neurobiology* (2006). doi:10.1016/j.conb.2006.03.014
7. Buzsáki, G. Large-scale recording of neuronal ensembles. *Nat. Neurosci.* **7,** 446–451 (2004).
8. Wallace, D. J. & Kerr, J. N. D. Chasing the cell assembly. *Current Opinion in Neurobiology* (2010). doi:10.1016/j.conb.2010.05.003
9. Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K. & Palm, G. Dynamics of Neuronal Firing Correlation: Modulation of " Effective Connectivity ". *JOURNALOFNEUROPHYSIOLOGY Print. i n U.S.A* **61,** (1989).
10. Abeles, M. & Gat, I. Detecting precise firing sequences in experimental data. *J. Neurosci. Methods* (2001). doi:10.1016/S0165-0270(01)00364-8
11. Baker, S. N. & Gerstein, G. L. Improvements to the Sensitivity of Gravitational Clustering for Multiple Neuron Recordings. *Neural Comput.* **12,** 2597–2620 (2000).
12. Grün, S., Diesmann, M. & Aertsen, A. Unitary Events in Multiple Single-Neuron Spiking Activity: II. Nonstationary Data. *Neural Comput.* (2002). doi:10.1162/089976602753284464
13. Rgy, G., Iki, B. & Chrobak, J. I. Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr. Opin. Neurobiol.* **5,** 504–510 (1995).
14. Humphries, M. D. Spike-Train Communities: Finding Groups of Similar Spike Trains. *J. Neurosci.* (2011). doi:10.1523/JNEUROSCI.2853-10.2011
15. Singer, W. & Gray, C. M. Visual Feature Integration and the Temporal Correlation Hypothesis. *Annu. Rev. Neurosci.* **18,** 555–586 (1995).
16. Singer, W. Neuronal Synchrony: A Versatile Code for the Definition of Relations? *Neuron* **24,** 49–65 (1999).
17. Von Der Malsburg, C. *The Correlation Theory of Brain Function*.
18. Brody, C. D. Correlations Without Synchrony. *Neural Comput.* **11,** 1537–1551 (1999).
19. Shadlen, M. N. & Movshon, J. A. Synchrony Unbound. *Neuron* **24,** 67–77 (1999).
20. Abeles, M. & Gerstein, G. L. Detecting Spatiotemporal Firing Patterns Among Simultaneously Recorded Single Neurons. *JOURNALOFNEUROPHYSIOLOGY Print. i n U.S.A* **60,** (1988).

21. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17,** 1500–9 (2014).
22. Machens, C. K., Romo, R. & Brody, C. D. Functional, But Not Anatomical, Separation of 'What' and 'When' in Prefrontal Cortex. *J. Neurosci.* **30,** 350–360 (2010).
23. Brody, C. D. Correlations Without Synchrony. *Neural Comput.* **11,** 1537–1551 (1999).
24. Truccolo, W., Hochberg, L. R. & Donoghue, J. P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13,** 105–111 (2010).
25. Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat. Commun.* **6,** 7759 (2015).
26. Churchland, M. M. & Shenoy, K. V. Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *J. Neurophysiol.* **97,** 4235–4257 (2007).
27. Chelaru, M. I. & Dragoi, V. Efficient coding in heterogeneous neuronal populations. *Proc. Natl. Acad. Sci. U. S. A.* **105,** 16344–9 (2008).
28. Lehky, S. R., Sereno, M. E. & Sereno, A. B. Population Coding and the Labeling Problem: Extrinsic Versus Intrinsic Representations. *Neural Comput.* **25,** 2235–2264 (2013).
29. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
30. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
31. Victor, J. D. & Purpura, K. P. Metric-space analysis of spike trains: theory, algorithms and application. *Netw. Comput. Neural Syst.* (1997). doi:10.1088/0954-898X_8_2_003
32. Zaidi, Q. *et al.* Perceptual spaces: mathematical structures to neural mechanisms. *J. Neurosci.* **33,** 17597–602 (2013).
33. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
34. Lehky, S. R., Sereno, M. E. & Sereno, A. B. Population Coding and the Labeling Problem: Extrinsic Versus Intrinsic Representations. *Neural Comput.* (2013). doi:10.1162/NECO_a_00486
35. Victor, J. D. & Purpura, K. P. Nature and Precision of Temporal Coding in Visual Cortex: A Metric-Space Analysis. *J. Neurophysiol. Pt-intd U.S.A* **76,** (1996).
36. Kriegeskorte, N. & Kievit, R. A. Representational geometry: integrating cognition, computation, and the brain. *Trends Cogn. Sci.* **17,** 401–12 (2013).
37. Van Der Maaten, L. & Hinton, G. *Visualizing Data using t-SNE. Journal of Machine Learning Research* **9,** (2008).
38. Grün, S., Diesmann, M. & Aertsen, A. Unitary Events in Multiple Single-Neuron Spiking Activity: II. Nonstationary Data. *Neural Comput.* **14,** 81–119 (2002).
39. Stuart, L., Walter, M. & Borisyuk, R. Visualisation of synchronous firing in multi-dimensional spike trains. in *BioSystems* (2002). doi:10.1016/S0303-2647(02)00084-9
40. Kohn, A., Smith, M. A. (2016). Utah array extracellular recordings of spontaneous

and visually evoked activity from anesthetized macaque primary visual cortex (V1). *CRCNS.org* (2016). Available at: http://dx.doi.org/10.6080/K0NC5Z4X.
41. Rao, N. G. & Donoghue, J. P. Cue to action processing in motor cortex populations. *J. Neurophysiol.* **111,** 441–53 (2014).
42. Pastalkova E, Wang Y, Mizuseki K, B. G. Simultaneous extracellular recordings from left and right hippocampal areas CA1 and right entorhinal cortex from a rat performing a left / right alternation task and other behaviors. (2015). Available at: http://dx.doi.org/10.6080/K0KS6PHF.
43. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220,** 149–166 (2013).
44. Kiani, R. *et al.* Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. *Neuron* **85,** 1359–73 (2015).
45. Gerstein, G. L. & Aertsen, A. M. Representation of cooperative firing activity among simultaneously recorded neurons. *J. Neurophysiol.* **54,** (1985).
46. Smith, M. A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* **28,** 12591–603 (2008).
47. Amirikian, B., Georgopoulos, A. P. & Georgopulos, A. P. Directional tuning profiles of motor cortical cells. *Neurosci. Res.* **36,** 73–9 (2000).
48. Rao, N. G. & Donoghue, J. P. Cue to action processing in motor cortex populations. *J. Neurophysiol.* **111,** 441–453 (2014).
49. Vargas-Irwin, C. E., Franquemont, L., Black, M. J. & Donoghue, J. P. Linking Objects to Actions: Encoding of Target Object and Grasping Strategy in Primate Ventral Premotor Cortex. *J. Neurosci.* **35,** 10888–97 (2015).
50. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. doi:10.1162/NECO_a_00684
51. Lillicrap, T. P. & Scott, S. H. Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics. *Neuron* **77,** 168–179 (2013).
52. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321,** 1322–7 (2008).
53. Mizuseki, K., Sirota, A., Pastalkova, E. & Buzsáki, G. Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* **64,** 267–280 (2009).
54. Park, E., Dvorak, D. & Fenton, A. A. Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS One* **6,** e22349 (2011).
55. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551,** 232–236 (2017).
56. Csicsvari, J. *et al.* Massively Parallel Recording of Unit and Local Field Potentials With Silicon-Based Electrodes. (2003). doi:10.1152/jn.00116.2003
57. Kim, D. H. *et al.* Pan-neuronal calcium imaging with cellular resolution in freely swimming zebrafish. *Nat. Methods* **14,** 1107–1114 (2017).
58. Park, E., Dvorak, D. & Fenton, A. A. Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS One* **6,** e22349 (2011).
59. Geiller, T., Fattahi, M., Choi, J.-S. & Royer, S. Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nat. Commun.* **8,** 14531 (2017).
60. Gerstein, G., Perkel, D. & Dayhoff, J. Cooperative firing activity in simultaneously

recorded populations of neurons: detection and measurement. *J. Neurosci.* **5,** (1985).
61. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220,** 149–166 (2013).
62. Humphries, M. D. Spike-Train Communities: Finding Groups of Similar Spike Trains. *J. Neurosci.* **31,** 2321–2336 (2011).
63. Fujisawa, S., Amarasingham, A., Harrison, M. T. & Buzsáki, G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11,** 823–833 (2008).
64. Wilson, N. R., Runyan, C. A., Wang, F. L. & Sur, M. Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature* **488,** 343–8 (2012).
65. Neske, G. T., Patrick, S. L. & Connors, B. W. Contributions of diverse excitatory and inhibitory neurons to recurrent network activity in cerebral cortex. *J. Neurosci.* **35,** 1089–105 (2015).
66. Paninski, L. Estimation of Entropy and Mutual Information. *Neural Comput.* (2003). doi:10.1162/089976603321780272
67. Granger, C. W. J. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica* **37,** 424–438 (1969).
68. Vicente, R. *et al.* Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *J Comput Neurosci* **30,** 45–67 (2011).
69. Kaminski, M. J. & Blinowska, K. J. A new method of the description of the information flow in the brain structures. *Biol. Cybern.* **65,** 203–210 (1991).
70. Blinowska, K. J. Review of the methods of determination of directed connectivity from multichannel data. *Medical and Biological Engineering and Computing* **49,** 521–529 (2011).
71. Ts'o, D. Y., Gilbert, C. D. & Wiesel, T. N. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* **6,** 1160–70 (1986).
72. Blanchard, T. C., Piantadosi, S. T. & Hayden, B. Y. Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. *J. Neurophysiol.* **119,** 1305–1318 (2018).
73. Chandrasekaran, C., Peixoto, D., Newsome, W. T. & Shenoy, K. V. Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nat. Commun.* **8,** 614 (2017).
74. Lloyd, S. P. *Least Squares Quantization in PCM*. *IEEE TRANSACTIONS ON INFORMATION THEORY* **28,** (1982).
75. Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **20,** 53–65 (1987).
76. Author, S., Kullback, S. & Leibler, R. A. *On Information*. *Source: The Annals of Mathematical Statistics* **22,** (1951).
77. Mantel, N. The detection of disease clustering and a generalized regression approach. *Cancer Res.* **27,** 209–20 (1967).
78. Perkel, D. H., Gerstein, G. L. & Moore, G. P. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys J* **7,** 419–440 (1967).
79. Amarasingham, A., Harrison, M. T., Hatsopoulos, N. G. & Geman, S. Conditional modeling and the jitter method of spike resampling. *J. Neurophysiol.* **107,** 517–31 (2012).
80. Gur, M. & Snodderly, D. M. Direction selectivity in V1 of alert monkeys: evidence for parallel pathways for motion processing. *J. Physiol.* **585,** 383–400 (2007).
81. Hawken, M. J., Parker, A. J. & Lund, J. S. Laminar organization and contrast

sensitivity of direction-selective cells in the striate cortex of the Old World monkey. *J. Neurosci.* **8,** 3541–8 (1988).

82. Skaggs, W. E., Skaggs, W. E., McNaughton, B. L., Gothard, K. M. & Markus, E. J. An Information-Theoretic Approach to Deciphering the Hippocampal Code. *IN* **5,** 1030--1037 (1993).
83. Bjerknes, T. L., Dagslott, N. C., Moser, E. I. & Moser, M.-B. Path integration in place cells of developing rats. *Proc. Natl. Acad. Sci. U. S. A.* **115,** E1637–E1646 (2018).
84. Jeffery, K. J., Donnett, J. G., Burgess, N. & O'Keefe, J. M. Directional control of hippocampal place fields. *Exp. brain Res.* **117,** 131–42 (1997).
85. Ayzenshtat, I., Jackson, J. & Yuste, R. Orientation Tuning Depends on Spatial Frequency in Mouse Visual Cortex. *eNeuro* **3,** (2016).

SIMNETS: a computationally efficient and scalable framework for identifying sub-networks of functionally similar neurons

Methods

Jacqueline B. Hynes [1, 2], **David M. Brandman**,[2, 3], **Jonas B. Zimmerman**[4], **John P. Donoghue**[1, 2, 4, 5], **Carlos Vargas-Irwin** [*, 1, 2]

1 Department of Neuroscience, Brown University, Providence, RI, USA.
2 Robert J. and Nancy D. Carney Institute for Brain Science, Brown University, Providence, RI, USA
3 Department of Surgery (Neurosurgery), Dalhousie University, Halifax, NS, Canada.
4 WYSS Institute, Chemin des Mines 9, CH-1202, Geneve, Switzerland.
5. Department of Engineering, Brown University, Providence, RI, USA.

∗ Correspondence to: carlos_vargas_irwin@brown.edu

Keywords: functional sub-networks, neural ensembles, clustering algorithm, dimensionality reduction, spike train distances

SIMNETS Algorithm Implementation

Here, we provide a short description of the steps in the algorithm and a detailed description of the methods used to implement each step of the SIMNETS framework:

Step 1: Calculate distances between the spike trains generated by each neuron
Pairwise distances are calculated between spike trains (S) generated by the same neuron using a spike train metric. This results in separate SxS Spike train Similarity (SSIM) matrices for each neuron. Here, we use the Victor-Purpura Spike train metric [31,35].

Step 2: Spike train similarity matrix correlation
Pairwise measures of correlation are calculated between all single-neuron SSIM matrices, resulting in a single NxN Neuron Similarity (NS) Matrix. Here, we use Pearson's Correlation.

Step 3: Dimensionality reduction
The high-dimensional, NxN Neuron Similarity Matrix is projected down into a desired number of d dimensions and visualized in a scatter plot, resulting in what we refer to as the Neuron Similarity Map. The dimensionality reduction step is carried out using t-distributed Stochastic Neighbor Embedding (t-SNE)[37].

Step 4: Cluster detection and statistical test
Putative functional ensembles are detected in the $N \times d$ Neuron Similarity Map using the unsupervised k-means clustering algorithm [74]. The number of clusters in the data is determined using a silhouette analyis[75] and the significance of the number of detected clusters is determined using a shuffle-based procedure.

The user selected parameters of the SIMNETS algorithm include: 1) the VP cost function constant, q, which operationally defines the temporal resolution over which the similarity of two spikes trains are tested, 2) the t-SNE parameter, perplexity, which influences the number of *effective* nearest neighbors (i.e., neurons) included in calculations that results in the low-dimensional neuron map, and 3) the dimensionality of the NS map.

A live code tutorial is available at: https://github.com/DonoghueLab/SIMNETS.git

Step 1: Victor-Purpura Spike Train Metric

The Victor-Purpura (VP) metric is a cost-based spike train distance function (D) that describes the similarity between pairs of spike trains (common neuron) in terms of their 'edit-distances'. A single distance value (d) is assigned to each pair of spike trains through a process that involves calculating the minimum total 'cost' (c) of the edit-steps needed to transform spike train A into spike train B:

$$\boldsymbol{d(A,B) = min\{\sum_{j=0}^{m-1} c\,(s_{j,}s_{j+1})\},} \qquad \textbf{(1)}$$

where $\{S_0, S_1, ..., S_m\}$ is the series of intermediate spike trains created after performing a single edit step. The list of possible edit-steps used in the VP transformation include: (1) inserting a spike, (2) deleting a spike, and (3) shifting a spike in time. Inserting or deleting a spike has a cost of $c = 1$, and shifting a single spike in time has a cost proportional to the amount of time that it is moved ($c = q\Delta t$). The set of edits-steps associated with the minimum total edit-cost defines the shortest path between two points (spike trains) in the neuron's spike train metric-space.

The q parameter influences the relative importance of spike count and spike time differences when assessing spike train similarities. When $q = 0$, the cost of shifting a spike to a desired location will always be cheaper than deleting and re-inserting a spike in a spike train. Thus, for $D_{[q=0]}$, the minimum cost is a function of the difference in the number of spikes between the spike trains. As the q value is increased beyond zero, spike time jitter begins to impact the cost of matching the spike trains. For example, if $q = 10$, shifting a spike by 0.15 s will have a cost of $c = 1.5$, which is still just under the cost of deleting and re-inserting spike, i.e., $c = 2$, making it the more cost effective option. However, if $q = 15$, the deleting and reinserting a spike will become the cheaper option. This means that if we were matching two temporally jittered spike trains with a similar number of spikes, the assigned spike train distance would jump from a small to a high value as q increases (due to the increasing cost associated with shifting a spike). On the other hand, if we were matching two spike trains that differed only by spike number, i.e., no temporal jitter, the cost of shifting a spike would not impact the total cost of matching the spike trains, and so we would not expect a jump in the assigned spike train distances with an increase in q value. In this way, q controls the temporal resolution of the spike train comparison. In the context of the SIMNETS algorithm, a low q parameter will bias the algorithm towards groupings neurons based on the 'information' encoded over coarse timescales, whereas a high q parameter will bias the algorithm towards groupings neurons based on the information encoded over coarse and fine timescales.

Step 2: Spike Train Similarity Matrix Correlation

We characterize the functional similarities between neurons by calculating pairwise measures of correlation between all pairs of SSIM matrices. We used Pearson's Correlation (r) to compare the SSIM matrices in this report because of its efficiently and empirical

success, however, other matrix correlation statistics or distances measures may also be used. The formula for calculating Pearson's r between a pair of SSIM matrices, $A = (a_{ij})$ and $B = (b_{ij})$ is given as:

$$r(\varphi) = \beta + h.exp\big(h.cos(\varphi - \mu)\big), \quad \textbf{(2)}$$

where *cov* is the covariance and σ is the standard deviation. This results in an NxN Correlation matrix, where each matrix entry corresponds to the correlation between a given pair of SSIM matrices, and each column (or row) of the matrix could be interpreted as the intrinsic coordinates of a single neuron in an N-dimensional space.

Step 3: t-SNE Dimensionality Reduction algorithm

In broad terms, the goal of the dimensionality reduction step is to reduce the number of variables required to represent each neuron's N dimensional correlation vector (*step 2*), i.e., its coordinates in the high-dimensional neuron space. This step improves clustering performance and allows us to visualize the relationships between neurons in a low-dimensional map. We used the t-distributed Stochastic Neighbor Embedding (t-SNE) dimensionality reduction algorithm, since it is designed to preserve local densities of the high-dimensional data, while revealing global structure such as the presence of clusters at several scales [37]. These properties make the algorithm particularly well suited for visualizing high-dimensional data with varying cluster densities.

A t-SNE transform is calculated through a process that involves 1) converting the sets of high- and low-dimensional correlation/distance measures into sets of joint probability distributions that describe the 'similarity' between the data points in the respective high and low dimensional spaces, and 2) minimizing the Kullback-Leibler divergence[76] between the sets of joint probabilities in the high-dimensional space and the low-dimensional map via gradient descent.

In the first step, the similarity of the data point w_j to w_i in the high dimensional space is modeled as the conditional probability, $p_{j|i}$, that w_i would pick w_j as its neighbor if neighbors were (stochastically) picked in proportion to their probability density, P_i, under a Gaussian kernel centered at w_i. Mathematically, the condition probability $p_{j|i}$ is given by:

$$p_{j|i} = exp\frac{\left(-||wi-wjl||^2/2\sigma_i^2\right)}{\sum_{k\neq i} exp\left(-||wi-wk||^2/2\sigma_i^2\right)} \quad \textbf{(3)}$$

where σ_i is the variance of the Gaussian centered on w_i. Importantly, the variance of the Gaussian kernel adapts to the local density of the data around each point to produce a probability density (P_i) with a fixed perplexity.

The *perplexity* hyper-parameter specifies the number of effective nearest neighbors included in the conditional probability calculations, where smaller *perplexity* values result in maps that are biased towards representing local relationships and larger values result in maps that represent local relationships with increasing consideration to any global structure that might

exist. More formally, perplexity is a measure of information describing how well a probability distribution predicts a sample and is defined as 2 $H^{(Pi)}$, where $H^{(Pi)}$ is the Shannon entropy of *Pi* measured in bits.

The low-dimensional probability (z_{ij}) distributions takes a similar form to p_{ji} except that a long-tailed Students t-distribution replaces the Gaussian distribution. The long-tail of the Students t-distribution ensures that moderately close points in the high-dimensional space are modeled by larger distance in the low-dimensional space, and as a result, eliminates any unwanted attractive forces between moderately dissimilar points that would have otherwise resulted in 'crowding' [37] in the low-dimensional representation between neighboring clusters with very different densities. Additionally, this particular form of the Students t-distribution (single degree of freedom) ensures that the low-dimensional representation is (mostly) scale invariant, meaning that clusters of points will interact in the same way as individual points. The effect is that the functional relationships between neurons are preserved across multiple scales of organization[37].

The overall aim of t-SNE is to find a low-dimensional data representation that minimizes any mismatch between the high-dimensional joint probability density, P, and the Students-t based joint probability distribution, Z. The minimization of the cost function is performed via gradient decent, with a gradient given by the equation:

$$\frac{\delta C}{\delta y_i} = 4 \sum_j (p_{ij} - z_{ij})\,(u_i - u_j)\left(1 + \|u_i - u_j\|^2\right)^{-1}. \quad \textbf{(4)}$$

In order to reduce computational complexity of this step, we perform a preliminary round of dimensionality reduction using principal component analysis (PCA) to project the $N \times N$ Neuron Similarity Matrix into smaller dimensional space (e.g., 50-*d*). The t-SNE algorithm then refines the resulting linear transform by minimizing the single Kullback-Leibler divergence between P and Q over multiple iterations. Seeding with a low-dimensional PCA projection also ensures that the algorithm converges to the same solution across repeated runs of the algorithm. This step results in the *d*x*N* Neuron Similarity map.

Step 4a: k-means Clustering Algorithm

k-means algorithm

The k-means algorithm is an unsupervised clustering method that partitions data into *k* clusters. We elected to use the k-means algorithm to cluster neurons in the NS map into putative functional ensembles because of its efficiency and its empirically evaluated performance in detecting functional groupings of neurons.

k-means clustering aims to partition the t-SNE outputs into *k* number of clusters, such that each data point belongs to a cluster with the nearest mean (see next section for selection of *k* value). The algorithm works iteratively to assign each data point (u_i) to one of the *C* centroids based on proximity, where the centroids have been initialized at the random locations C_1, C_2..., C_m. After all points are assigned, new centroids are calculated from the assigned data points. This procedure is repeated for a specific number of iterations, e.g., 100,

or until the centroids no longer move between iterations. The algorithm aims to minimize the sum of the squared error (SSE) between each data point:

$$SSE = \sum_{j=1}^{m} \sum_{i=1}^{l} \left\| u_i^{(j)} - c_j \right\|^2 \quad \textbf{(5)}$$

Step 4b: SIMNETS Silhouette Analysis and Significance Test

Silhouette Analysis We used silhouette analysis to assess the quality of the k-means clustering partitions across a range of values of k [75], with the goal of finding an optimal partition number for the data. A silhouette value (h_i) is a measure of how similar point y_i is to other data points in its assigned cluster c_j as compared to other clusters:

$$h_i = \frac{(b_i - a_i)}{max(a_i, b_i)}, \quad \textbf{(6)}$$

where a_i is the average distance from u_i to other points in its assigned cluster c_j, and b_i is the average distance from u_i to points in the other clusters, minimized over all possible cluster configurations. An optimal number of clusters $\hat{k}$ is the value of k that maximizes the average silhouette $(\hat{h})$ value for k = 2..., k_f. Silhouette values ranges from −1 to +1, where a high value indicates that u_i is well matched to its own cluster and poorly matched to neighboring clusters. In general, a maximized average silhouette below 0.25 indicates data that are not structured while a value below 0.5 would indicate poor or potentially spurious clusters [75]. In the next section, we outline a procedure for testing the statistical significance of the cluster number to determine if the data can be partitioned into statistically meaningful clusters.

Shuffle-based Significance Test We developed a significance test for the purpose of determining the likelihood of detecting a given number of clusters by chance under the null hypothesis that there is no genuine covariation relationship between the inherent structures of the SSIM matrices. The significance test involves generating a null-distribution of silhouette values based on shuffled data across a range of k values. In SIMNETS, functional similarities are captured by the pairwise measures of correlation between the single neuron SSIM matrices. Our test relies on a shuffling procedure that destroys the pairwise dependencies between the SSIM matrices, and subsequently, any significant measures of correlation in the Neuron Correlation Matrix.

Our approach is inspired by the Mantel test [77], a permutation based procedure that tests the significance of the observed correlation between two symmetrical matrices. The intuition of a Mantel test is that if a significant relationship exists between the values of matrix A and matrix B, then randomizing the rows and columns of one matrix will destroy any existing dependencies. As a result, the correlation between the shuffled matrix pair will tend to be lower than the original correlation value observed between the un-shuffled matrix pair. The probability of observing $r_{A.B}$ is then calculated as the proportion of permutations for which the shuffled correlation measures are smaller than or equal to $r_{A.B.}$ Here, we carry out a similar permutation operation on the SSIM matrices, in that we destroy any dependencies that exist between the matrices; however, we use the average silhouette value as the test statistic, rather than the correlation values, as is the case with the Mantel test.

The procedure involves a symmetrically shuffling of the rows/column of each N SSIM matrix separately, and re-calculating the pairwise correlations between the SSIM matrices to generate a new NxN Correlation Matrix. This NxN correlation matrix is then transformed into a new NS map using t-SNE (i.e., step 3), and a new set of silhouette values is calculated (i.e., step 4) for the range of tested k values. This procedure – SSIM matrix shuffling, steps 3 and step 4 from SIMNETS – are repeated to generate a null-distribution of average silhouette values that is approximately normally distributed (e.g., 1000+ iterations). If the observed maximized average silhouette value $\hat{h}$ falls above the empirically calculated $(1-\alpha)100\%$ confidence interval, then the detected number of $\hat{k}$ clusters is considered statistically meaningful.

Explanation of relevant symbols associated with SIMNETS algorithm

Table 1: explanation of relevant symbols

Symbol	Description	Calculation/Selection
q	temporal sensitivity value for Victor-Purpura spike train metric	hyper-parameter (user specified)
perplexity	perplexity: a measure of information that controls the number of effective nearest neighbors in t-SNE dimensionality reduction algorithm	hyper-parameter (user specified)
d	desired dimensions for low-dimensional Neuron Similarity Map	hyper-parameter (user specified)
k	k-means partition number/ range of cluster numbers to test for during k-means clustering	range of k-values (user specified)
$\hat{k}$	Optimal number of detected clusters Estimated from Silhouette analysis over a range of k values	output (eqt. 6)
$\hat{h}$	maximized average silhouette value	output (eqt. 7)

Hardware, Software, and Processing Time

All analyses were run on a Dell PC with an Intel Xeon® Processor and 24 GB of RAM. All analyses were run using MATLAB® software from MathWorks, version 9.4, R2018. Armadillo, a C ++ linear algebra library (called from within MATLAB) performed some of the main matrix operations.

On this hardware, analysis run-time for a dataset of 100 neurons (100 one-second spike trains per neuron) takes approximately 3.0 seconds, while 1000 neurons takes approximately 4 minutes (Supplementary Fig. 9). By comparison, calculating the pairwise cross-correlation values (without a jitter/shuffle correction procedure) for 100 or 1000 neurons takes approximately 6 minutes and 6 hours, respectively, on the same hardware. The computational complexity of the SIMNETS algorithm scales almost linearly with neuron number and exponentially with the number of spike trains Introducing a new neuron only requires generating a single new SSIM matrix, however, adding a new trial requires generating a new SSIM matrix for each neuron. The low computational cost of adding new

neurons means that datasets with large numbers of neurons could be functionally categorized and clustered in a reasonable amount of time (< 1 hr for 5,000 neurons).

Comparison Method Implementation:

Direct Comparison (DC) Method

We emphasize that the SIMNETS algorithm does not directly compare the firing patterns *between* different neurons. Instead, pairwise comparisons are performed between common spike trains of a single neuron, on a neuron-by-neuron basis. The between neuron comparisons are then made between all pairs of the single neuron SSIM matrices. This allows the algorithm to find neurons that generate a set of spike trains with common signature spike train geometries (i.e., set of distances), rather than grouping neurons based on the degree of coordination between their moment-to-moment firing patterns. In order to evaluate the effectiveness of this strategy, we compared the performance of SIMNETS to a representation of traditional approaches that directly compare the spike trains of different neurons.

The Direct Comparison (DC) method computes pairwise spike train similarities between matching 'trials' between neuron pairs. In contrast to the SIMNETS method, the t-SNE dimensionality reduction step is applied to an NxN matrix of distance values, rather than SSIM correlation values. That is, for a set of neurons N = $\{n_1, n_2, \ldots n_k\}$, the DC method builds a NxN matrix, M, where each $M_{x,y}$ entry is the sum of the spike train distances between the spike trains of neuron n_x, S = $\{S_1, S_2, \ldots, S_j\}$, and the spike trains of neuron n_y, U = $\{U_1, U_2, \ldots, S_j\}$:

$$M_{x,y} = \sum_{i=1}^{j} D\left(S_{i,}, U_i\right), \tag{7}$$

where D is a vector of Victor-Purpura spike train distance of length j (equation 1).

Firing Rate Covariation (FRC) Method

Each matrix element in the Kiani et al (2015) neuron correlation matrix is calculated as the Pearson's correlation between the trial-by-trial firing rate outputs of a neuron pair. The FRC method outlined here is ultimately inspired by the Kiani et al, however, in an effort to present a fair assessment of the SIMNETS algorithm, we adjusted the Kiani et al (2015) algorithm to make it more directly comparable to the SIMNETS algorithm. Specifically, after generating an NxN neuron correlation matrix by calculating the single trial firing rate correlation statistics for each neuron pair, we applied the t-SNE dimensionality reduction methods to the correlation matrix, as outlined above in step 3 of the SIMNETS algorithm. The FRC method *perplexity* values were selected to match those used when applying SIMNETS to the various datasets.

Cross-Correlation Histogram (CCH) Method

A spike train Cross-Correlation histogram[78] (CCH) is a function that indicates the probability of observing a spike in neuron A given a spike in neuron B across multiple time-lags. Here, we calculated CCHs using 1ms time bins at lags spanning ± 200 ms. Histograms were smoothed with a 5 ms kernel (kernel = [0.05 0.25 0.40 0.25 0.05]), jitter-corrected to account for spurious peaks in the histograms[79], and normalized by the geometric mean of each neuron's firing rate (see Kohn et al. (2008) for details). To facilitate comparisons across methods, the temporal resolution of the jitter correction was set to match the temporal resolution of the SIMNETS temporal accuracy parameter, q, for a given dataset, e.g., when q = 20, jitter = 50 ms (1/q).

A single CCH NxN neuron correlation matrix was generated for the V1, M1, and CA1 datasets. Each matrix element was calculated as the area under the curve above a 2 STD threshold (average of first and last 50 ms in histogram) and between +/- 10 ms (see Supplementary Figure 8 for examples). Each CCH neuron correlation matrix was then transformed into a low-dimensional NS map using t-SNE (*perplexity* was similar to that used for SIMNETS analyses). For more details on jitter method and cross-correlation analysis, see Amarasingham et al. (2012) and Smith and Kohn, (2008).

DM, FRC, CCH Shuffle-based statistical test A shuffle-based statistical test was used to assess the significance of the detected clusters in the DM, FRC, and CCH NS maps. In all cases, the trials labels were shuffled before calculating the NxN neuron similarities measure that make up their respective NxN correlation/distances matrices. The NxN matrices were transformed into a new NS map using t-SNE (i.e., SIMNETS step 3), and a new set of silhouette values is calculated (i.e., SIMNETS step 4) for the range of tested k values. This procedure is repeated to generate a null-distribution of average silhouette values that is approximately normally distributed (e.g., 1000+ iterations). If the observed maximized average silhouette value $\hat{h}$ falls above the empirically calculated $(1-\alpha)\ 100\%$ confidence interval, then the detected number of $\hat{k}$ clusters is considered statistically meaningful.

Synthetic Dataset – Data Simulation and Analysis

Spike train Simulation We simulated the spiking activity of a population of N = 180 synthetic neurons that consisted of 3 functionally distinct 'ensembles' (E_1, E_2, E_3) of 60 neurons. Each functional ensemble was designed to produce similar spike-trains for two non-modulating conditions, referred to as the 'baseline' conditions, and a different pattern for a third condition, referred to as the 'modulating' condition. For example, ensemble E_1 was modulated during condition A and exhibited the same baseline activity spike pattern during both conditions B and C, whereas ensemble E_2 was modulated during condition B and exhibited the same baseline spike pattern during conditions A and C, etc. Each E_i ensemble was further divided into three sub-groups of n = 20 neurons, where each sub-group altered

their spike-train patterns between the active and baseline states according to one of three different encoding strategies:

(1) Rate coding: firing rate increased by 50% for the modulating condition (all spike times were randomly chosen)
(2) Temporal coding: the two baseline conditions and the modulating condition were associated with specific (randomly generated) temporal sequences of spikes. The number of spikes was kept constant across baseline and modulating conditions. Spike times were jittered by +/- 50 ms for each trial.
(3) Mixed temporal/rate coding: Similar to the temporal coding, but the spikes were jittered in a temporal window of 5 ms. Additionally, the modulating condition included 25% more spikes.

In order to simulate stochastic variation in spiking patterns, 50% of the spikes were randomly removed for each condition. A total of 30 seconds of simulated recording time was generated, with the trial condition changing every second between A, B, and C patterns.

SIMNETS Cluster Characterization We demonstrate SIMNETS' ability to cluster functionally similar neurons in the synthetic dataset by comparing the pairwise similarity measures between and within ground-truth functional ensembles. We compare the distributions of the pairwise correlation values from the Neuron Similarity/Distance Matrices for neurons from the same functional ensemble ('Within' ensemble neuron pairs) and different functional ensembles ('Between' ensemble neuron pairs). Data on the y-axis was plotted as the normalized percentage of neurons for the Within-ensemble neurons pairs and the Between-ensemble neuron pairs. A rank-sum statistical test was carried out on the Within and Between distributions of similarity values, using an alpha value of $\alpha = .001$.

Neural Datasets – Task Description and Data Analysis

Primate Primary Visual Cortex

Task Description We analyzed a previously described dataset of 112 primary visual (V1) single-units (which we refer to as neurons) recorded in an anesthetized *Macaca fascicularis* using a 96-channel microelectrode array [40,46]Briefly, sinusoidal gratings were presented at 6 different orientations θ = {0°, 30°, 60°, 90°, 120°, 150°} and 2 drift directions (rightward and leftward drift, orthogonal to orientation). Each stimulus was presented 112 times for 1.28 seconds. The position and size of the stimuli was sufficient to cover the receptive fields of all recorded neurons. For more details on the data processing and task design, see Smith and Kohn (2008) and Kohn and Smith, (2016).

Single Neuron Functional Characterization We characterized the preferred orientation of each V1 neuron by fitting a Gaussian distribution to the firing rate function R:

$$R(\theta) = \hat{A}.exp^{\frac{-(\theta-\mu)^2}{2\sigma^2}}, \tag{9}$$

where θ is the stimulus orientation, $\hat{A}$ is the peak response, μ the mean, and σ^2 is variance of

the Gaussian. The function takes on a maximum value at $\theta = \mu$, for $\theta = [0, 180)$, which corresponds to the neuron's preferred orientation. We chose not to use drift-direction preferences when characterizing the functional properties of the neurons[80,81] as only a very small percentage exhibited significant differences in the magnitude of their peak responses for drift-direction.

V1 SIMNETS Analysis We extracted 1 second of spiking data from the first 30 repetitions of each stimulus ($S = 360$, spike trains), starting 0.28 seconds after stimulus onset. Only a small fraction of the total number of recorded trials was used in the analysis (25%) as we wanted to demonstrate SIMNETS ability to clusters neurons in datasets where only a small number of trials are available.

V1 SIMNETS Map Features We used a circular-linear correlation (r_{cl}) analysis to assess SIMNETS' ability to group neurons according to their functional similarities. The correlation between each neuron's preferred orientation and its location along each dimension in the low dimensional map y_i was calculated using:

$$r_{\theta,y} = \frac{cov(\theta_i, y_i)}{\sigma_A \sigma_B}, \tag{10}$$

where σ_A and σ_B are the standard deviation of the neurons' preferred orientations and y represents the neurons' locations in the map. A high correlation value indicates a strong relationship between a neuron's preferred orientation/direction and map location and demonstrates that functionally similar neurons were mapped to nearby regions of the map. The r_{cl} value for the dimension with the highest value was reported.

V1 SIMNETS Cluster Features We then characterized the functional properties of the detected sub-nets by calculating ensemble tuning functions (ETFs). ETFs were calculated by normalizing and averaging the joint firing rates across all neurons in each detected SIMNETS cluster. A bootstrap resampling method was used to determine the significance of the peaks in the ensemble tuning function. The null distribution describing the probability of getting the observed peak response by chance was computed from the EFTs of multiple sub-sets of neurons, randomly sampled from the population across 10,000 iterations. The sample size on each iteration was set to match the number of neurons in the given detected cluster. A response that falls above (or below) the 99% confidence interval is considered significant.

Primate Primary Motor Cortex

Task Description SIMNETS was applied to previously described dataset of 103 *Macaca mulatta* primary motor (M1) cortex neurons (i.e., single-units) recorded during a planar 8-direction reaching task [33,41]The single-unit activity was simultaneously recorded from the upper limb area of primary motor cortex using a chronically implanted microelectrode array. The monkey was operantly trained to move a cursor that matched its hand location to targets projected onto a horizontal reflective surface. A visual cue was used to signal movement direction during a variable duration instructed delay period (1 – 1.6 s) to one of eight radially distributed targets on the screen with the associated reach angles of φ = {0°, 45°, 90°, 135°, 180°, 225°, 270°, 315°}. At the end of the instructed delay period, the central target was extinguished, instructing the monkey to reach towards the previously cued target.

M1 SIMNETS Analysis We analyzed 1 second of neural data from correct trials (S = 114, trials), starting 0.1 second before movement onset. Characterization of the detected SIMNETS clusters is similar to that described in section 5.1.

We characterized the preferred direction of each M1 neuron by fitting a von Mises distribution [47] to the firing rate function R:

$$r(\varphi) = \beta + h.exp\big(h.cos(\varphi - \mu)\big) \tag{11}$$

where β is the offset of the function, h is the depth of the tuning, is the reach angle and μ is preferred reach direction of the cell. The function takes on a maximum value at $\varphi = \mu$, which corresponds to the neuron's preferred reach angle.

M1 SIMENTS map and cluster fea

Rat Hippocampal CA1

Task Description We applied SIMNETS to a previously described dataset of rat hippocampal neurons[52,53] made publicly available by the Collaborative Research in Computational Neuroscience (CRCNS) data-sharing repository [42]. The neurons were simultaneously recorded from the CA1 hippocampal region using multi-site silicon probes while the rat performed a spatial navigation task in a maze. Briefly, the rat was trained to run through the arms of a 'figure-8' maze in a left/right alternating manner in order to receive a reward. The left/right track runs were interleaved with a wheel-run period that functionally served as a memory delay-period. The rat performed T = 17 correct trials (T_r = 8, left trials; T_l = 9, right trials), taking on average 4.3 seconds to reach the rewards located at either end of the arms. The rat's path along each arm of the track was linearized and divided into small (50 cm) or large (325 cm) spatial bins for the spatial firing field analysis or SIMNETS analysis, respectively (see next section for more details).

CA1 Single Neuron Functional Characterization The rat's path along each arm of the track was linearized and divided into 50 mm spatial bins when generating the spatial firing field maps. Bins corresponding to reward locations and the inter-trial activity were excluded from the analysis, leaving a total of 39 bins for each of the left and right trajectories, where the first 19 spatial bins were common to both trajectories. We generated a separate spatial firing map for the left and right trajectories of each neuron by dividing the number of spikes in the i-th bin by the rat's occupancy time t_i, and used a Gaussian kernel (width = 3 bins/150 mm) to smooth across the firing rates in each bin. Neurons that did not exhibit a 5 Hz firing rate in at least 1 spatial bin were not included in the analysis, leaving a total of N = 80 neurons. We characterized the neurons as non-place cells (n = 20, nPC) or place cells (n = 60, PC) based on their spatial firing properties and an information-theoretic measure of the spatial information in their spikes [52,82,83]. Neurons were classified as having place cell-like activity if the firing rate in three contiguous bins exceeded the mean of all other firing fields by 20% [83,84] (using 2.5 STD of the out-of-field firing rate produced similar results[52]) and if their information content exceeded 0.5 bits/spike[82] on either the left or right trajectories. The

spatial information metric, I_{spike}, is a measure of the extent to which a neuron's spiking activity can be used to predict the rat's position along the track. The spatial information content of the neuron (measured in bits/spike) is defined as:

$$I_{spike} = \sum_{i=1}^{l} P_i \frac{v_i}{V} log_2 \frac{v_i}{V}, \tag{12}$$

where P_i is the occupancy probability, v_i is the firing rate in the *i*-th bin, and *V* is the overall mean firing rate of the cell across all bins in trajectory.

CA1 SIMNETS Analysis We divided the *T* linearized trajectories into six 325 cm spatial bins and extracted 0.75 s spike trains beginning at the time that rat entered a given bin. The time window duration was selected to capture the smaller receptive fields ~ 0.6 s but still include a large portion of the average place field width (1 s – see Pastalkova et al. 2008). The spatial bin size corresponds to the approximate distance travelled in this time window. This resulted in S = 108 spike train events. The SIMNETS algorithm was applied to the resulting N x S spike train

CA1 NS Map Features - Distance Analysis We compared the distances between neurons characterized as non-PCs and PCs in order to demonstrate the ability of the four tested methods to cluster the non-PCs to a specific region of the NS map. The Euclidean distances were calculated between all pairs of non-PCs (referred to as 'Within' pairs) and between pairs of non-PCs and PCs ('Between' pairs) in the low dimensional NS map. A rank-sum statistical test was carried out on the Within and Between distance distributions using a significance threshold of p = 0.001.

CA1 Ensemble Firing Rate Maps were generated for each of the three example SIMNETS clusters by averaging across the normalized single neuron firing rate spatial maps for all neurons in the *i*-th cluster. A bootstrap resampling method was used to test for significant peaks (p = .01) in the ensemble spatial firing map (i.e., place-field cell-assemblies). The procedure involved randomly sampling a subset of neurons from the neuron population, where each subset was equal in size to the number of neurons in the *i*-th cluster. A new ensemble tuning function was calculated over repeated iterations of this procedure (10000 iterations). The resulting null distribution describes the probability of getting the observed peak response if the detected neuron clusters were selected at random. A response that falls above (or below) the 99% confidence interval is considered significant.

CA1 Ensemble Activity Similarity Maps are low-dimensional neural activity state-space maps that capture the relationships between neural ensemble activity patterns on individual trials. This method of visualizing low-dimensional projections of ensemble activity has previously been described by this group [33,49] Generating the low-dimensional Ensemble Activity Similarity maps consists of three steps. 1) Calculate the pairwise spike train distances between the *i*-th spike train event of neuron *j* and all other spike trains belonging to that neuron $S_j = \{S_{1,j}, S_{2,j}, \ldots, S_{m,j}\}$. This results in a vector of pairwise spike-train distances $d(S_{i,j})$ of length *m*. 2) the pair-wise similarity vectors for the *i*-th trial of all *n* neurons of a given ensemble are concatenated and combined into a matrix, resulting in an *m* x *mn* matrix (D_{ensemble}). The resulting *m* x *mn* pairwise distance matrix constitutes the relational embedding of the entire data set. 3) The final step consists of projecting the high-dimensional *m* x *mn* distance matrix down into a *m* x *d* matrix, where *d* is the desired dimension of the projection.

Table 2: Summary of datasets

Symbol	Synthetic Dataset	V1 Dataset	M1 Dataset	CA1 Dataset
N	180	112	103	80
T	30	360	114	18
S	30	360	114	119
***S*pike train duration (s)**	1	1	1	0.75

Abbreviations: *N*, neuron number; *T*, trial number; *S*, spike train events number.

Table 3: Summary of SIMNETS Inputs/Outputs for each dataset

Symbol	Synthetic Dataset	V1 Dataset	M1 Dataset	CA1 Dataset
q	[0, 10, 200]	20	10	50
perplexity	50	30	20	15
d	3	3	3	3
$\hat{k}$	4,3,3	3	3	6
k	2-15+	2-15+	2-15+	2-15+
$\hat{h}$	0.5, 0.99, 0.95	0.71	0.68	0.78

Abbreviations: *q*, temporal accuracy parameter; *perplexity*, t-SNE 'number of nearest neighbors' parameter; *d*, dimensions; $\hat{k}$, optimal number of statistically significant clusters; $\hat{h}$, peak average silhouette value.

Bibliography

1. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13,** 51–62 (2011).
2. Zagha, E., Ge, X. & McCormick, D. A. Competing Neural Ensembles in Motor Cortex Gate Goal-Directed Motor Output. *Neuron* **88,** 565–577 (2015).
3. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321,** 1322–7 (2008).
4. Harris, K. D. Opinion: Neural signatures of cell assembly organization. *Nat. Rev. Neurosci.* (2005). doi:10.1038/nrn1669
5. Brown, E. N., Kass, R. E. & Mitra, P. P. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat. Neurosci.* (2004). doi:10.1038/nn1228
6. Briggman, K. L., Abarbanel, H. DI & Kristan, W. B. From crawling to cognition: analyzing the dynamical interactions among populations of neurons. *Current Opinion in Neurobiology* (2006). doi:10.1016/j.conb.2006.03.014
7. Buzsáki, G. Large-scale recording of neuronal ensembles. *Nat. Neurosci.* **7,** 446–451 (2004).
8. Wallace, D. J. & Kerr, J. N. D. Chasing the cell assembly. *Current Opinion in Neurobiology* (2010). doi:10.1016/j.conb.2010.05.003
9. Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K. & Palm, G. Dynamics of Neuronal Firing Correlation: Modulation of " Effective Connectivity ". *JOURNALOFNEUROPHYSIOLOGY Print. i n U.S.A* **61,** (1989).
10. Abeles, M. & Gat, I. Detecting precise firing sequences in experimental data. *J. Neurosci. Methods* (2001). doi:10.1016/S0165-0270(01)00364-8
11. Baker, S. N. & Gerstein, G. L. Improvements to the Sensitivity of Gravitational Clustering for Multiple Neuron Recordings. *Neural Comput.* **12,** 2597–2620 (2000).
12. Grün, S., Diesmann, M. & Aertsen, A. Unitary Events in Multiple Single-Neuron Spiking Activity: II. Nonstationary Data. *Neural Comput.* (2002). doi:10.1162/089976602753284464
13. Rgy, G., Iki, B. & Chrobak, J. I. Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr. Opin. Neurobiol.* **5,** 504–510 (1995).
14. Humphries, M. D. Spike-Train Communities: Finding Groups of Similar Spike Trains. *J. Neurosci.* (2011). doi:10.1523/JNEUROSCI.2853-10.2011
15. Singer, W. & Gray, C. M. Visual Feature Integration and the Temporal Correlation Hypothesis. *Annu. Rev. Neurosci.* **18,** 555–586 (1995).
16. Singer, W. Neuronal Synchrony: A Versatile Code for the Definition of Relations? *Neuron* **24,** 49–65 (1999).
17. Von Der Malsburg, C. *The Correlation Theory of Brain Function*.
18. Brody, C. D. Correlations Without Synchrony. *Neural Comput.* **11,** 1537–1551 (1999).
19. Shadlen, M. N. & Movshon, J. A. Synchrony Unbound. *Neuron* **24,** 67–77 (1999).
20. Abeles, M. & Gerstein, G. L. Detecting Spatiotemporal Firing Patterns Among Simultaneously Recorded Single Neurons. *JOURNALOFNEUROPHYSIOLOGY Print. i n U.S.A* **60,** (1988).
21. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural

recordings. *Nat. Neurosci.* **17,** 1500–9 (2014).
22. Machens, C. K., Romo, R. & Brody, C. D. Functional, But Not Anatomical, Separation of 'What' and 'When' in Prefrontal Cortex. *J. Neurosci.* **30,** 350–360 (2010).
23. Brody, C. D. Correlations Without Synchrony. *Neural Comput.* **11,** 1537–1551 (1999).
24. Truccolo, W., Hochberg, L. R. & Donoghue, J. P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13,** 105–111 (2010).
25. Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat. Commun.* **6,** 7759 (2015).
26. Churchland, M. M. & Shenoy, K. V. Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *J. Neurophysiol.* **97,** 4235–4257 (2007).
27. Chelaru, M. I. & Dragoi, V. Efficient coding in heterogeneous neuronal populations. *Proc. Natl. Acad. Sci. U. S. A.* **105,** 16344–9 (2008).
28. Lehky, S. R., Sereno, M. E. & Sereno, A. B. Population Coding and the Labeling Problem: Extrinsic Versus Intrinsic Representations. *Neural Comput.* **25,** 2235–2264 (2013).
29. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
30. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
31. Victor, J. D. & Purpura, K. P. Metric-space analysis of spike trains: theory, algorithms and application. *Netw. Comput. Neural Syst.* (1997). doi:10.1088/0954-898X_8_2_003
32. Zaidi, Q. *et al.* Perceptual spaces: mathematical structures to neural mechanisms. *J. Neurosci.* **33,** 17597–602 (2013).
33. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
34. Lehky, S. R., Sereno, M. E. & Sereno, A. B. Population Coding and the Labeling Problem: Extrinsic Versus Intrinsic Representations. *Neural Comput.* (2013). doi:10.1162/NECO_a_00486
35. Victor, J. D. & Purpura, K. P. Nature and Precision of Temporal Coding in Visual Cortex: A Metric-Space Analysis. *J. Neurophysiol. Pt-intd U.S.A* **76,** (1996).
36. Kriegeskorte, N. & Kievit, R. A. Representational geometry: integrating cognition, computation, and the brain. *Trends Cogn. Sci.* **17,** 401–12 (2013).
37. Van Der Maaten, L. & Hinton, G. *Visualizing Data using t-SNE. Journal of Machine Learning Research* **9,** (2008).
38. Grün, S., Diesmann, M. & Aertsen, A. Unitary Events in Multiple Single-Neuron Spiking Activity: II. Nonstationary Data. *Neural Comput.* **14,** 81–119 (2002).
39. Stuart, L., Walter, M. & Borisyuk, R. Visualisation of synchronous firing in multi-dimensional spike trains. in *BioSystems* (2002). doi:10.1016/S0303-2647(02)00084-9
40. Kohn, A., Smith, M. A. (2016). Utah array extracellular recordings of spontaneous and visually evoked activity from anesthetized macaque primary visual cortex (V1).

CRCNS.org (2016). Available at: http://dx.doi.org/10.6080/K0NC5Z4X.
41. Rao, N. G. & Donoghue, J. P. Cue to action processing in motor cortex populations. *J. Neurophysiol.* **111,** 441–53 (2014).
42. Pastalkova E, Wang Y, Mizuseki K, B. G. Simultaneous extracellular recordings from left and right hippocampal areas CA1 and right entorhinal cortex from a rat performing a left / right alternation task and other behaviors. (2015). Available at: http://dx.doi.org/10.6080/K0KS6PHF.
43. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220,** 149–166 (2013).
44. Kiani, R. *et al.* Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. *Neuron* **85,** 1359–73 (2015).
45. Gerstein, G. L. & Aertsen, A. M. Representation of cooperative firing activity among simultaneously recorded neurons. *J. Neurophysiol.* **54,** (1985).
46. Smith, M. A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* **28,** 12591–603 (2008).
47. Amirikian, B., Georgopoulos, A. P. & Georgopulos, A. P. Directional tuning profiles of motor cortical cells. *Neurosci. Res.* **36,** 73–9 (2000).
48. Rao, N. G. & Donoghue, J. P. Cue to action processing in motor cortex populations. *J. Neurophysiol.* **111,** 441–453 (2014).
49. Vargas-Irwin, C. E., Franquemont, L., Black, M. J. & Donoghue, J. P. Linking Objects to Actions: Encoding of Target Object and Grasping Strategy in Primate Ventral Premotor Cortex. *J. Neurosci.* **35,** 10888–97 (2015).
50. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. doi:10.1162/NECO_a_00684
51. Lillicrap, T. P. & Scott, S. H. Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics. *Neuron* **77,** 168–179 (2013).
52. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321,** 1322–7 (2008).
53. Mizuseki, K., Sirota, A., Pastalkova, E. & Buzsáki, G. Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* **64,** 267–280 (2009).
54. Park, E., Dvorak, D. & Fenton, A. A. Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS One* **6,** e22349 (2011).
55. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551,** 232–236 (2017).
56. Csicsvari, J. *et al.* Massively Parallel Recording of Unit and Local Field Potentials With Silicon-Based Electrodes. (2003). doi:10.1152/jn.00116.2003
57. Kim, D. H. *et al.* Pan-neuronal calcium imaging with cellular resolution in freely swimming zebrafish. *Nat. Methods* **14,** 1107–1114 (2017).
58. Park, E., Dvorak, D. & Fenton, A. A. Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS One* **6,** e22349 (2011).
59. Geiller, T., Fattahi, M., Choi, J.-S. & Royer, S. Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nat. Commun.* **8,** 14531 (2017).
60. Gerstein, G., Perkel, D. & Dayhoff, J. Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. *J. Neurosci.* **5,** (1985).

61. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220,** 149–166 (2013).
62. Humphries, M. D. Spike-Train Communities: Finding Groups of Similar Spike Trains. *J. Neurosci.* **31,** 2321–2336 (2011).
63. Fujisawa, S., Amarasingham, A., Harrison, M. T. & Buzsáki, G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11,** 823–833 (2008).
64. Wilson, N. R., Runyan, C. A., Wang, F. L. & Sur, M. Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature* **488,** 343–8 (2012).
65. Neske, G. T., Patrick, S. L. & Connors, B. W. Contributions of diverse excitatory and inhibitory neurons to recurrent network activity in cerebral cortex. *J. Neurosci.* **35,** 1089–105 (2015).
66. Paninski, L. Estimation of Entropy and Mutual Information. *Neural Comput.* (2003). doi:10.1162/089976603321780272
67. Granger, C. W. J. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica* **37,** 424–438 (1969).
68. Vicente, R. *et al.* Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *J Comput Neurosci* **30,** 45–67 (2011).
69. Kaminski, M. J. & Blinowska, K. J. A new method of the description of the information flow in the brain structures. *Biol. Cybern.* **65,** 203–210 (1991).
70. Blinowska, K. J. Review of the methods of determination of directed connectivity from multichannel data. *Medical and Biological Engineering and Computing* **49,** 521–529 (2011).
71. Ts'o, D. Y., Gilbert, C. D. & Wiesel, T. N. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* **6,** 1160–70 (1986).
72. Blanchard, T. C., Piantadosi, S. T. & Hayden, B. Y. Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. *J. Neurophysiol.* **119,** 1305–1318 (2018).
73. Chandrasekaran, C., Peixoto, D., Newsome, W. T. & Shenoy, K. V. Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nat. Commun.* **8,** 614 (2017).
74. Lloyd, S. P. *Least Squares Quantization in PCM*. *IEEE TRANSACTIONS ON INFORMATION THEORY* **28,** (1982).
75. Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **20,** 53–65 (1987).
76. Author, S., Kullback, S. & Leibler, R. A. *On Information*. *Source: The Annals of Mathematical Statistics* **22,** (1951).
77. Mantel, N. The detection of disease clustering and a generalized regression approach. *Cancer Res.* **27,** 209–20 (1967).
78. Perkel, D. H., Gerstein, G. L. & Moore, G. P. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys J* **7,** 419–440 (1967).
79. Amarasingham, A., Harrison, M. T., Hatsopoulos, N. G. & Geman, S. Conditional modeling and the jitter method of spike resampling. *J. Neurophysiol.* **107,** 517–31 (2012).
80. Gur, M. & Snodderly, D. M. Direction selectivity in V1 of alert monkeys: evidence for parallel pathways for motion processing. *J. Physiol.* **585,** 383–400 (2007).
81. Hawken, M. J., Parker, A. J. & Lund, J. S. Laminar organization and contrast sensitivity of direction-selective cells in the striate cortex of the Old World monkey.

J. Neurosci. **8,** 3541–8 (1988).
82. Skaggs, W. E., Skaggs, W. E., McNaughton, B. L., Gothard, K. M. & Markus, E. J. An Information-Theoretic Approach to Deciphering the Hippocampal Code. *IN* **5,** 1030--1037 (1993).
83. Bjerknes, T. L., Dagslott, N. C., Moser, E. I. & Moser, M.-B. Path integration in place cells of developing rats. *Proc. Natl. Acad. Sci. U. S. A.* **115,** E1637–E1646 (2018).
84. Jeffery, K. J., Donnett, J. G., Burgess, N. & O'Keefe, J. M. Directional control of hippocampal place fields. *Exp. brain Res.* **117,** 131–42 (1997).
85. Ayzenshtat, I., Jackson, J. & Yuste, R. Orientation Tuning Depends on Spatial Frequency in Mouse Visual Cortex. *eNeuro* **3,** (2016).

SIMNETS: a computationally efficient and scalable framework for identifying sub-networks of functionally similar neurons

Supplementary Information

Jacqueline B. Hynes [1, 2], **David M. Brandman**,[2, 3], **Jonas B. Zimmerman**[4], **John P. Donoghue**[1, 2, 4, 5], **Carlos E. Vargas-Irwin** [*, 1, 2,3]

1 Department of Neuroscience, Brown University, Providence, RI, USA.
2 Robert J. and Nancy D. Carney Institute for Brain Science, Brown University, Providence, RI, USA
3 Department of Surgery (Neurosurgery), Dalhousie University, Halifax, NS, Canada.
4 WYSS Institute, Chemin des Mines 9, CH-1202, Geneve, Switzerland.
5. Department of Engineering, Brown University, Providence, RI, USA.

∗Correspondence to: carlos_vargas_irwin@brown.edu

Keywords: functional sub-networks, neural ensembles, clustering algorithm, dimensionality reduction, spike train distances

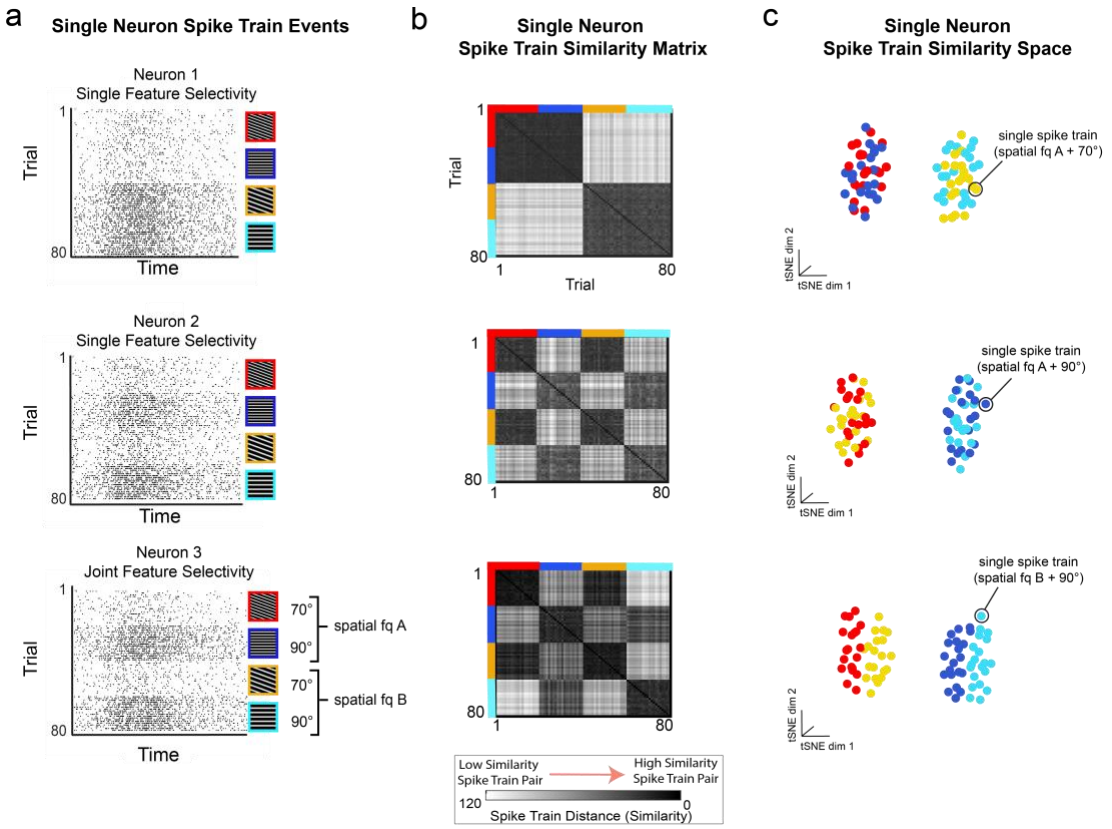

Supplementary Figure 1 | Representing the computational properties of three synthetic V1 neurons using SSIM matrices.

a. Spike raster plots showing simulated spike train outputs for three artificial V1 neurons (N1, N2, and N3) in response to repeated presentations of four different sinusoidal grating stimuli. Neuron receptive fields were modeled as a 2D symmetric Gabors[85], such that neurons N1 and N2 were sensitive to differences in either the stimulus spatial frequencies or the stimulus orientations, respectively, whereas neuron N3 was sensitive to differences in both the stimulus spatial frequencies and orientation. Each grating stimulus has a different spatial frequency and orientation combination: stimulus 1 (indicated by red square) and stimulus 2 (blue square) had the same spatial frequency ('spatial fq. A') but different orientations (i.e., 70° or 90°) and stimulus 3 (yellow square) and stimulus 4 (cyan square) had the same spatial frequency ('spatial fq. B') but different orientations. **b**. Single neuron Spike Train Similarity (SSIM) matrices for each of the simulated neurons. SSIM matrices describe the intrinsic relationship between a neuron's spike train outputs in terms of distance: similar spike trains correspond to lower distance values (dark pixels) and dissimilar spike-trains correspond to higher distance values (light pixels). Colored bars indicate the stimulus pairings associated with each pixel (similar notation to a). Each SSIM matrix can be interpreted as an abstract representation of the neuron's functional input-output mapping across the set of sampled conditions. Note that the ordering of the trials in the SSIM matrix is arbitrary, but needs to be consistent across neurons in order to compare one neuron to another in the data set. The correlation between pairs of SSIM matrices, can be used to quantify the relative functional similarity between the neuron pairs (e.g., Pearson's, $r_{n1.n3} = 0.13$; Pearson's, $r_{n2.n3} = 0.81$). **c**, Low-dimensional representation of the single neuron SSIM matrices. Dimensionality reduction tools can transform the high-dimensional SSIM data into a low-dimensional 'spike train relational maps'. Individual spike train events are represented as colored points and the similarity between spike trains corresponds to the distance between points in the space. The topology of the maps captures the essence of how each neuron categorizes the different stimulus inputs.

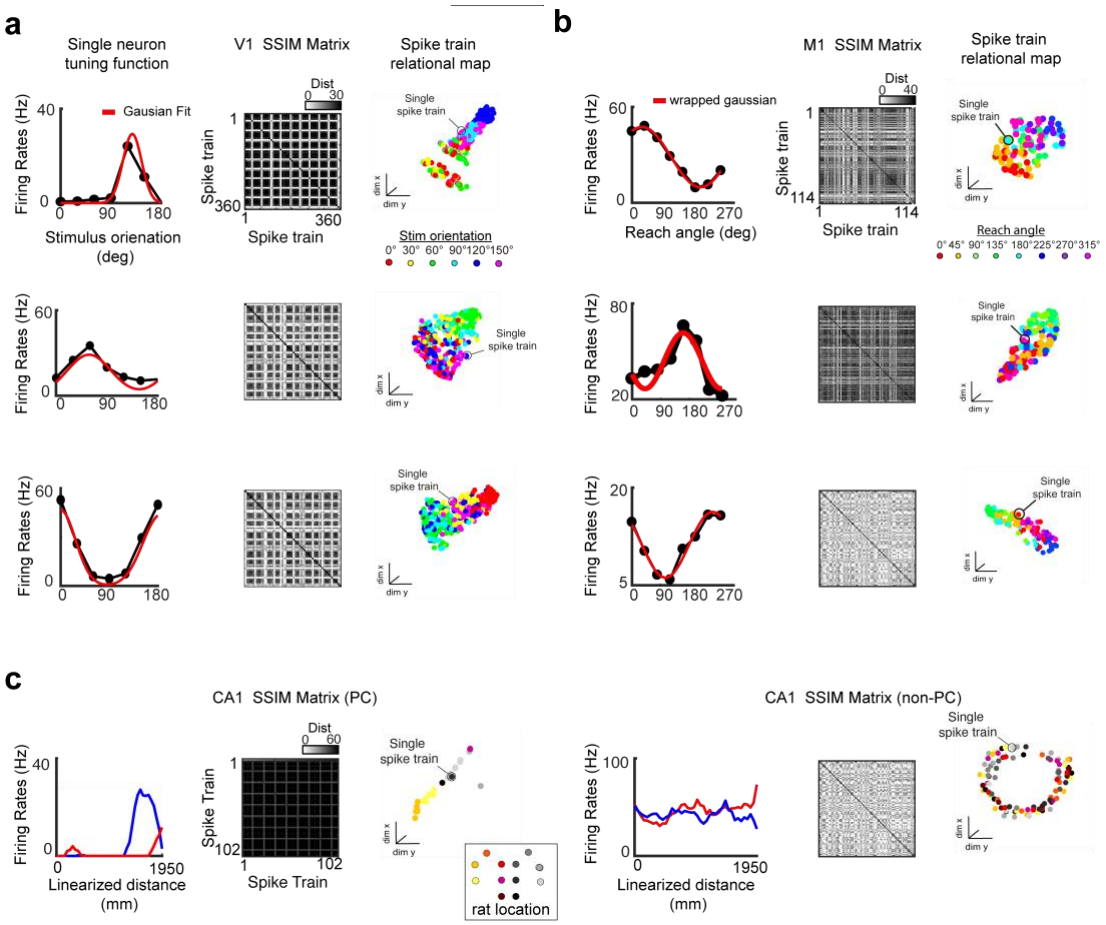

Supplementary Figure 2 | The information processing properties of example V1, M1, and Hippocampal CA1 neurons are captured in the structure of single neuron SSIM matrices and SSIM spaces.

a-b. Three single-neuron tuning functions (*left column*), single neuron SSIM matrices (*middle column*), and SSIM maps (*right column*) for example neurons from V1 (**a**) and M1(**b**) data sets. Each single neuron SSIM map, which represents the spike train outputs of a single neuron across each trial trials, was generated by applying t-SNE to the SSIM matrices[30]. Each point in the SSIM map represents an individual spike train generated by the neuron on a trial and the colors indicates the visual stimulus conditions (**a**) or movement conditions (**b**) corresponding to that trial. **c**. Place-dependent firing vectors, single neuron SSIM matrices, and single neuron SSIM map for an example PC and non-PC in CA1 dataset. Red and blue lines represent the neuron's average firing rate for the left and right trials, respectively. Each point in the spike train relational map represents individual spike trains and the colors indicate the rat's location during the left or right track runs.

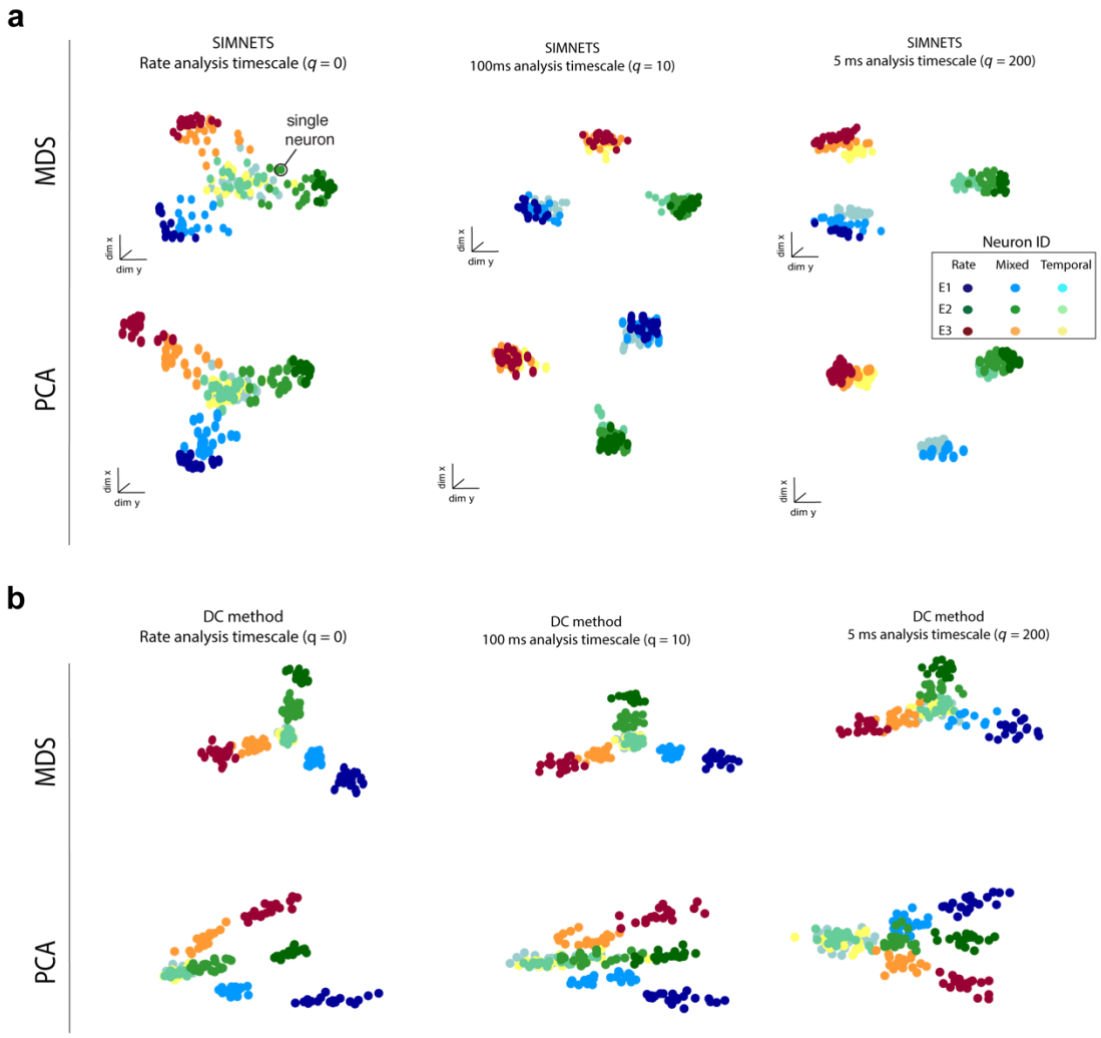

Supplementary Figure 3| SIMNETS and Direct Comparison (DC) Neuron Similarity (NS) maps for synthetic neuron population generate using alternative dimensionality reduction methods.

a. SIMNETS NS maps for three tested temporal accuracy values (columns) generated using Principal Component Analysis (PCA) (top row) and Multi-dimensional Scaling (MDS) (bottom row). Individual points represent single neurons and the colors indicate their ground-truth ensemble assignments/coding properties. **b**. Direct Comparison (DC) NS Maps for three tested temporal accuracy values (columns) generated using PCA (top row) and MDS (bottom row).

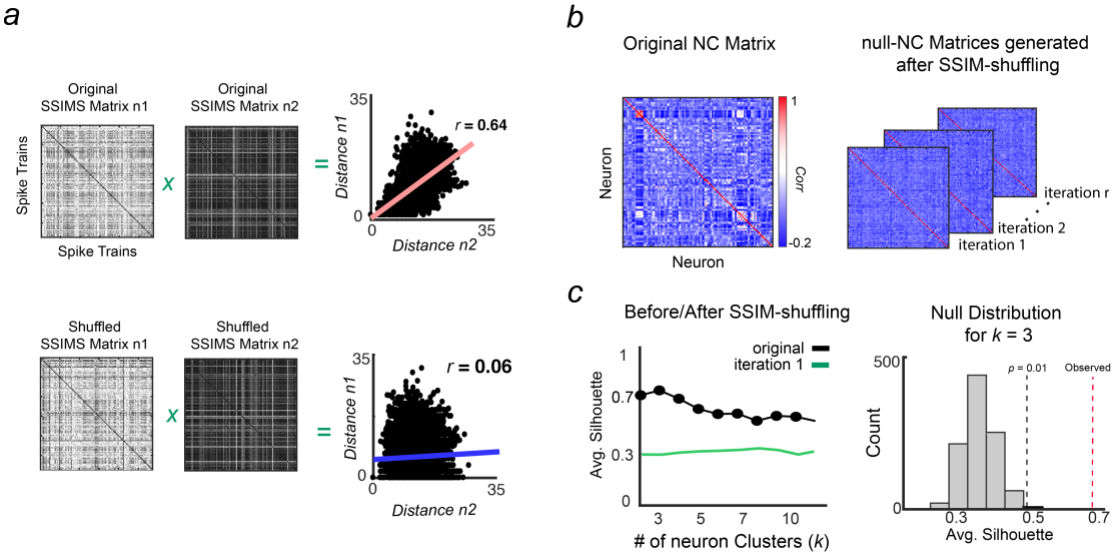

Supplementary Figure 4 | SIMNETS shuffle-based statistical test for assessing clustering number significance.

a. Two single-neuron SSIM matrices (n1 and n2) and scatter plots illustrating SSIM matrix correlation values before (top) and after (bottom) shuffling procedure (randomized rows and columns). Red line indicates a high correlation SSIM matrix pair (i.e., original) and blue line indicates a low-correlation SSIM matrix pair (i.e., shuffled). **b.** Original NS matrix and three example 'null' NS matrices. Each null NS matrix was generated from a single iteration of the SSIMS shuffling procedure (independently shuffle all *N* SSIM matrices) and recalculating the *N*x*N* pairwise SSIM matrix correlations. **c.** Left: original (black) and 'null' (green) cluster silhouette values for a single iteration of the shuffling procedure. The null silhouette values were calculated from a low-dimensional projection of NS matrix (i.e., NS map). Note: null silhouette values fall below the 0.5 mark (indicative of poor cluster separation) and peak around k = 3 in original values (black line) has disappeared. *Right*: distribution of null silhouette values and original silhouette value (red broken line) for k = 3 clusters.

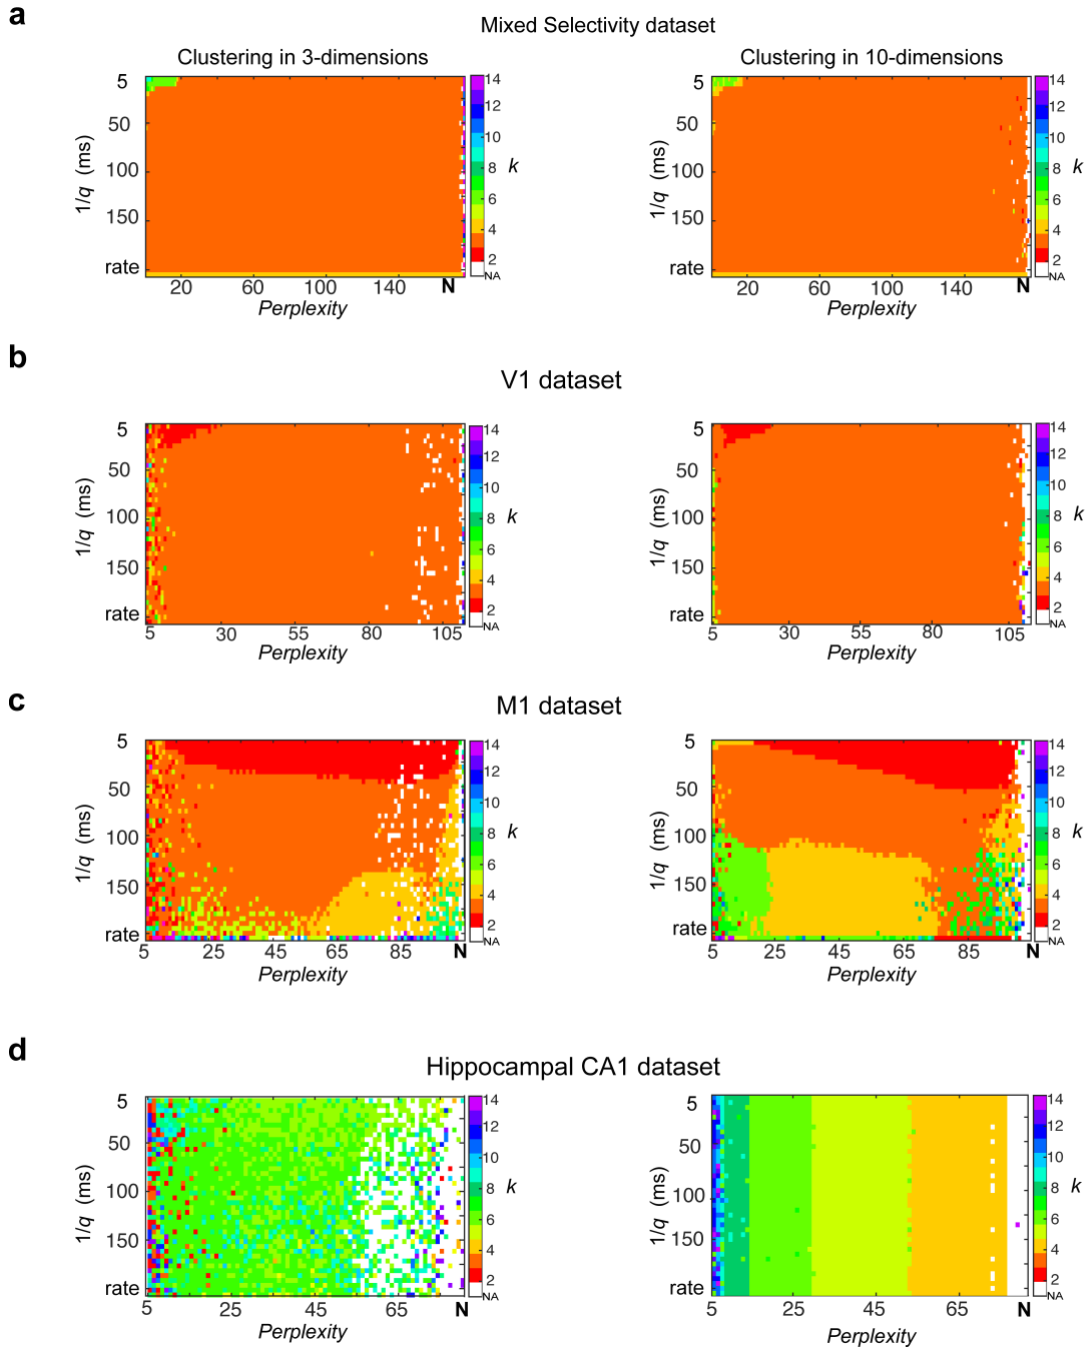

a.

Supplementary Figure 5 | Parameter sweep: number of SIMNETS clusters as function of *perplexity* and spike train comparison temporal accuracy parameter, *q*, in a high and low dimensional space.

a-d. SIMNETS clusters (*k*) as function of *perplexity* and temporal accuracy (1/*q*) in a 3-*d* t-SNE projection space (*left*) and a 10-*d* t-SNE projection space (*right*) for the synthetic (**a**), V1 (**b**), M1 (**c**), and Hippocampal CA1 dataset (**d**). Color bars indicate the number of detected SIMNETS clusters (determined through a silhouette analysis) for a given set of parameters

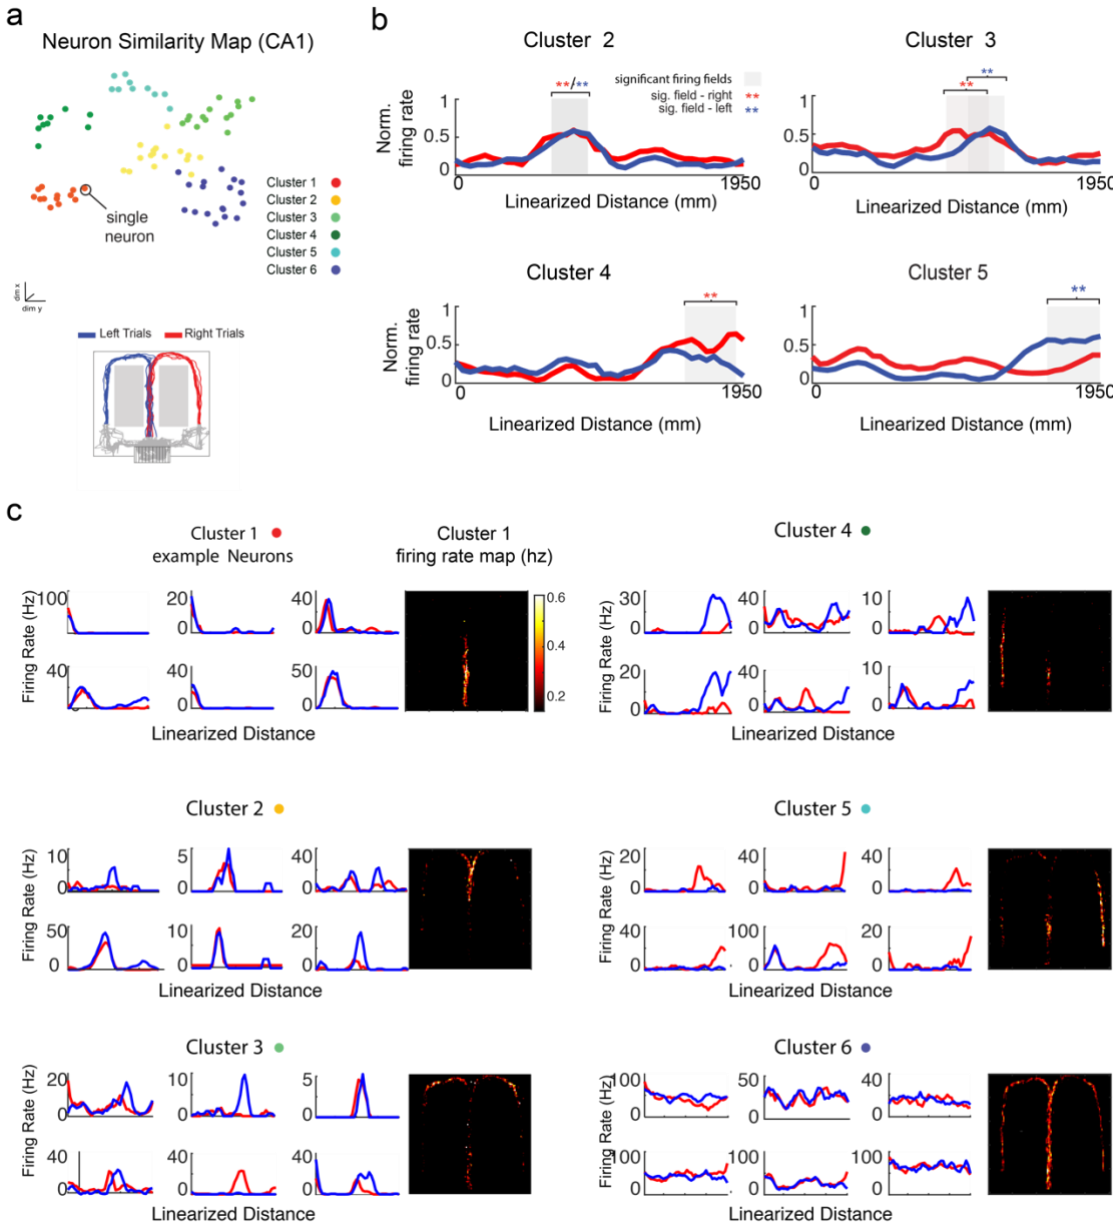

Supplementary Figure 6 | SIMNETS clusters detected in Hippocampal CA1 dataset.

a. SIMNETS NS Map (*top*) and the rat's location in the maze across left (blue line) and right (red line) track runs (*bottom*). **b**. Place-dependent firing rates for SIMNETS clusters 2-5 (see Fig. 7 for clusters 1 and 6). **c.** Example place-dependent firing field for individual neurons in each of the detected SIMNETS clusters, with 2D spatial firing maps (*right column*). Single neuron firing rate maps were average across neurons and normalized (see Methods for more details).

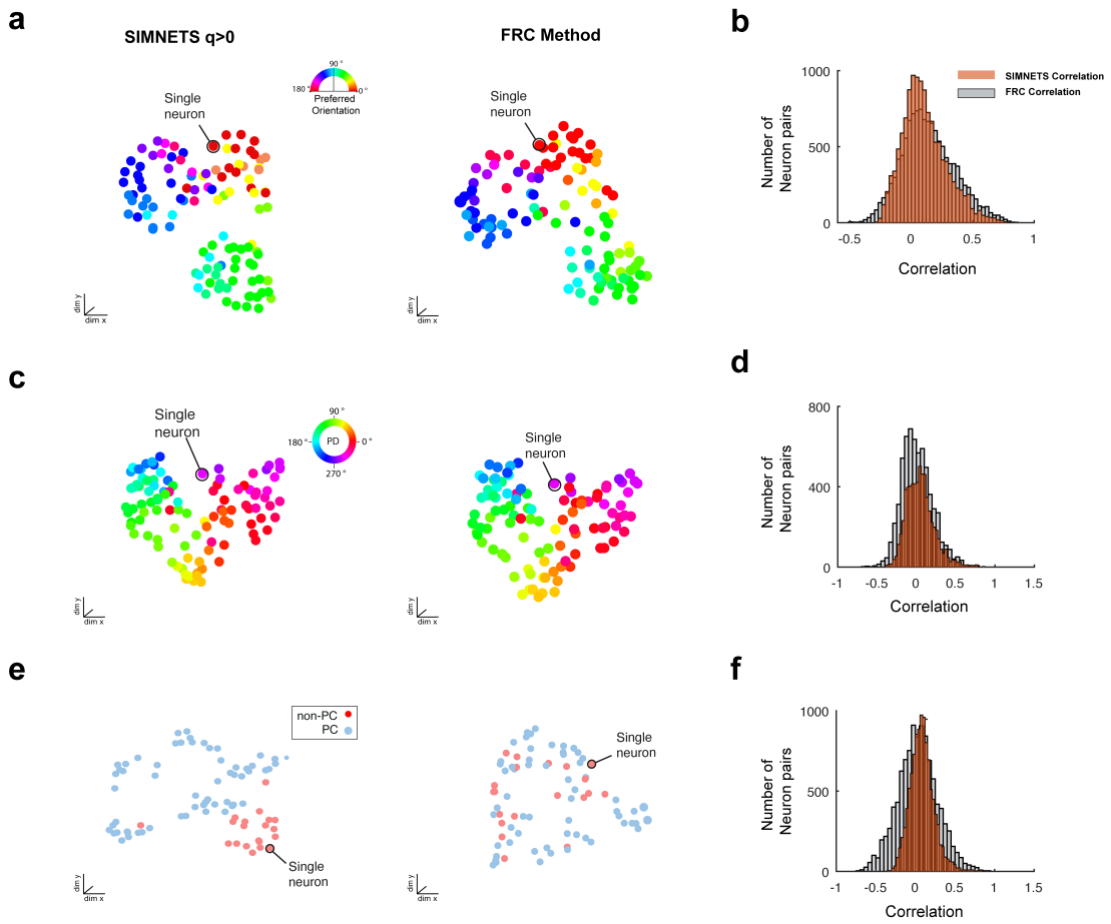

Supplementary Figure 7 | Comparison of SIMNETS and Firing Rate Covariation (FRC) method performance in organizing neurons according to their estimated functional properties.

a - d SIMNETS NS maps (*left column*) and FRC NS maps (*right column*) for the Synthetic (**a**), V1 (**b**), M1(**c**), and CA1 neuron datasets (**d**). Each dot represents a single neuron and the different colors indicate the simulated (**a**) or estimated information processing properties of the neurons (**b-d**). The FRC method failed to organize the Synthetic (**a**) and CA1 (**d**) neurons according to their functional properties. SIMNETS and FRC methods performed similarly in their ability to organize the M1(**c**) neuron populations according to their functional properties (SIMNETS $r$cl = 0.92, $p < 0.001$; FRC $r$cl = 0.93, $p < 0.001$), with SIMNETS marginally outperforming the FRC method for the V1 dataset (**b**) (SIMNETS $r$cl = 0.89, $p < 0.001$; FRC $r$cl = 0.86, $p < 0.001$). **e - h** Histograms showing the distributions for the SIMNETS and FRC NxN correlation values (Kruskal Wallis; p <0.001).

Supplementary Table 1: summary of statistical evaluation of distributions of 'non-PCs vs non-PCs' distances (*within*) and 'non-PCs vs PCs' distances (*between*) in the SIMNETS, DM, FRC, and CCH Neuron Similarity Maps for the CA1 dataset shows that SIMNETS is the only method capable of detecting statistically significant clusters of non-PCs vs. PCs.

Symbol	**SIMNETS**	**DM**	**FRC**	**CCH**
Detected statistically significant clusters of non-PCs and PCs	Yes	No	No	No
M±STD: *within* distances	45.5 ± 27.8	18.1 ± 11.8	37.0± 15.8	156.4 ± 102.6
M±STD: *between* distances	92.8±25.4	34.2±15.5	41.4±7.7	166.9±88.1
Ratio of the means: *between*/*within* distances	2.04 (best separation)	1.89	1.12	1.06 (worst separation)
Ranksum test *within* vs. *between* distances	$p = 6.223^{-64}$	$p = 3.681^{-29}$	$p = 0.002$	$p = 0.042$
$\hat{h}$	0.78 ($p < 0.01$)	0.55 (**n.s.**)	0.70 ($p < 0.01$)	0.65 ($p < 0.01$)
$\hat{k}$	6	-	5	2
r_s location (best dim) vs. avg. firing rate	0.03 ($p = 0.79$)	0.97 ($p < 0.001$)	-0.4 ($p = 0.69$)	-0.4361 ($p = 0.96$)

Abbreviations: $\hat{h}$, peak silhouette value; $\hat{k}$, optimal number of statistically significant clusters; n.s., non-significant, r_s, Spearman's correlation statistic.

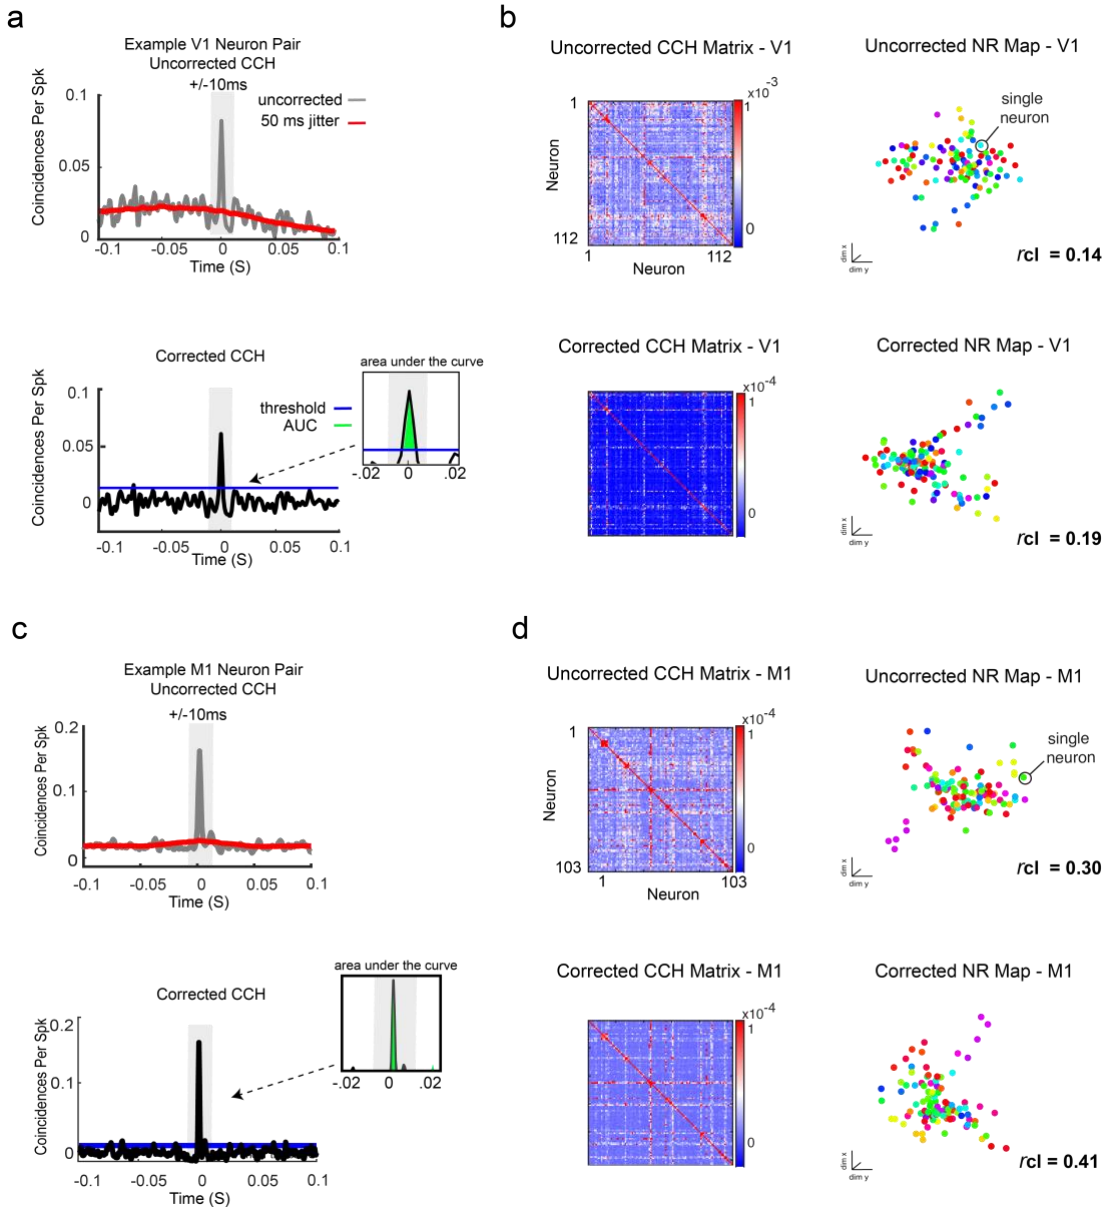

Supplementary Figure 8 | Poor recapitulation of the estimated V1 and M1 functional properties when t-SNE is applied to *N*x*N* cross-correlation matrix.

a, **c**. Multiple time-lag cross-correlation histograms (CCHs) for example V1 (**a**) and M1 (**c**) neuron pair, showing uncorrected CCHs (*top*) and jitter-corrected CCHs (*bottom*). Red line indicates the expected (spurious) cross-correlation values based on slow time-scale and trial-to-trial firing rate covariations, calculated using a jitter resampling procedure[79]. The AUC (area under the curve; inset) above the 2 STD threshold (blue line) and between +/- 10ms for each neuron for the uncorrected and corrected pairwise CCHs. **b**,**d**. *left*: uncorrected and corrected CCH matrices for V1 (**b**) and M1(**d**) neuron populations were transformed into low-dimensional NS maps (right). Individual points represent neurons and the colors correspond to the neurons' preferred stimulus orientation (**b**) or preferred reach angle (**d**). Color notation similar to Fig. 5d and 6d. For more details on jitter correction and cross-correlation analysis, see Methods, Amarasingham et al. (2012), and Smith and Kohn, (2008).

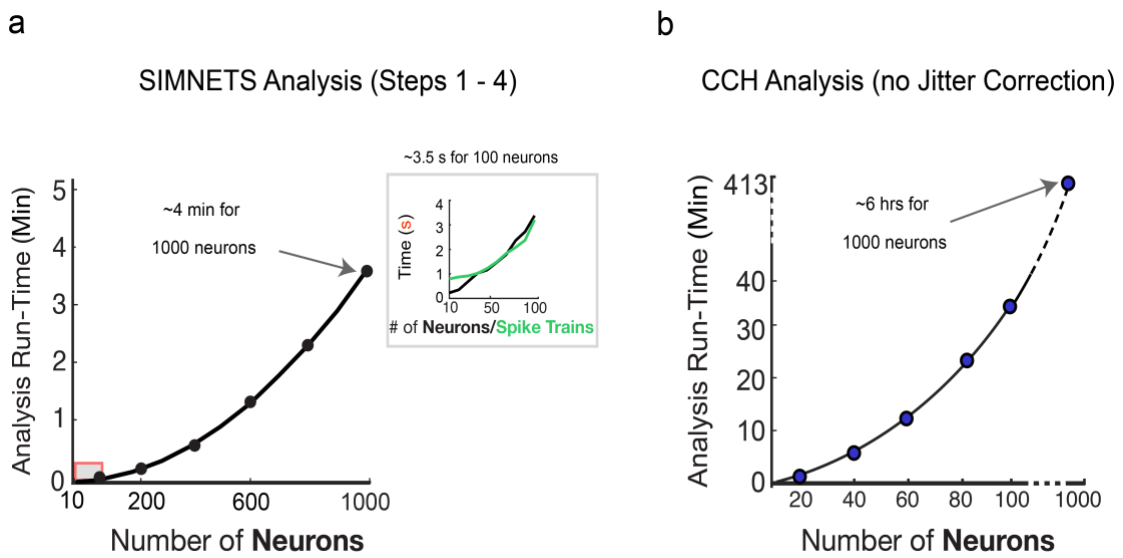

Supplementary Figure 9 | Comparison of SIMNETS and cross-correlation analysis (CCH) run-time as a function of neuron number

a. SIMNETS analysis run-time (min) for steps 1 – 4 (without statistical test) for M1 neuron population as a function of neuron number. Smaller subsets of the population were used for analyses involving 100 neurons or less and multiple concatenated sets of the population were pooled together to generate larger neuron populations for analyses involving > 100 neurons. *Inset*: SIMNETS analysis run-time (s) as a function of neuron number (*black line*) or spike train number (*green line*). **b**. CCH analysis run-time (not including Jitter correction) as a function of neuron number. Dash line indicates an interval discontinuity in the *x*-/*y*-axes. The SIMNETS has a computational complexity that is near-linear: $O(n^{1+\varepsilon})$, where ε is infinitesimally small.

Bibliography

1. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13,** 51–62 (2011).
2. Zagha, E., Ge, X. & McCormick, D. A. Competing Neural Ensembles in Motor Cortex Gate Goal-Directed Motor Output. *Neuron* **88,** 565–577 (2015).
3. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321,** 1322–7 (2008).
4. Harris, K. D. Opinion: Neural signatures of cell assembly organization. *Nat. Rev. Neurosci.* (2005). doi:10.1038/nrn1669
5. Brown, E. N., Kass, R. E. & Mitra, P. P. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat. Neurosci.* (2004). doi:10.1038/nn1228
6. Briggman, K. L., Abarbanel, H. DI & Kristan, W. B. From crawling to cognition: analyzing the dynamical interactions among populations of neurons. *Current*

Opinion in Neurobiology (2006). doi:10.1016/j.conb.2006.03.014
7. Buzsáki, G. Large-scale recording of neuronal ensembles. *Nat. Neurosci.* **7,** 446–451 (2004).
8. Wallace, D. J. & Kerr, J. N. D. Chasing the cell assembly. *Current Opinion in Neurobiology* (2010). doi:10.1016/j.conb.2010.05.003
9. Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K. & Palm, G. Dynamics of Neuronal Firing Correlation: Modulation of " Effective Connectivity ". *JOURNALOFNEUROPHYSIOLOGY Print. i n U.S.A* **61,** (1989).
10. Abeles, M. & Gat, I. Detecting precise firing sequences in experimental data. *J. Neurosci. Methods* (2001). doi:10.1016/S0165-0270(01)00364-8
11. Baker, S. N. & Gerstein, G. L. Improvements to the Sensitivity of Gravitational Clustering for Multiple Neuron Recordings. *Neural Comput.* **12,** 2597–2620 (2000).
12. Grün, S., Diesmann, M. & Aertsen, A. Unitary Events in Multiple Single-Neuron Spiking Activity: II. Nonstationary Data. *Neural Comput.* (2002). doi:10.1162/089976602753284464
13. Rgy, G., Iki, B. & Chrobak, J. I. Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr. Opin. Neurobiol.* **5,** 504–510 (1995).
14. Humphries, M. D. Spike-Train Communities: Finding Groups of Similar Spike Trains. *J. Neurosci.* (2011). doi:10.1523/JNEUROSCI.2853-10.2011
15. Singer, W. & Gray, C. M. Visual Feature Integration and the Temporal Correlation Hypothesis. *Annu. Rev. Neurosci.* **18,** 555–586 (1995).
16. Singer, W. Neuronal Synchrony: A Versatile Code for the Definition of Relations? *Neuron* **24,** 49–65 (1999).
17. Von Der Malsburg, C. *The Correlation Theory of Brain Function*.
18. Brody, C. D. Correlations Without Synchrony. *Neural Comput.* **11,** 1537–1551 (1999).
19. Shadlen, M. N. & Movshon, J. A. Synchrony Unbound. *Neuron* **24,** 67–77 (1999).
20. Abeles, M. & Gerstein, G. L. Detecting Spatiotemporal Firing Patterns Among Simultaneously Recorded Single Neurons. *JOURNALOFNEUROPHYSIOLOGY Print. i n U.S.A* **60,** (1988).
21. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17,** 1500–9 (2014).
22. Machens, C. K., Romo, R. & Brody, C. D. Functional, But Not Anatomical, Separation of 'What' and 'When' in Prefrontal Cortex. *J. Neurosci.* **30,** 350–360 (2010).
23. Brody, C. D. Correlations Without Synchrony. *Neural Comput.* **11,** 1537–1551 (1999).
24. Truccolo, W., Hochberg, L. R. & Donoghue, J. P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13,** 105–111 (2010).
25. Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat. Commun.* **6,** 7759 (2015).
26. Churchland, M. M. & Shenoy, K. V. Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *J. Neurophysiol.* **97,** 4235–4257 (2007).
27. Chelaru, M. I. & Dragoi, V. Efficient coding in heterogeneous neuronal populations. *Proc. Natl. Acad. Sci. U. S. A.* **105,** 16344–9 (2008).
28. Lehky, S. R., Sereno, M. E. & Sereno, A. B. Population Coding and the Labeling

Problem: Extrinsic Versus Intrinsic Representations. *Neural Comput.* **25,** 2235–2264 (2013).
29. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
30. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
31. Victor, J. D. & Purpura, K. P. Metric-space analysis of spike trains: theory, algorithms and application. *Netw. Comput. Neural Syst.* (1997). doi:10.1088/0954-898X_8_2_003
32. Zaidi, Q. *et al.* Perceptual spaces: mathematical structures to neural mechanisms. *J. Neurosci.* **33,** 17597–602 (2013).
33. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. *Neural Comput.* **27,** 1–31 (2015).
34. Lehky, S. R., Sereno, M. E. & Sereno, A. B. Population Coding and the Labeling Problem: Extrinsic Versus Intrinsic Representations. *Neural Comput.* (2013). doi:10.1162/NECO_a_00486
35. Victor, J. D. & Purpura, K. P. Nature and Precision of Temporal Coding in Visual Cortex: A Metric-Space Analysis. *J. Neurophysiol. Pt-intd U.S.A* **76,** (1996).
36. Kriegeskorte, N. & Kievit, R. A. Representational geometry: integrating cognition, computation, and the brain. *Trends Cogn. Sci.* **17,** 401–12 (2013).
37. Van Der Maaten, L. & Hinton, G. *Visualizing Data using t-SNE*. *Journal of Machine Learning Research* **9,** (2008).
38. Grün, S., Diesmann, M. & Aertsen, A. Unitary Events in Multiple Single-Neuron Spiking Activity: II. Nonstationary Data. *Neural Comput.* **14,** 81–119 (2002).
39. Stuart, L., Walter, M. & Borisyuk, R. Visualisation of synchronous firing in multi-dimensional spike trains. in *BioSystems* (2002). doi:10.1016/S0303-2647(02)00084-9
40. Kohn, A., Smith, M. A. (2016). Utah array extracellular recordings of spontaneous and visually evoked activity from anesthetized macaque primary visual cortex (V1). *CRCNS.org* (2016). Available at: http://dx.doi.org/10.6080/K0NC5Z4X.
41. Rao, N. G. & Donoghue, J. P. Cue to action processing in motor cortex populations. *J. Neurophysiol.* **111,** 441–53 (2014).
42. Pastalkova E, Wang Y, Mizuseki K, B. G. Simultaneous extracellular recordings from left and right hippocampal areas CA1 and right entorhinal cortex from a rat performing a left / right alternation task and other behaviors. (2015). Available at: http://dx.doi.org/10.6080/K0KS6PHF.
43. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220,** 149–166 (2013).
44. Kiani, R. *et al.* Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. *Neuron* **85,** 1359–73 (2015).
45. Gerstein, G. L. & Aertsen, A. M. Representation of cooperative firing activity among simultaneously recorded neurons. *J. Neurophysiol.* **54,** (1985).
46. Smith, M. A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* **28,** 12591–603 (2008).
47. Amirikian, B., Georgopoulos, A. P. & Georgopulos, A. P. Directional tuning profiles of motor cortical cells. *Neurosci. Res.* **36,** 73–9 (2000).

48. Rao, N. G. & Donoghue, J. P. Cue to action processing in motor cortex populations. *J. Neurophysiol.* **111,** 441–453 (2014).
49. Vargas-Irwin, C. E., Franquemont, L., Black, M. J. & Donoghue, J. P. Linking Objects to Actions: Encoding of Target Object and Grasping Strategy in Primate Ventral Premotor Cortex. *J. Neurosci.* **35,** 10888–97 (2015).
50. Vargas-Irwin, C. E., Brandman, D. M., Zimmermann, J. B., Donoghue, J. P. & Black, M. J. Spike Train SIMilarity Space (SSIMS): A Framework for Single Neuron and Ensemble Data Analysis. doi:10.1162/NECO_a_00684
51. Lillicrap, T. P. & Scott, S. H. Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics. *Neuron* **77,** 168–179 (2013).
52. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321,** 1322–7 (2008).
53. Mizuseki, K., Sirota, A., Pastalkova, E. & Buzsáki, G. Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* **64,** 267–280 (2009).
54. Park, E., Dvorak, D. & Fenton, A. A. Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS One* **6,** e22349 (2011).
55. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551,** 232–236 (2017).
56. Csicsvari, J. *et al.* Massively Parallel Recording of Unit and Local Field Potentials With Silicon-Based Electrodes. (2003). doi:10.1152/jn.00116.2003
57. Kim, D. H. *et al.* Pan-neuronal calcium imaging with cellular resolution in freely swimming zebrafish. *Nat. Methods* **14,** 1107–1114 (2017).
58. Park, E., Dvorak, D. & Fenton, A. A. Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS One* **6,** e22349 (2011).
59. Geiller, T., Fattahi, M., Choi, J.-S. & Royer, S. Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nat. Commun.* **8,** 14531 (2017).
60. Gerstein, G., Perkel, D. & Dayhoff, J. Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. *J. Neurosci.* **5,** (1985).
61. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220,** 149–166 (2013).
62. Humphries, M. D. Spike-Train Communities: Finding Groups of Similar Spike Trains. *J. Neurosci.* **31,** 2321–2336 (2011).
63. Fujisawa, S., Amarasingham, A., Harrison, M. T. & Buzsáki, G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11,** 823–833 (2008).
64. Wilson, N. R., Runyan, C. A., Wang, F. L. & Sur, M. Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature* **488,** 343–8 (2012).
65. Neske, G. T., Patrick, S. L. & Connors, B. W. Contributions of diverse excitatory and inhibitory neurons to recurrent network activity in cerebral cortex. *J. Neurosci.* **35,** 1089–105 (2015).
66. Paninski, L. Estimation of Entropy and Mutual Information. *Neural Comput.* (2003). doi:10.1162/089976603321780272
67. Granger, C. W. J. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica* **37,** 424–438 (1969).
68. Vicente, R. *et al.* Transfer entropy—a model-free measure of effective connectivity

for the neurosciences. *J Comput Neurosci* **30,** 45–67 (2011).
69. Kaminski, M. J. & Blinowska, K. J. A new method of the description of the information flow in the brain structures. *Biol. Cybern.* **65,** 203–210 (1991).
70. Blinowska, K. J. Review of the methods of determination of directed connectivity from multichannel data. *Medical and Biological Engineering and Computing* **49,** 521–529 (2011).
71. Ts'o, D. Y., Gilbert, C. D. & Wiesel, T. N. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* **6,** 1160–70 (1986).
72. Blanchard, T. C., Piantadosi, S. T. & Hayden, B. Y. Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. *J. Neurophysiol.* **119,** 1305–1318 (2018).
73. Chandrasekaran, C., Peixoto, D., Newsome, W. T. & Shenoy, K. V. Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nat. Commun.* **8,** 614 (2017).
74. Lloyd, S. P. *Least Squares Quantization in PCM*. *IEEE TRANSACTIONS ON INFORMATION THEORY* **28,** (1982).
75. Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **20,** 53–65 (1987).
76. Author, S., Kullback, S. & Leibler, R. A. *On Information*. *Source: The Annals of Mathematical Statistics* **22,** (1951).
77. Mantel, N. The detection of disease clustering and a generalized regression approach. *Cancer Res.* **27,** 209–20 (1967).
78. Perkel, D. H., Gerstein, G. L. & Moore, G. P. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys J* **7,** 419–440 (1967).
79. Amarasingham, A., Harrison, M. T., Hatsopoulos, N. G. & Geman, S. Conditional modeling and the jitter method of spike resampling. *J. Neurophysiol.* **107,** 517–31 (2012).
80. Gur, M. & Snodderly, D. M. Direction selectivity in V1 of alert monkeys: evidence for parallel pathways for motion processing. *J. Physiol.* **585,** 383–400 (2007).
81. Hawken, M. J., Parker, A. J. & Lund, J. S. Laminar organization and contrast sensitivity of direction-selective cells in the striate cortex of the Old World monkey. *J. Neurosci.* **8,** 3541–8 (1988).
82. Skaggs, W. E., Skaggs, W. E., McNaughton, B. L., Gothard, K. M. & Markus, E. J. An Information-Theoretic Approach to Deciphering the Hippocampal Code. *IN* **5,** 1030--1037 (1993).
83. Bjerknes, T. L., Dagslott, N. C., Moser, E. I. & Moser, M.-B. Path integration in place cells of developing rats. *Proc. Natl. Acad. Sci. U. S. A.* **115,** E1637–E1646 (2018).
84. Jeffery, K. J., Donnett, J. G., Burgess, N. & O'Keefe, J. M. Directional control of hippocampal place fields. *Exp. brain Res.* **117,** 131–42 (1997).
85. Ayzenshtat, I., Jackson, J. & Yuste, R. Orientation Tuning Depends on Spatial Frequency in Mouse Visual Cortex. *eNeuro* **3,** (2016).