



Stimuli Reduce the Dimensionality of Cortical Activity

Luca Mazzucato¹, Alfredo Fontanini^{1,2} and Giancarlo La Camera^{1,2*}

¹ Department of Neurobiology and Behavior, State University of New York at Stony Brook, Stony Brook, NY, USA, ² Graduate Program in Neuroscience, State University of New York at Stony Brook, Stony Brook, NY, USA

The activity of ensembles of simultaneously recorded neurons can be represented as a set of points in the space of firing rates. Even though the dimension of this space is equal to the ensemble size, neural activity can be effectively localized on smaller subspaces. The dimensionality of the neural space is an important determinant of the computational tasks supported by the neural activity. Here, we investigate the dimensionality of neural ensembles from the sensory cortex of alert rats during periods of ongoing (inter-trial) and stimulus-evoked activity. We find that dimensionality grows linearly with ensemble size, and grows significantly faster during ongoing activity compared to evoked activity. We explain these results using a spiking network model based on a clustered architecture. The model captures the difference in growth rate between ongoing and evoked activity and predicts a characteristic scaling with ensemble size that could be tested in high-density multi-electrode recordings. Moreover, we present a simple theory that predicts the existence of an upper bound on dimensionality. This upper bound is inversely proportional to the amount of pair-wise correlations and, compared to a homogeneous network without clusters, it is larger by a factor equal to the number of clusters. The empirical estimation of such bounds depends on the number and duration of trials and is well predicted by the theory. Together, these results provide a framework to analyze neural dimensionality in alert animals, its behavior under stimulus presentation, and its theoretical dependence on ensemble size, number of clusters, and correlations in spiking network models.

Keywords: gustatory cortex, dimensionality, hidden markov models, ongoing activity, mean field theory, spiking network model, metastable dynamics

OPEN ACCESS

Edited by:

Ruben Moreno-Bote,
Universitat Pompeu Fabra, Spain

Reviewed by:

Martin Paul Nawrot,
Universität zu Köln, Germany
Iñigo Arandia-Romero,
Universitat Pompeu Fabra, Spain

*Correspondence:

Giancarlo La Camera
giancarlo.lacamera@stonybrook.edu

Received: 31 August 2015

Accepted: 02 February 2016

Published: 17 February 2016

Citation:

Mazzucato L, Fontanini A and La Camera G (2016) Stimuli Reduce the Dimensionality of Cortical Activity. *Front. Syst. Neurosci.* 10:11. doi: 10.3389/fnsys.2016.00011

INTRODUCTION

Understanding the dynamics of neural activity and how it is generated in cortical circuits is a fundamental question in Neuroscience. The spiking activity of ensembles of simultaneously recorded neurons can be represented in terms of sequences of firing rate vectors, as shown e.g., in frontal (Abeles et al., 1995; Seidemann et al., 1996; Durstewitz et al., 2010), gustatory (Jones et al., 2007; Mazzucato et al., 2015), motor (Kemere et al., 2008), premotor and somatosensory cortex (Ponce-Alvarez et al., 2012). The dimension of each firing rate vector is equal to the number of ensemble neurons N and the collection of rate vectors across trials takes the form of a set of points in the N -dimensional space of firing rates. Such points may not fill the whole space, but be restricted to lie inside a lower-dimensional subspace (see Ganguli et al., 2008). Roughly, dimensionality is the minimal number of dimensions necessary to provide an accurate description

of the neural dynamics. If ensemble neurons are independent of each other, neural activities at different times will scatter around in the space of firing rate, filling a large portion of the space. In this case, dimensionality will be maximal and equal to the size of the ensemble N . At the other extreme, if all neurons are strongly correlated, ensemble activity localizes along a line. In this case, dimensionality is minimal and equal to one. These simple examples suggest that dimensionality captures information about the structure of a cortical circuit and the functional relations among the simultaneously recorded neurons, such as their firing rates correlation computed over timescales of hundreds of milliseconds.

Different definitions of dimensionality have been introduced for different tasks and across neural systems (Ganguli et al., 2008; Churchland et al., 2010a; Abbott et al., 2011; Ganguli and Sompolinsky, 2012; Cadieu et al., 2013; Rigotti et al., 2013; Gao and Ganguli, 2015). Such measures of dimensionality can shed light on the underlying neural computation; for example, they can predict the onset of an error trial in a recall task (Rigotti et al., 2013), or can allow the comparison of classification accuracy between different brain areas (e.g., IT vs. V4) and synthetic algorithms (Cadieu et al., 2013). Here, we investigate a measure of dimensionality closely related to the firing rate correlations of simultaneously recorded neurons (Abbott et al., 2011); such correlations may provide a signature of feature-based attention (Cohen and Maunsell, 2009) and other top-down cognitive factors (Nienborg et al., 2012). We elucidate the dependence of dimensionality on experimental parameters, such as ensemble size and interval length, and we show that it varies across experimental conditions. We address these issues by comparing recordings of ensembles of neurons from the gustatory cortex (GC) of alert rats to a biologically plausible network model based on neural clusters with recurrent connectivity. This model captures neural activity in GC during periods of ongoing and stimulus-evoked activity, explaining how the spatiotemporal dynamics of ensemble activity is organized in sequences of metastable states and how single-neuron firing rate distributions are modulated by stimulus presentation (Mazzucato et al., 2015). Here, we show that the same model expounds the observed dependence of dimensionality on ensemble size and how such dependence is reduced by the presentation of a stimulus. By comparing the clustered network model with a homogeneous network without clusters, we find that the clustered network has a larger dimensionality that depends on the number of clusters and the firing rate correlations among ensemble neurons. A simple theory explains these results and allows extrapolating the scaling of dimensionality to very large ensembles. Our theory shows that recurrent networks with clustered connectivity provide a substrate for high-dimensional neural representations, which may lead to computational advantages.

METHODS

Experimental Procedures

Adult female Long Evans rats were used for this study (Samuelsen et al., 2012; Mazzucato et al., 2015). Animals received *ad lib*.

access to food and water, unless otherwise mentioned. Movable bundles of 16 microwires attached to a “mini-microdrive” (Fontanini and Katz, 2006; Samuelsen et al., 2012) were implanted in GC (AP 1.4, ML \pm 5 from bregma, DV -4.5 from dura). After electrode implantation, intra-oral cannulae (IOC) were inserted bilaterally (Phillips and Norgren, 1970; Fontanini and Katz, 2005). At the end of the surgery a positioning bolt for restraint was cemented in the acrylic cap. Rats were given at least 7 days for recovery before starting the behavioral procedures outlined below. All experimental procedures were approved by the Institutional Animal Care and Use Committee of Stony Brook University and complied with University, state, and federal regulations on the care and use of laboratory animals. More details can be found in Samuelsen et al. (2012).

Rats were habituated to being restrained and receiving fluids through IOCs, and then trained to self-deliver water by pressing a lever following a 75 dB auditory cue at a frequency of 4 KHz. The interval at which lever-pressing delivered water was progressively increased to 40 ± 3 s (ITI). During experimental sessions additional tastants were automatically delivered at random times near the middle of the ITI, at random trials and in the absence of the anticipatory cue. A computer-controlled, pressurized, solenoid-based system delivered $\sim 40 \mu\text{l}$ of fluids (opening time ~ 40 ms) directly into the mouth through a manifold of 4 polyimide tubes slid into the IOC. The following four tastants were delivered: 100 mM NaCl, 100 mM sucrose, 100 mM citric acid, and 1 mM quinine HCl. Water ($\sim 50 \mu\text{l}$) was delivered to rinse the mouth clean through a second IOC 5 s after the delivery of each tastant. Each tastant was delivered for at least 6 trials in each condition. Upon termination of each recording session the electrodes were lowered by at least $150 \mu\text{m}$ so that a new ensemble could be recorded.

Evoked activity periods were defined as the interval after tastant delivery (time $t = 0$ in our figures) and before water rinse (time $t = 5$ s). Only trials in which the tastants were automatically delivered were considered for the analysis of evoked activity, to minimize the effects of cue-related expectations (Samuelsen et al., 2012). Ongoing activity periods were defined as the 5 s-long intervals at the end of each inter-trial period.

The behavioral state of the rat was monitored during the experiment for signs of disengagement. Erratic lever pressing, inconstant mouth movements and fluids dripping from the mouth indicated disengagement and led to the termination of the experiment. In addition, since disengagement from the task is also reflected in the emergence of high power μ oscillations in local field potentials, occurrences of such periods were removed offline and not analyzed further (Fontanini and Katz, 2008).

Data Analysis

Single neuron action potentials were amplified, bandpass filtered (at 300–8 KHz), digitized and recorded to a computer (Plexon, Dallas, TX). Single units of at least 3:1 signal-to-noise ratio were isolated using a template-matching algorithm, cluster cutting techniques and examination of inter-spike interval plots (Offline Sorter, Plexon, Dallas, TX). All data analyses and model simulations were performed using custom software

written in Matlab (Mathworks, Natick, MA, USA), Mathematica (Wolfram Research, Champaign, IL), and C. Starting from a pool of 299 single neurons in 37 sessions, neurons with peak firing rate lower than 1 Hz (defined as silent) were excluded from further analysis, as well as neurons with a large peak around the 6–10 Hz in the spike power spectrum, which were considered somatosensory (Katz et al., 2001; Samuelsen et al., 2012; Horst and Laubach, 2013). Only ensembles with 3 or more simultaneously recorded neurons were further analyzed (167 non-silent, non-somatosensory neurons from 27 ensembles). We analyzed ongoing activity in the 5 s interval preceding either the auditory cue or taste delivery, and evoked activity in the 5 s interval following taste delivery in trials without anticipatory cue, wherein significant taste-related information is present (Jezzini et al., 2013).

Hidden Markov Model (HMM) Analysis

Here we briefly outline the procedure used in Mazzucato et al. (2015), see this reference and (Jones et al., 2007; Escola et al., 2011; Ponce-Alvarez et al., 2012) for further details. Under the HMM, a system of N recorded neurons is assumed to be in one of a predetermined number of hidden (or latent) states (Rabiner, 1989; Zucchini and MacDonald, 2009). Each state m is defined as a vector of N firing rates $v_i(m)$, $i = 1, \dots, N$, one for each simultaneously recorded neuron. In each state, the neurons were assumed to discharge as stationary Poisson processes (Poisson-HMM). We matched the model to the data segmented in 1-ms bins (see below). In such short bins, we found that typically at most one spike was emitted across all simultaneously recorded neurons. If more than one neuron fired an action potential in a given bin, only one (randomly chosen) was kept for further analysis (this only occurred in a handful of bins per trial; Escola et al., 2011). We denote by $y_i(t)$ the spiking activity of the i -th neuron in the interval $[t, t + dt]$, $y_i(t) = 1$ if the neuron emitted a spike and $y_i(t) = 0$ otherwise. Denoting with S_t the hidden state of the ensemble at time t , the probability of having a spikes from neuron i in a given state m in the interval $[t, t + dt]$ is given by $p(y_i(t) = 1 | S_t = m) = 1 - e^{-v_i(m)dt}$.

The firing rates $v_i(m)$ completely define the states and are also called “emission probabilities” in HMM parlance. The emission and transition probabilities were found by maximization of the log-likelihood of the data given the model via the expectation-maximization (EM), or Baum-Welch, algorithm (Rabiner, 1989), a procedure known as “training the HMM.” For each session and type of activity (ongoing vs. evoked), ensemble spiking activity from all trials was binned at 1 ms intervals prior to training assuming a fixed number of hidden states M (Jones et al., 2007; Escola et al., 2011). For each given number of states M , the Baum-Welch algorithm was run 5 times, each time with random initial conditions for the transition and emission probabilities. The range of hidden states M for the HMM analyses were $M_{min} = 10$ and $M_{max} = 20$ for spontaneous activity, and $M_{min} = 10$ and $M_{max} = 40$ for evoked activity. Such numbers were based on extensive exploration of the parameter space and previous studies (Jones et al., 2007; Miller and Katz, 2010; Escola et al., 2011; Ponce-Alvarez et al., 2012; Mazzucato et al., 2015). For evoked activity, each HMM was trained on all four

tastes simultaneously. Of the models thus obtained, the one with largest total likelihood M^* was taken as the best HMM match to the data, and then used to estimate the probability of the states given the model and the observations in each bin of each trial (a procedure known as “decoding”). During decoding, only those hidden states with probability exceeding 80% in at least 50 consecutive bins were retained (henceforth denoted simply as “states”). State durations were approximately exponentially distributed with median duration 0.60 s (95% CIs: 0.07–4.70) during ongoing activity and 0.30 s (0.06–2.80) during evoked activity (Mazzucato et al., 2015).

The firing rate fits $v_i(m)$ in each trial were obtained from the analytical solution of the maximization step of the Baum-Welch algorithm,

$$v_i(m) = -\frac{1}{dt} \ln \left(1 - \frac{\sum_{t=1}^T r_m(t) y_i(t)}{\sum_{t=1}^T r_m(t)} \right). \quad (1)$$

Here, $[y_i(1), \dots, y_i(T)]$ is the spike train of the i -th neuron in the current trial, and T is the total duration of the trial. $r_m(t) = P(S_t = m | y(1), \dots, y(T))$ is the probability that the hidden state S_t at time t is m , given the observations.

Dimensionality Measure

We defined the dimensionality of the neural activity as

$$d = \frac{1}{\sum_{i=1}^N \tilde{\lambda}_i^2}, \quad (2)$$

where the $\tilde{\lambda}_i$ are the principal eigenvalues expressed as fractions of the total amount of variance explained, i.e., $\tilde{\lambda}_i = \lambda_i / (\sum_j \lambda_j)$, where λ_j are the eigenvalues of the covariance matrix of the firing rates (see below).

The dimensionality can be computed exactly in some relevant special cases. The calculation is simplified by the observation that Equation (2) is equivalent to

$$d = \frac{[Tr(C_f)]^2}{Tr(C_f^2)},$$

where C_f is the true covariance matrix of the firing rate vectors, $Tr(A) \equiv \sum_{i=1}^N A_{ii}$ is the trace of matrix A , and $Tr(A^2) = \sum_{i,j=1}^N A_{ij}A_{ji}$. We consider in the following only the case of firing rates in equal bins, hence we can replace C_f with the covariance matrix of the spike counts C in the definition of d :

$$d = \frac{[Tr(C)]^2}{Tr(C^2)} = \frac{b_N^2}{c_N + a_N}, \quad (3)$$

where for later convenience we have introduced the notation

$$a_N = \sum_{i=1}^N C_{ii}^2, \quad b_N = \sum_{i=1}^N C_{ii}, \quad c_N = \sum_{i \neq j} C_{ij}C_{ji}. \quad (4)$$

Note that d does not depend on the distribution of firing rates, but only on their covariance, up to a common scaling factor.

Dimensionality in the case of uniform pair-wise correlations

When all the pair-wise correlations r_{ij} are identical, $r_{ij} = \rho$ for all $i \neq j$,

$$r_{ij} = \begin{pmatrix} 1 & \rho & \dots & \rho \\ \rho & 1 & & \vdots \\ \vdots & & \ddots & \rho \\ \rho & \dots & \rho & 1 \end{pmatrix} \quad (5)$$

we have $C_{ij} = \rho\sqrt{\sigma_i^2\sigma_j^2}$ for $i \neq j$, where $\sigma_i^2 = C_{ii}$ is the spike count variance. In this case, we find from Equation (4) that

$$a_N = \sum_{i=1}^N \sigma_i^4, \quad b_N = \sum_{i=1}^N \sigma_i^2, \quad c_N = \rho^2(b_N^2 - a_N). \quad (6)$$

and the dimensionality, Equation (3), is given by

$$d = \frac{1}{\rho^2 + (1 - \rho^2)g(N)}, \quad (7)$$

where

$$g_N = \frac{a_N}{b_N^2} = \frac{\sum_{i=1}^N \sigma_i^4}{\left(\sum_{i=1}^N \sigma_i^2\right)^2},$$

Note that since both a_N and b_N scale as N when N is large, in general $g_N \sim \frac{1}{N}$ for large N .

If all spike counts have equal variance, $\sigma_i = \sigma$, we find exactly $g_N = \frac{1}{N}$:

$$d = \frac{1}{\rho^2 + \frac{(1 - \rho^2)}{N}} = \frac{N}{N\rho^2 + (1 - \rho^2)}, \quad (8)$$

and the dependence of d on the variance drops out. Note that for uncorrelated spike counts ($\rho = 0$) this formula gives $d = N$, whereas for any finite correlation we find the upper bound $d = 1/\rho^2$. For $N > 1$, the dimensionality is inversely related to the amount of pair-wise correlation ρ .

Consider the case where spike counts have variances σ_i^2 drawn from a probability distribution with mean $E[\sigma_i^2] = \sigma^2$ and variance $Var[\sigma_i^2] = \delta\sigma^4$, and the pair-wise correlation coefficients r_{ij} , for $i \neq j$, are drawn from a distribution with mean $E[r_{ij}] = \rho$ and variance $Var[r_{ij}] = \delta\rho^2$. In such a case one can evaluate Equation (3) approximately by its Taylor expansion around the mean values of the quantities in Equation (4). At leading order in N one finds

$$E[d] \approx \frac{E[b_N^2]}{E[c_N] + E[a_N]} = \frac{N\sigma^4 + \delta\sigma^4}{(N - 1)\sigma^4(\rho^2 + \delta\rho^2) + \sigma^4 + \delta\sigma^4}, \quad (9)$$

where $E[.]$ denotes expectation. To obtain this result we have used the definitions in Equation (4), from which

$$\begin{aligned} E[a_N] &= N(\sigma^4 + \delta\sigma^4), \quad E[b_N^2] = N^2\sigma^4 + N\delta\sigma^4, \\ E[c_N] &= (N^2 - N)\sigma^4(\rho^2 + \delta\rho^2), \end{aligned} \quad (10)$$

and the fact that, given a random vector X_i with mean μ_i and covariance C_{ij} , and a constant symmetric matrix A_{ij} , the expectation value of the quadratic form $\sum_{i,j} X_i A_{ij} X_j$ is

$$E\left[\sum_{i,j} X_i A_{ij} X_j\right] = \sum_{ij} (A_{ij} C_{ji} + \mu_i A_{ij} \mu_j). \quad (11)$$

In the case of uncorrelated spike counts ($\rho = 0, \delta\rho = 0$), dimensionality still depends linearly on the ensemble size N , but with a smaller slope $\frac{\sigma^4}{\sigma^4 + \delta\sigma^4} < 1$ compared to the case of equal variances (Equation 8 with $\rho = 0$).

Dimensionality in the Case of Neural Clusters

Given an ensemble of N neurons arranged in Q clusters (motivated by the model network described later in section ‘‘Spiking neuron model’’), we created ensembles of uncorrelated spike trains for $N \leq Q$ and correlated within each cluster for $N > Q$. Thus, if $N \leq Q$ the correlation matrix is the $N \times N$ identity matrix. If $N > Q$, the $(Q+1)$ th neuron was added to the first cluster, with correlation ρ with the other neuron of the cluster, and uncorrelated to the neurons in the remaining clusters. The $(Q+2)$ th neuron was added to the second cluster, with correlation ρ with the other neuron of the second cluster, and uncorrelated to the neurons in the remaining clusters, and so on. Similarly, the $(2Q+p)$ th neuron ($p \leq Q$) was added to the p -th cluster, with pair-wise correlation ρ with the other neurons of the same cluster, but no correlation with the neurons in the remaining clusters; and so on. In general, for $N = mQ + p$ neurons (where $m = \lfloor \frac{N}{Q} \rfloor \geq 1$ is the largest integer smaller than $\frac{N}{Q}$), the procedure picked $m + 1$ neurons per cluster for the first p cluster and m neurons per cluster for the remaining $Q - p$ clusters, with uniform pair-wise correlations ρ in the same cluster while neurons from different clusters were uncorrelated. The resulting correlation matrix r was block diagonal

$$r = \text{diag}(R_1, \dots, R_Q),$$

where each of the Q blocks contains the correlations of neurons from the same cluster. Inside each block R_i , the off-diagonal terms are equal to the uniform within-cluster correlation ρ :

$$R_i = \begin{pmatrix} 1 & \rho & \dots & \rho \\ \rho & 1 & & \vdots \\ \vdots & & \ddots & \rho \\ \rho & \dots & \rho & 1 \end{pmatrix}$$

The first p blocks have size $(m + 1) \times (m + 1)$ and the last $Q - p$ blocks have size $m \times m$, so that $(m + 1)p + m(Q - p) = N$. The remaining elements of matrix r (representing pair-wise correlations of neurons belonging to different clusters) were all zero. Recalling that $C_{ij} = r_{ij}\sigma_i\sigma_j$, one finds $Tr(C) = pb_{m+1} + (Q - p)b_m$ and $Tr(C^2) = \rho^2[pb_{m+1}^2 + (Q - p)b_m^2] +$

$(1 - \rho^2)[pa_{m+1} + (Q - p)a_m]$, where a_n and b_n are defined in Equation (6), from which one obtains

$$d = \begin{cases} b_N^2/a_N & , N \leq Q \\ \frac{[pb_{m+1} + (Q-p)b_m]^2}{\rho^2[pb_{m+1}^2 + (Q-p)b_m^2] + (1-\rho^2)[pa_{m+1} + (Q-p)a_m]} & , N > Q \end{cases} \quad (12)$$

In the approximation where all neurons have the same variance this simplifies to

$$d = \begin{cases} N & , N \leq Q \\ \frac{N}{1+m\rho^2[1-(Q-p)/N]} & , N > Q. \end{cases} \quad (13)$$

Recall that in the formulae above m and p depend on N . For finite ρ , Equation (13) predicts the bound $d \leq Q/\rho^2$ for any $N > 1$, with this value reached asymptotically for large N . When single neuron variances σ_i^2 are drawn from a distribution with mean $E[\sigma_i^2] = \sigma^2$ and variance $Var[\sigma_i^2] = \delta\sigma^4$, an expression for the dimensionality can be obtained from Equation (12) at leading order in the expectation values of the quantities in Equation (4) (not shown), with a procedure similar to that used to obtain Equation (9).

Pair-Wise Correlations

Given neuron i and neuron j 's spike trains, we computed the spike count correlation coefficient r_{ij}

$$r_{ij} = \frac{S_{ij}}{\sqrt{S_{ii}S_{jj}}},$$

where S is the sample covariance matrix of the spike counts estimated as

$$S_{ij} = \frac{1}{N_b N_T - 1} \sum_{b,s=1}^{N_b, N_T} (n_i(b, s) - \langle n_i \rangle) (n_j(b, s) - \langle n_j \rangle), \quad (14)$$

where $n_i(b, s)$ is the spike count of neuron i in bin b and trial s . The sum goes over all N_b bins and over all N_T trials in a session, whereas $\langle n_i \rangle$ is the average across trials and bins for neuron i . In the main text and figures we present results obtained with a bin size of 200 ms, but have performed the same analyses with bin sizes varying from 10 ms to 5 s (see Results for details).

Significance of the correlation was estimated as follows (Renart et al., 2010): $N_{shuffle} = 200$ trial-shuffled correlation coefficients r'_{ij} were computed, then a p -value was determined as the fraction of shuffled coefficients r'_{ij} whose absolute value exceeded the absolute value of the experimental correlation, $p = \frac{\#(|r'_{ij}| > |r_{ij}|)}{N_{shuffle}}$. For example, a correlation r was significant at $p = 0.05$ confidence level if no more than 10 shuffled correlation coefficients out of 200 exceeded r .

The pair-wise correlations of firing rates vectors computed in bins of fixed duration T were given by Equation (14) with $n_i(b, s)$ replaced by $n_i(b, s)/T$. Instead, correlations of firing rates vectors inside hidden states (which have variable duration) were estimated after replacing $n_i(b, s)$ in Equation (14) with $v_i(m, s)$, the firing rate of neuron i in state m in trial s . For each trial s , this quantity was computed according to Equation (1).

Estimation of Dimensionality

The eigenvalues λ_j in Equation (2) were found with a standard Principal Component Analysis (PCA) of the set of all firing rate vectors (Chapin and Nicolelis, 1999). The firing rate vectors were obtained via the HMM analysis (see Equation 1); all data from either ongoing or evoked activity were used. For the analysis of **Figure 3E**, where the duration and number of trials were varied, only the firing rate vectors of the HMM states present in the given trial snippet were used (even if present for only a few ms). When firing rate vectors in hidden states were not available (mainly, in “shuffled” datasets and in asynchronous homogeneous networks, see below for details), the firing rates were computed as spike counts in $T = 200$ ms bins divided by T , $n_i(b, s)/T$, where $n_i(b, s)$ is as defined in Equation (14) (**Figures 3E, G, 6E, 7D, 9A**). Dimensionality values were averaged across 20 simulated sessions for each ensemble size N ; in each session, 40 trials of 5 s duration, resulting in $N_T = 1,000$ bins, were used (using bin widths of 50–500 ms did not change the results). Note that for the purpose of computing the dimensionality (Equation 3), it is equivalent to use either the binned firing rate $n_i(b, s)/T$ or the spike count $n_i(b, s)$.

In our data, d roughly corresponded to the number of principal components explaining between 80 and 90% of the variance. However, note that all eigenvalues are retained in our definition of dimensionality given in Equation (2) above.

Shuffled Datasets

The dimensionality of the data as a function of ensemble size N was validated against surrogate datasets constructed by shuffling neurons across different sessions while matching the empirical distribution of ensemble sizes. Comparison analyses between empirical and shuffled ensembles were trial-matched using the minimal number of trials per condition across ensembles, and then tested for significant difference with the Mann-Whitney test on samples obtained from 20 bootstrapped ensembles. Neurons whose firing rate variance exceeded the population average by two standard deviations were excluded (8/167 of non-silent, non-somatosensory neurons).

Dependence on the Number of Trials: Simulations (Figures 7E, 8A)

The estimate of d from data depends on the number and duration of the trials (**Figure 3E** and Equation 16 below). To investigate this phenomenon in a simple numerical setting we generated $N \times N_T$ “nominal” firing rates, thought of as originating from N neurons, each sampled N_T times (trials). The single firing rates were sampled according to a log-normal distribution with equal means and covariance leading to Equation (7), i.e., $C_{ij} = \rho\sigma_i\sigma_j(1 - \delta_{ij}) + \sigma_i^2\delta_{ij}$, with $\delta_{ij} = 1$ if $i = j$, and zero otherwise (note that the actual distribution used is immaterial since the dimensionality only depends on the covariance matrix, see Equation 3). We considered the two cases of equal variance for all ensemble neurons, $\sigma_i = \sigma$ for all i (**Figure 8A**) or variances σ_i sampled from a log-normal distribution (**Figure 8A** and “+” in **Figure 7E**). The same N and N_T as used for the analysis of the model simulations in **Figure 7D** were used (where the “trials” were N_T bins of 200 ms in 40 intervals of 5 second

duration for each ensemble size N). The covariance of the data thus generated was estimated according to Equation (14), based on which the dimensionality Equation (3) was computed. The estimated dimensionality depends on N and N_T and was averaged across 100 values of d , each obtained as explained above. Note that in this simplified setting increasing the duration of each trial is equivalent to adding more trials, i.e., the effect of having a trial 400 ms long producing 2 firing rates (one for each 200 ms bin) is equivalent to having two trials of 200 ms duration. In the general case, the effect of trial duration on d will depend on how trial duration affects the variance and correlations of the firing rates.

Dependence on the Number of Trials: Theory

The dependence of dimensionality on the number of trials can be computed analytically under the assumption that N ensemble neurons generate spike counts n_i , for $i = 1, \dots, N$, distributed according to a multivariate Gaussian. Since we are interested in the spike-count covariance Equation (14), we can assume the spike-count distribution to have zero mean and true covariance C_{ij} . The matrix $M^{(N_T)} = (N_T - 1) \cdot S^{(N_T)}$, where $S^{(N_T)}$ is the covariance matrix Equation (14) sampled from N_T trials, is distributed according to a Wishart distribution $W_N(C_{ij}, N_T - 1)$ with $N_T - 1$ degrees of freedom (Mardia et al., 1979). Since the variance of the Wishart distribution,

$$\text{Var}(M_{ij}) = (N_T - 1)(C_{ij}^2 + C_{ii}C_{jj}),$$

is proportional to N_T , we obtain the variance of the entries of the sample covariance as

$$\text{Var}(S_{ij}^{(N_T)}) = \frac{C_{ij}^2 + C_{ii}C_{jj}}{N_T - 1}, \quad (15)$$

to be used in the estimator of d (from Equation 3)

$$\hat{d} = \frac{[\text{Tr}(S)]^2}{\text{Tr}(S^2)} = \frac{\widehat{b_N^2}}{\widehat{c_N} + \widehat{a_N}},$$

where $\widehat{a_N}$, $\widehat{c_N}$, $\widehat{b_N^2}$ are given by Equation (4) with C replaced by S . With a calculation similar to that used to obtain Equation (9), to leading order in N and N_T one finds

$$E[\hat{d}] \approx \frac{E[\widehat{b_N^2}]}{E[\widehat{c_N}] + E[\widehat{a_N}]},$$

with

$$E[\widehat{a_N}] = N(\sigma^4 + \delta\sigma^4) + \frac{2N\sigma^4}{N_T - 1},$$

$$E[\widehat{b_N^2}] = N^2\sigma^4 + N\delta\sigma^4 + \frac{2N\sigma^4}{N_T - 1},$$

$$E[\widehat{c_N}] = (N^2 - N)(\rho^2 + \delta\rho^2)\sigma^4 + (N^2 - N)\frac{1 + \rho^2 + \delta\rho^2}{N_T - 1}\sigma^4,$$

where we also used Equations (10) and (11), with $\text{Var}[\sigma_i^2] = \delta\sigma^4$ and $\text{Var}[r_{ij}] = \delta\rho^2$, for $i \neq j$. In conclusion, one finds

$$E[\hat{d}] = \frac{\left(N^2 + \frac{2N}{N_T - 1}\right)\sigma^4 + N\delta\sigma^4}{(N^2 - N)\left(\rho^2 + \delta\rho^2 + \frac{1 + \rho^2 + \delta\rho^2}{N_T - 1}\right)\sigma^4 + N\left(1 + \frac{2}{N_T - 1}\right)\sigma^4 + N\delta\sigma^4}. \quad (16)$$

Model Fitting

The dependence of the data's dimensionality on ensemble size N was fitted by a straight line via standard least-squares,

$$d = \beta_1 \cdot N + \beta_0,$$

separately for ongoing and evoked activity (**Figures 3B–D, 6B–D**). Comparison between the dimensionality of evoked and ongoing activity was carried out with a 2-way ANOVA with condition (evoked vs. ongoing) and ensemble size (N) as factors. Since d depends on the number and duration of the trials used to estimate the covariance matrix (**Figure 3E** and Equation 16), we matched both the number of trials and trial length in comparisons of ongoing and evoked dimensionality. If multiple tastes were used, the evoked trials were each matched to a random subset of an equal number of ongoing trials.

The dependence of dimensionality d on ensemble size N in a surrogate dataset of Poisson spike trains with mean pairwise correlation ρ (generated according to the algorithm described in the next section) was modeled as Equation (16) with $\delta\rho^2 = \alpha\rho^2$ and $\delta\sigma^4 = \sigma^4 = \beta$ (**Figure 7D**, dashed lines); N_T was fixed to 1000 (40 trials of 5 s each, segmented in 200 ms bins). The parameters α , β were tuned to fit all Poisson trains simultaneously on datasets with $N = 5, 10, \dots, 100$ and $\rho = 0, 0.01, 0.05, 0.1, 0.2$, with 20 ensembles for each value (**Figure 7D**; only the fits for $\rho = 0, 0.1, 0.2$ are shown). A standard non-linear least-squares procedure was used (Holland and Welsch, 1977).

Generation of Correlated Poisson Spike Trains

Ensembles of independent and correlated Poisson spike trains were generated for the analysis of **Figure 7**. Ensembles of independent stationary Poisson spike trains with given firing rates v_i were generated by producing their interspike intervals according to an exponential distribution with parameter v_i . Stationary Poisson spike trains with fixed pairwise correlations (but no temporal correlations) were generated according to the method reported in Macke et al. (2009), that we briefly outline below.

We split each trial into 1 ms bins and consider the associated binary random variable $X_i(t) = 1$ if the i -th neuron emitted a spike in the t -th bin, and $X_i(t) = 0$ if no spike was emitted. These samples were obtained by first drawing a sample from an auxiliary N -dimensional Gaussian random variable $\mathcal{U} \sim \mathcal{N}(\gamma, \Lambda)$ and then thresholding it into 0 and 1: $X_i = 1$ if $\mathcal{U}_i > 0$, and $X_i = 0$ otherwise. Here, $\gamma = \{\gamma_1, \gamma_2, \dots, \gamma_N\}$ is the mean vector and $\Lambda = \{\Lambda_{ij}\}$ is the covariance matrix of the N -dimensional Gaussian variable \mathcal{U} . For appropriately chosen

parameters γ_i and Λ_{ij} the method generates correlated spike trains with the desired firing rates v_i and pairwise spike count correlation coefficients r_{ij} .

The prescription for γ_i and Λ_{ij} is most easily expressed as a function of the desired probabilities μ_i of having a spike in a bin of width dt , $\mu_i = P(X_i(t) = 1)$, and the pairwise covariance c_{ij} of the random binary vectors $X_i(t)$ and $X_j(t)$, from which γ_i and Λ_{ij} can be obtained by inverting the following relationships:

$$\begin{aligned}\mu_i &= \Phi(\gamma_i), \\ c_{ii} &= \Phi(\gamma_i)\Phi(-\gamma_i), \\ c_{ij} &= \Phi_2(\gamma_i, \gamma_j, \Lambda_{ij}) - \Phi(\gamma_i)\Phi(\gamma_j), \quad i \neq j.\end{aligned}$$

Here, $\Phi(x)$ is the cumulative distribution of a univariate Gaussian with mean 0 and variance 1 evaluated at x , and $\Phi_2(x, y, \Lambda)$ is the cumulative distribution of a bivariate Gaussian with means 0, variances 1 and covariance Λ evaluated at (x, y) (note that the distributions Φ and Φ_2 are unrelated to the N -dimensional Gaussian $\mathcal{U} \sim \mathcal{N}(\gamma, \Lambda)$). Without loss of generality we imposed unit variances for \mathcal{U}_i , i.e., $\Lambda_{ii} = 1$.

We related the spike probabilities μ_i to the firing rates v_i as $\mu_i = 1 - e^{-v_i dt}$, with $(1 - \mu_i)$ being the probability of no spikes in the same bin. When dt approaches zero, $\mu_i \approx v_i dt$ and the spike trains generated as vectors of binary random variables by sampling $\mathcal{U} \sim \mathcal{N}(\gamma, \Lambda)$ will approximate Poisson spike trains ($dt = 1$ ms bins were used). In order to have a fair comparison with the data generated by the spiking network model (described in the next section), the mean firing rates of the Poisson spike trains were matched to the average firing rates obtained from the simulated data.

Since γ and Λ were the same in all bins, values of $X_i(t)$ and $X_i(s)$ were independent for $t \neq s$ (i.e., the spike trains had no temporal correlations). As a consequence, the random binary vectors have the same pair-wise correlations as the spike counts, and the c_{ij} are related to the desired r_{ij} by $c_{ij} = r_{ij} \sqrt{\mu_i(1 - \mu_i)\mu_j(1 - \mu_j)}$, where $\mu_i(1 - \mu_i)$ is the variance of X_i . See Macke et al. (2009) for further details.

Spiking Network Model

We modeled the data with a recurrent spiking network of $N = 5000$ randomly connected leaky integrate-and-fire (LIF) neurons, of which 4000 excitatory (E) and 1000 inhibitory (I). Connection probability $p_{\beta\alpha}$ from neurons in population $\alpha \in E, I$ to neurons in population $\beta \in E, I$ were $p_{EE} = 0.2$ and $p_{EI} = p_{IE} = p_{II} = 0.5$; a fraction $f = 0.9$ of excitatory neurons were arranged into Q different clusters, with the remaining neurons belonging to an unstructured (“background”) population (Amit and Brunel, 1997). Synaptic weights $J_{\beta\alpha}$ from neurons in population $\alpha \in E, I$ to neurons in population $\beta \in E, I$ scaled with N as $J_{\beta\alpha} = j_{\beta\alpha}/\sqrt{N}$, with $j_{\beta\alpha}$ constants having the following values (units of mV): $j_{EI} = 3.18$, $j_{IE} = 1.06$, $j_{II} = 4.24$, $j_{EE} = 1.77$. Within an excitatory cluster synaptic weights were potentiated, i.e., they took average values of $\langle J \rangle_+ = J_+ j_{EE}$ with $J_+ > 1$, while synaptic weights between units belonging to different clusters were depressed to average values $\langle J \rangle_- = J_- j_{EE}$, with $J_- = 1 - \gamma f (J_+ - 1) < 1$, with $\gamma = 0.5$. The latter relationship between

J_+ and J_- helps to maintain balance between overall potentiation and depression in the network (Amit and Brunel, 1997).

Below spike threshold, the membrane potential V of each LIF neuron evolved according to

$$\tau_m \frac{dV}{dt} = -V + \tau_m (I_{rec} + I_{ext} + I_{stim})$$

with a membrane time constant $\tau_m = 20$ ms for excitatory and 10 ms for inhibitory units. The input current was the sum of a recurrent input I_{rec} , an external current I_{ext} representing an ongoing afferent input from other areas, and an external stimulus I_{stim} representing e.g., a delivered taste during evoked activity only. In our units, a membrane capacitance of 1nF is set to 1. A spike was said to be emitted when V crossed a threshold V_{thr} , after which V was reset to a potential $V_{reset} = 0$ for a refractory period of $\tau_{ref} = 5$ ms. Spike thresholds were chosen so that, in the unstructured network (i.e., with $J_+ = J_- = 1$), the E and I populations had average firing rates of 3 and 5 spikes/s, respectively (Amit and Brunel, 1997). The recurrent synaptic input I_{rec}^i to unit i evolved according to the dynamical equation

$$\tau_s \frac{dI_{rec}^i}{dt} = -I_{rec}^i + \sum_{j=1}^N J_{ij} \sum_k \delta(t - t_k^j),$$

where t_k^j was the arrival time of k -th spike from the j -th pre-synaptic unit, and τ_s was the synaptic time constant (3 and 2 ms for E and I units, respectively), resulting in an exponential post-synaptic current in response to a single spike, $\frac{J_{ij}}{\tau_s} \exp(-t/\tau_s) \Theta(t)$, where $\Theta(t) = 1$ for $t \geq 0$, and $\Theta(t) = 0$ otherwise. The ongoing external current to a neuron in population α was constant and given by $I_{ext} = N_{ext} p_{\alpha 0} J_{\alpha 0} v_{ext}$, where $N_{ext} = n_E N$, $p_{\alpha 0} = p_{EE}$, $J_{\alpha 0} = \frac{j_{\alpha 0}}{\sqrt{N}}$ with $j_{E0} = 0.3$, $j_{I0} = 0.1$, and $v_{ext} = 7$ spikes/s. During evoked activity, stimulus-selective units received an additional input representing one of the four incoming stimuli. The stimuli targeted combinations of neurons as observed in the data. Specifically, the fractions of neurons responsive to $n = 1, 2, 3$ or all 4 stimuli were 17% (27/162), 22% (36/162), 26% (42/162), and 35% (57/162) (Jezzini et al., 2013; Mazzucato et al., 2015). Each stimulus had constant amplitude v_{stim} ranging from 0 to 0.5 v_{ext} . In the following we measure the stimulus amplitude as percentage of v_{ext} (e.g., “10%” corresponds to $v_{stim} = 0.1 v_{ext}$). The onset of each stimulus was always $t = 0$, the time of taste delivery. The stimulus current to a unit in population α was constant and given by $I_{stim} = N_{ext} p_{\alpha 0} J_{\alpha 0} v_{stim}$.

Mean Field Analysis of the Model

The stationary states of the spiking network model in the limit of large N were found with a mean field analysis (Amit and Brunel, 1997; Brunel and Hakim, 1999; Fusi and Mattia, 1999; Curti et al., 2004; Mazzucato et al., 2015). Under typical conditions, each neuron of the network receives a large number of small post-synaptic currents (PSCs) per integration time constant. In such a case, the dynamics of the network can be analyzed under the diffusion approximation within the population density approach. The network has $\alpha = 1, \dots, Q + 2$ sub-populations, where the

first Q indices label the Q excitatory clusters, $\alpha = Q + 1$ labels the “background” units, and $\alpha = Q + 2$ labels the homogeneous inhibitory population. In the diffusion approximation (Tuckwell, 1988; Lánský and Sato, 1999; Richardson, 2004), the input to each neuron is completely characterized by the infinitesimal mean μ_α and variance σ_α^2 of the post-synaptic potential (see Mazzucato et al., 2015 for the expressions of the infinitesimal mean and variance for all subpopulations).

Parameters were chosen so that the network with $J_+ = J_- = 1$ (where all $E \rightarrow E$ synaptic weights are equal) would operate in the balanced asynchronous regime (van Vreeswijk and Sompolinsky, 1996, 1998; Renart et al., 2010), where incoming contributions from excitatory and inhibitory inputs balance out, neurons fire irregular spike trains with weak pair-wise correlations.

The unstructured network has only one dynamical state, i.e., a stationary point of activity where all E and I neurons have constant firing rate v_E and v_I , respectively. In the structured network (where $J_+ > 1$), the network undergoes continuous transitions among a repertoire of states, as shown in the main text. To avoid confusion between network activity states and HMM states, we refer to the former as network “configurations” instead of states. Admissible networks configurations must satisfy the $Q + 2$ self-consistent mean field equations (Amit and Brunel, 1997)

$$v_\alpha = F_\alpha(\mu_\alpha(\vec{v}), \sigma_\alpha^2(\vec{v})),$$

where $\vec{v} = [v_1, \dots, v_Q, v_E^{bg}, v_I]$ is the firing rate vector and $F_\alpha(\mu_\alpha, \sigma_\alpha^2)$ is the current-to-rate response function of the LIF neurons. For fast synaptic times, i.e., $\frac{\tau_s}{\tau_m} \ll 1$, $F_\alpha(\mu_\alpha, \sigma_\alpha^2)$ is well approximated by (Brunel and Sergi, 1998; Fourcaud and Brunel, 2002)

$$F_\alpha(\mu_\alpha, \sigma_\alpha) = \left(\tau_{ref} + \tau_{m,\alpha} \sqrt{\pi} \int_{H_{eff,\alpha}}^{\Theta_{eff,\alpha}} e^{u^2} [1 + \text{erf}(u)] \right)^{-1},$$

where

$$\Theta_{eff,\alpha} = \frac{V_{thr,\alpha} - \mu_\alpha}{\sigma_\alpha} + ak_\alpha,$$

$$H_{eff,\alpha} = \frac{V_{reset,\alpha} - \mu_\alpha}{\sigma_\alpha} + ak_\alpha,$$

where $k_\alpha = \sqrt{\tau_{s,\alpha}/\tau_{m,\alpha}}$ is the square root of the ratio of synaptic time constant to membrane time constant, and $a = \frac{|\xi(1/2)|}{\sqrt{2}} \sim 1.03$. This theoretical response function has been fitted successfully to the firing rate of neocortical neurons in the presence of *in vivo*-like fluctuations (Rauch et al., 2003; Giugliano et al., 2004; La Camera et al., 2006, 2008).

The fixed points \vec{v}^* of the mean field equations were found with Newton’s method (Press et al., 2007). The fixed points

can be either stable (attractors) or unstable depending on the eigenvalues λ_α of the stability matrix

$$S_{\alpha\beta} = \frac{1}{\tau_{s,\alpha}} \left(\frac{\partial F_\alpha(\mu_\alpha(\vec{v}), \sigma_\alpha^2(\vec{v}))}{\partial v_\beta} \right) - \left(\frac{\partial F_\alpha(\mu_\alpha(\vec{v}), \sigma_\alpha^2(\vec{v}))}{\partial \sigma_\alpha^2} \frac{\partial \sigma_\alpha^2}{\partial v_\beta} - \delta_{\alpha\beta} \right),$$

evaluated at the fixed point \vec{v}^* (Mascaro and Amit, 1999). If all eigenvalues have negative real part, the fixed point is stable (attractor). If at least one eigenvalue has positive real part, the fixed point is unstable. Stability is meant with respect to an approximate linearized dynamics of the mean and variance of the input current:

$$\tau_{s,\alpha} \frac{dm_\alpha}{dt} = -m_\alpha + \mu_\alpha(\vec{v})$$

$$\frac{\tau_{s,\alpha}}{2} \frac{ds_\alpha^2}{dt} = -s_\alpha^2 + \sigma_\alpha^2(\vec{v})$$

$$v_\alpha(t) = F_\alpha(m_\alpha(\vec{v}), s_\alpha^2(\vec{v})),$$

where μ_α and σ_α^2 are the stationary values for fixed \vec{v} given earlier. For fast synaptic dynamics in the asynchronous balanced regime, these rate dynamics are in very good agreement with simulations (La Camera et al., 2004—see Renart et al., 2004; Giugliano et al., 2008 for more detailed discussions).

Metastable Configurations in the Network Model

The stable configurations of a network with an infinite number of neurons were obtained in the mean field approximation of the previous section and are shown in **Figure 4B** for $Q = 30$ and a range of values of the relative potentiation parameter J_+ . Above the critical point $J_+ = 4.2$, stable configurations characterized by a finite number of active clusters emerge (gray lines; the number of active clusters is reported next to each line). For a given J_+ , the firing rate is the same in all active clusters and is inversely proportional to the total number of active clusters. Stable patterns of firing rates are also found in the inhibitory population (red lines), in the inactive clusters (having low firing rates; gray dashed lines), and in the unstructured excitatory population (dashed blue lines). For a fixed value of J_+ , multiple stable configurations coexist with different numbers of active clusters. For example, for $J_+ = 5.3$, configurations with up to 7 active clusters are stable, each configuration with different firing rates. This generates multistable firing rates in single neurons, i.e., the property, also observed in the data, that single neurons can attain more than 2 firing rates across states (Mazzucato et al., 2015). Note that if $J_+ \leq 5.15$ an alternative stable configuration of the network with all clusters inactive (firing rates < 10 spikes/s) is also possible (single brown line).

Strictly speaking, the configurations in **Figure 4B** are stable only in a network containing an infinite number of uncorrelated neurons. In a finite network (or when neurons are strongly correlated) these configurations can lose stability due to strong fluctuations, which ignite transitions among the different configurations. Full details are reported in Mazzucato et al. (2015).

Model Simulations and Analysis of Simulated Data

The dynamical equations of the LIF neurons were integrated with the Euler algorithm with a time step of $dt = 0.1$ ms. We simulated 20 different networks (referred to as “sessions” in the following) during both ongoing and evoked activity. We chose four different stimuli per session during evoked activity (to mimic taste delivery). Trials were 5 s long. The HMM analyses for **Figures 2, 5** were performed on ensembles of randomly selected excitatory neurons with the same procedure used for the data (see previous section “Hidden Markov Model (HMM) analysis”). The ensemble sizes were chosen so as to match the empirical ensemble sizes (3–9 randomly selected neurons). For the analysis of **Figure 9A**, ensembles of increasing size (from 5 to 100 neurons) were used from simulations with $Q = 30$ clusters. When the ensemble size was less than the number of clusters ($N \leq Q$), each neuron was selected randomly from a different cluster; when ensemble size was larger than the number of clusters, one neuron was added to each cluster until all clusters were represented, and so on until all N neurons had been chosen. To allow comparison with surrogate Poisson spike trains, the dimensionality of the simulated data was computed from the firing rate vectors in $T = 200$ ms bins as explained in section “Dimensionality measure.” For control, the dimensionality was also computed from the firing rate vectors in hidden states obtained from an HMM analysis, obtaining qualitatively similar results.

RESULTS

Dimensionality of the Neural Activity

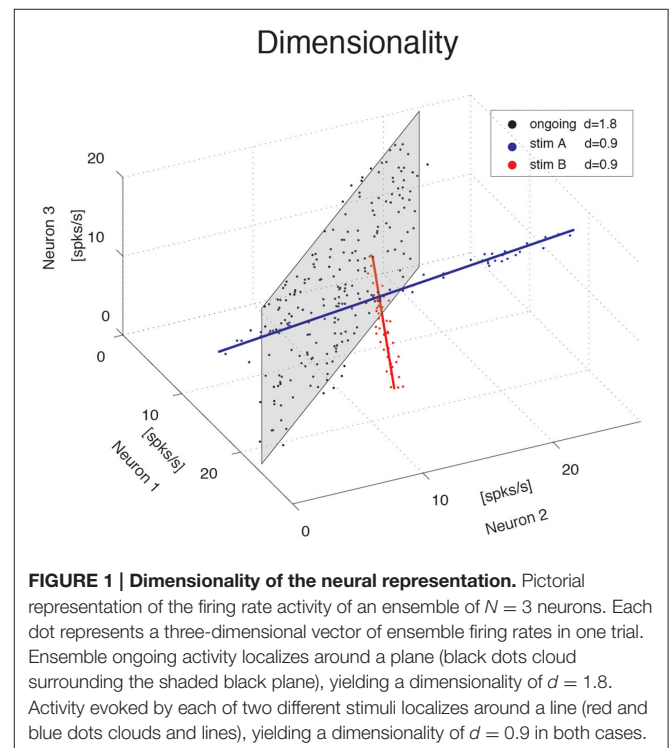
We investigate the dimensionality of sequences of firing rate vectors generated in the GC of alert rats during periods of ongoing or evoked activity (see Methods). To provide an intuitive picture of the meaning of dimensionality adopted in this paper, consider the firing rate vectors from N simultaneously recorded neurons. These vectors can occupy, a priori, the entire N -dimensional vector space minimally required to describe the population activity of N independent neurons. However, the sequence of firing rate vectors generated by the neural dynamics may occupy a subspace that is spanned by a smaller number $m < N$ of coordinate axes. For example, the data obtained by the ensemble of three simulated spike counts in **Figure 1** mostly lie on a 2D space, the plane shaded in gray. Although 3 coordinates are still required to specify all data points, a reduced representation of the data, such as that obtained from PCA, would quantify the dimension of the relevant subspace as being close to 2. To quantify this fact we use the following definition of

dimensionality (Abbott et al., 2011)

$$d = \left(\sum_{i=1}^N \tilde{\lambda}_i^2 \right)^{-1},$$

where N is the ensemble size and $\tilde{\lambda}_i$ are the normalized eigenvalues of the covariance matrix, each expressing the fraction of the variance explained by the corresponding principal component (see Methods for details). According to this formula, if the first n eigenvalues express each a fraction $1/n$ of the variance while the remaining eigenvalues vanish, the dimensionality is $d = n$. In less symmetric situations, d reflects roughly the dimension of the linear subspace explaining most variance about all data points. In the example of the data on the gray plane of **Figure 1**, $d = 1.8$, which is close to 2, as expected. Similarly, data points lying mostly along the blue and red straight lines in **Figure 1** have a dimensionality of 0.9, close to 1. In all cases, $d > 0$ and $d \leq N$, where N is the ensemble size.

The blue and red data points in **Figure 1** were obtained from a fictitious scenario where neuron 1 and neuron 2 were selective to surrogate stimuli A and B, respectively, and are meant to mimic two possible evoked responses. The subspace containing responses to both stimuli A and B would have a dimensionality $d_{A+B} = 1.7$, similar to the dimensionality of the data points distributed on the gray plane (meant instead to represent spike counts during ongoing activity in the same fictitious scenario). Thus, a dimensionality close to 2 could originate from different patterns of activity, such as occupying a plane or two straight lines. Other and more complex scenarios are, of course, possible. In general, the dimensionality



will reflect existing functional relationships among ensemble neurons (such as pair-wise correlations) as well as the response properties of the same neurons to external stimuli. The pictorial example of **Figure 1** caricatures a stimulus-induced reduction of dimensionality, as found in the activity of simultaneously recorded neurons from the GC of alert rats, as we show next.

Dimensionality is Proportional to Ensemble Size

We computed the dimensionality of the neural activity of ensembles of 3–9 simultaneously recorded neurons in the gustatory cortex of alert rats during the 5 s inter-trial period preceding (ongoing activity) and following (evoked activity) the delivery of a taste stimulus (said to occur at time $t = 0$; see Methods). Ensemble activity in single trials during both ongoing (**Figure 2A**) and evoked activity (**Figure 2B**) could be characterized in terms of sequences of metastable states, where each state is defined as a collection of firing rates across simultaneously recorded neurons (Jones et al., 2007; Mazzucato et al., 2015). Transitions between consecutive states were detected via a Hidden Markov Model (HMM) analysis, which provides the probability that the network is in a certain state at every 1 ms bin (**Figure 2**, color-coded lines superimposed to raster plots). The ensemble of spike trains was considered to be in a given state if the posterior probability of being in that state exceeded 80% in at least 50 consecutive 1-ms bins (**Figure 2**, color-coded shaded areas). Transitions among states were triggered by the co-modulation of a variable number of ensemble neurons and occurred at seemingly random times (Mazzucato et al., 2015). For this reason, the dimensionality of the neural activity was computed based on the firing rate vectors in each HMM state (one firing rate vector per state per trial; see Methods for details).

The average dimensionality of ongoing activity across sessions was $d_{\text{ongoing}} = 2.6 \pm 1.2$ (mean \pm SD; range: [1.2, 5.0]; 27 sessions). An example of the eigenvalues for a representative ensemble of eight neurons is shown in **Figure 3A**, where $d = 4.42$. The dimensionality of ongoing activity was approximately linearly related to ensemble size (**Figure 3B**, linear regression, $r = 0.4$, slope $b_{\text{ongoing}} = 0.26 \pm 0.12$, $p = 0.04$). During evoked activity dimensionality did not differ across stimuli (one-way ANOVA, no significant difference across tastants, $p > 0.8$), hence all evoked data points were combined for further analysis. An example of the eigenvalue distribution of the ensemble in **Figure 2B** is shown in **Figure 3C**, where $d_{\text{evoked}} = 1.3 \sim 1.7$ across 4 different taste stimuli. Across all sessions, dimensionality was overall smaller ($d_{\text{evoked}} = 2.0 \pm 0.6$, mean \pm SD, range: [1.1, 3.9]) and had a reduced slope as a function of N compared to ongoing activity (**Figure 3D**, linear regression, $r = 0.39$, slope $b_{\text{evoked}} = 0.13 \pm 0.03$, $p < 10^{-4}$). However, since dimensionality depends on the number and duration of the trials used for its estimation (**Figure 3E**), a proper comparison requires matching trial number and duration for each data point, as described next.

Stimulus-Induced Reduction of Dimensionality

We matched the number and duration of the trials for each data point and ran a two-way ANOVA with condition (ongoing

vs. evoked) and ensemble size as factors. Both the main dimensionality [$F_{(1, 202)} = 11.93$, $p < 0.001$] and the slope were significantly smaller during evoked activity [test of interaction, $F_{(6, 202)} = 5.09$, $p < 10^{-4}$]. There was also a significant effect of ensemble size [$F_{(6, 202)} = 18.72$, $p < 10^{-14}$], confirming the results obtained with the separate regression analyses. These results suggest that stimuli induce a reduction of the effective space visited by the firing rate vector during evoked activity. This was confirmed by a paired sample analysis of the individual dimensionalities across all $27 \times 4 = 108$ ensembles (27 ensemble times 4 gustatory stimuli; $p < 0.002$, Wilcoxon signed-rank test).

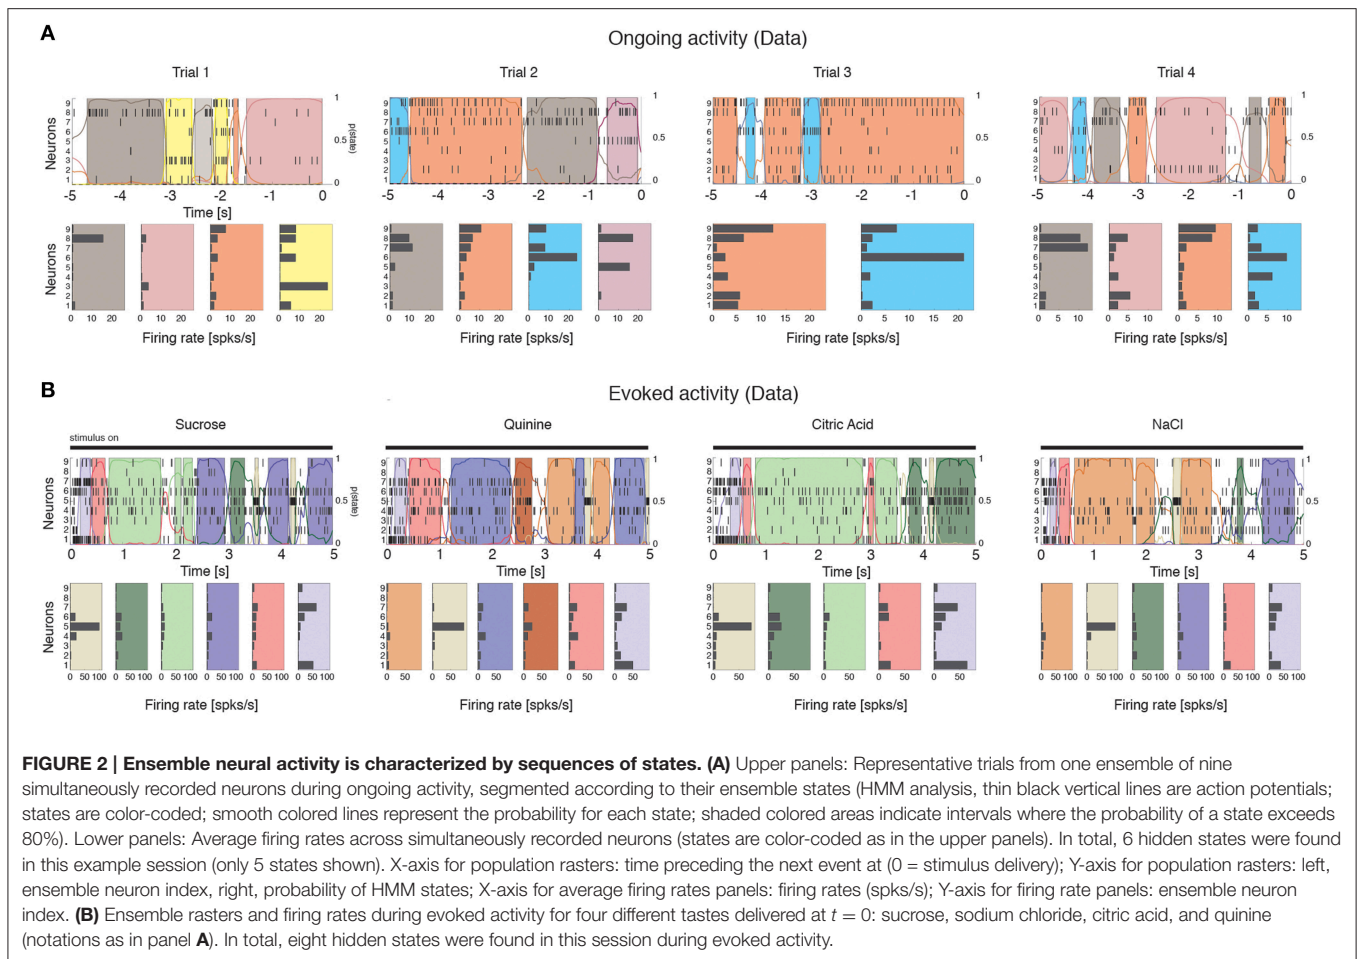
Dimensionality is larger in Ensembles of Independent Neurons

The dimensionality depends on the pair-wise correlations of simultaneously recorded neurons. Shuffling neurons across ensembles would destroy the correlations (beyond those expected by chance), and would give a measure of how different the dimensionality of our datasets would be compared to sets of independent neurons. We measured the dimensionality of surrogate datasets obtained by shuffling neurons across sessions; because shuffling destroys the structure of the hidden states, firing rates in bins of fixed duration (200 ms) were used to estimate the dimensionality (see Methods for details). As expected, the slope of d vs. N was larger in the shuffled datasets compared to the simultaneously recorded ensembles (not shown) during both ongoing activity ($b_{\text{shuff}} = 0.67 \pm 0.06$ vs. $b_{\text{data}} = 0.60 \pm 0.01$; mean \pm SD, Mann-Whitney test, $p < 0.001$, 20 bootstraps), and evoked activity ($b_{\text{shuff}} = 0.36 \pm 0.07$ vs. $b_{\text{data}} = 0.29 \pm 0.01$; $p < 0.001$). Especially during ongoing activity, this result was accompanied by a narrower distribution of pair-wise correlations in the shuffled datasets compared to the simultaneously recorded datasets (**Figure 3G**), and is consistent with an inverse relationship between dimensionality and pair-wise correlations (see Equation 9).

Time Course of Dimensionality as a Function of Ensemble Size

Unlike ongoing activity, the dependence of dimensionality on ensemble size (the slope of the linear regression of d vs. N) was modulated during different epochs of the post-stimulus period [**Figure 3F**, full lines; two-way ANOVA; main effect of time $F_{(4, 495)} = 3.80$, $p < 0.005$; interaction time \times condition: $F_{(4, 495)} = 4.76$, $p < 0.001$]. In particular, the dependence of d on the ensemble size N almost disappeared immediately after stimulus presentation in the simultaneously recorded, but not in the shuffled ensembles (trial-matched slope in the first evoked second: $b_{\text{evoked}} = 0.07 \pm 0.01$ vs $b_{\text{shuff}} = 0.19 \pm 0.07$) and converged to a stable value after approximately 1 second (slope after the first second $b_{\text{evoked}} = 0.38 \pm 0.01$; compare with a stable average slope during ongoing activity of $b_{\text{ongoing}} = 0.57 \pm 0.01$, **Figure 3F**).

Note that the dimensionality is larger when the firing rate is computed in bins (as in **Figure 3F**) rather than in HMM states (as in **Figures 3B–D**, where the slopes are about half than in **Figure 3F**). The reason is that firing rates and correlations are



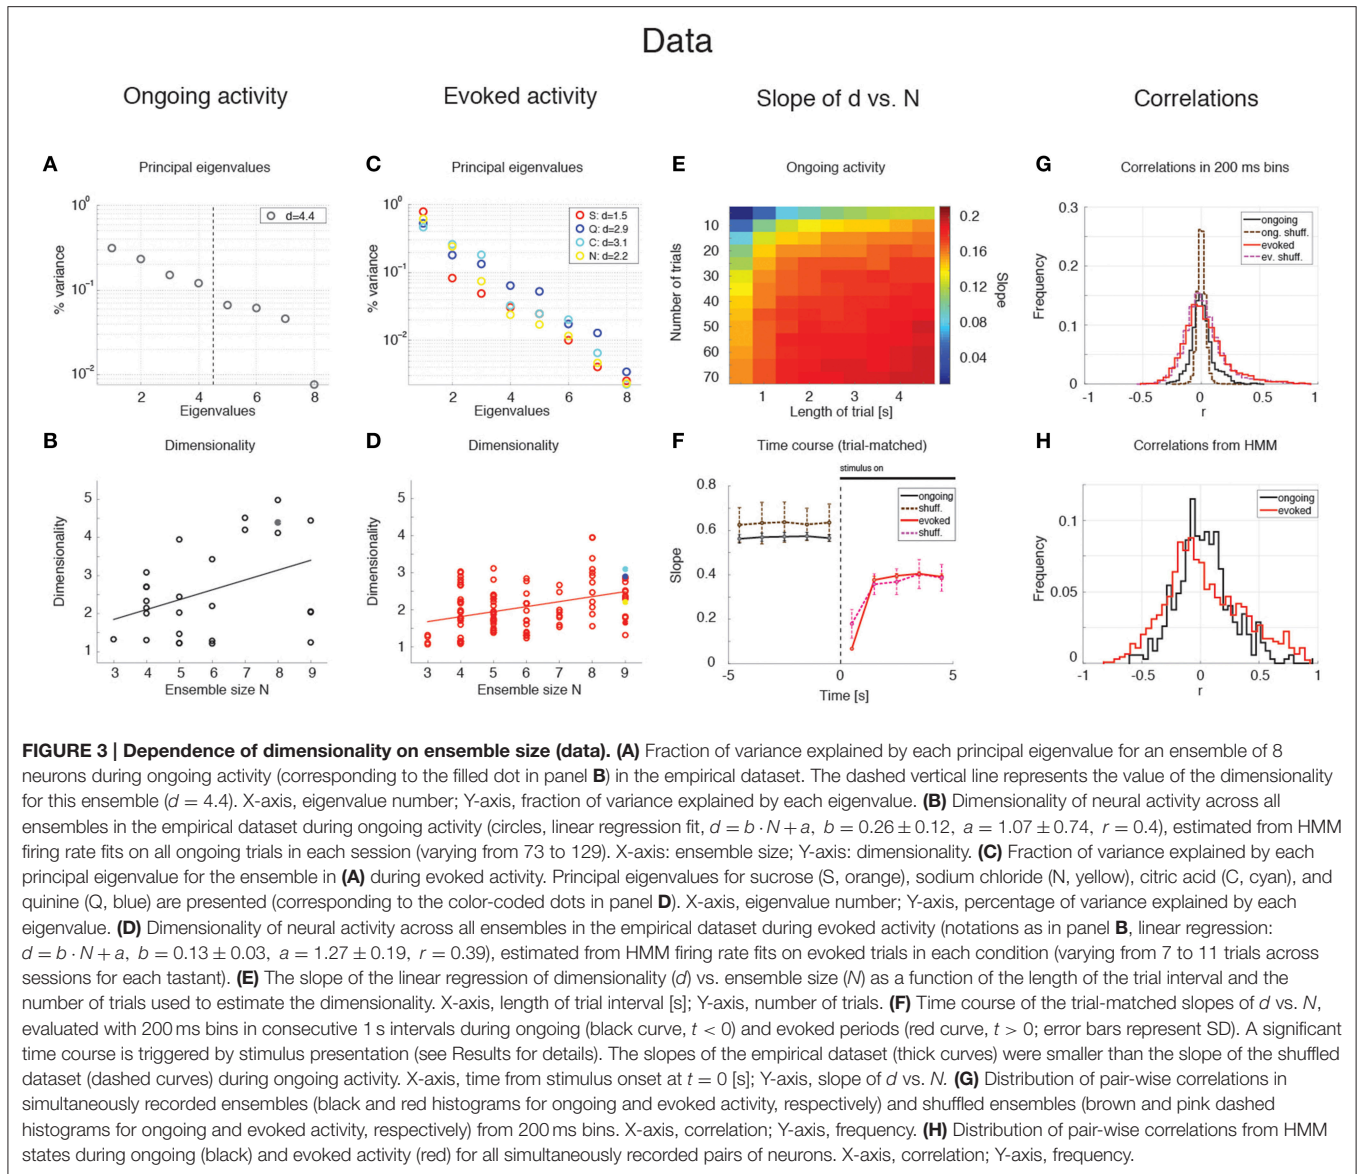
approximately constant during the same HMM state, whereas they may change when estimated in bins of fixed duration that include transitions among hidden states. These changes tend to dilute the correlations resulting in higher dimensionality as predicted e.g., by Equation (9). A comparison of the pairwise correlations of binned firing rates (Figure 3G) vs. those of firing rates in HMM states (Figure 3H) confirmed this hypothesis. Also, if the argument above is correct, one would expect a dependence of dimensionality on (fixed) bin duration. We computed the correlations and dimensionality of binned firing rates for various bin durations and found that r increases and d decreases for increasing bin durations (not shown). However, the slope of d vs. N is always larger in ongoing than in evoked activity regardless of bin size (ranging from 10 ms to 5 s; not shown). This confirms the generality of the results of Figures 3B–D, which were obtained using firing rate vectors in hidden states.

To summarize our main results so far, we found that dimensionality depends on ensemble size during both ongoing and evoked activity, and such dependence is significantly reduced in the post-stimulus period. This suggests that while state sequences during ongoing activity explore a large portion of the available firing rate space, the presentation of a stimulus initially collapses the state sequence along a more stereotyped and

lower-dimensional response (Katz et al., 2001; Jezzini et al., 2013). During both ongoing and evoked activity, the dimensionality is also different than expected by chance in a set of independent neurons (shuffled datasets).

Clustered Spiking Network Model of Dimensionality

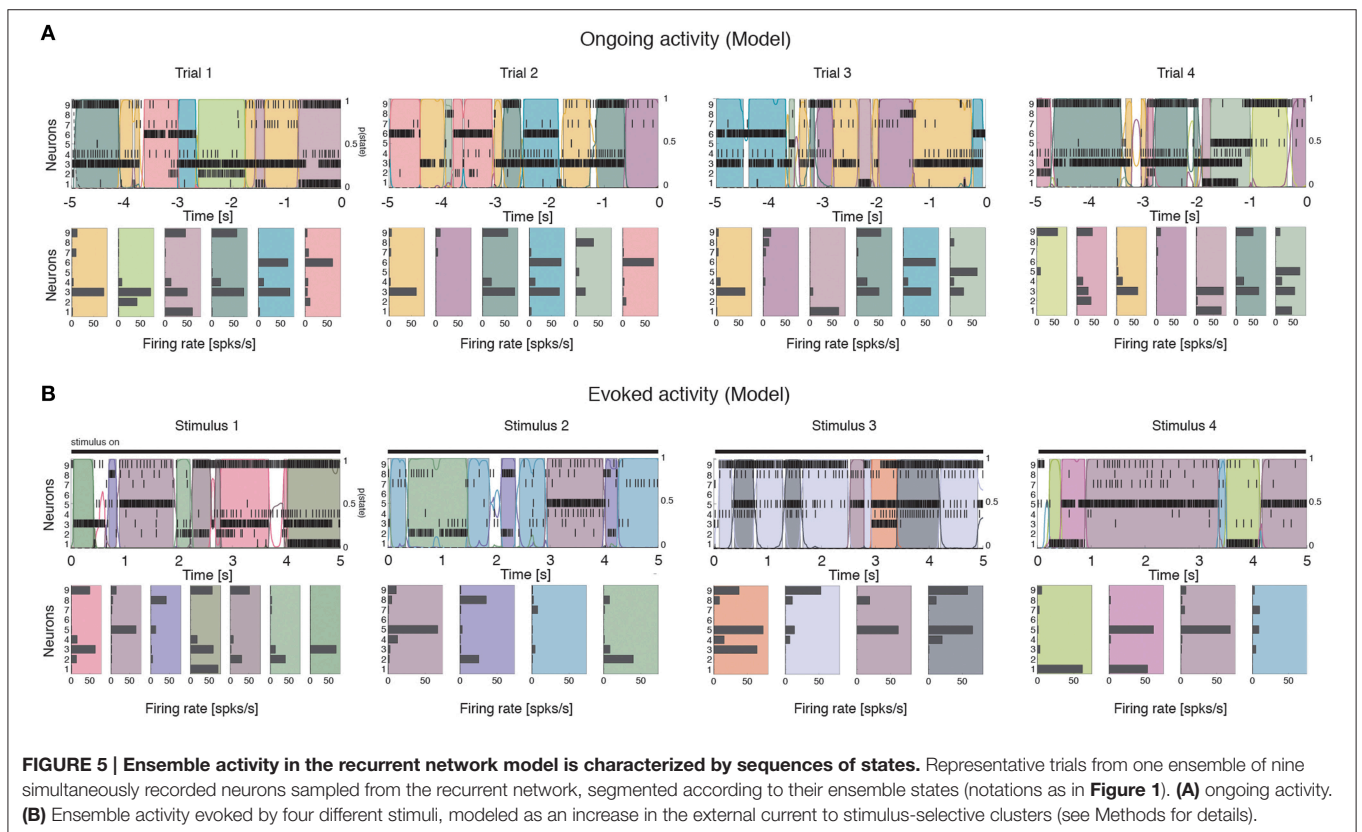
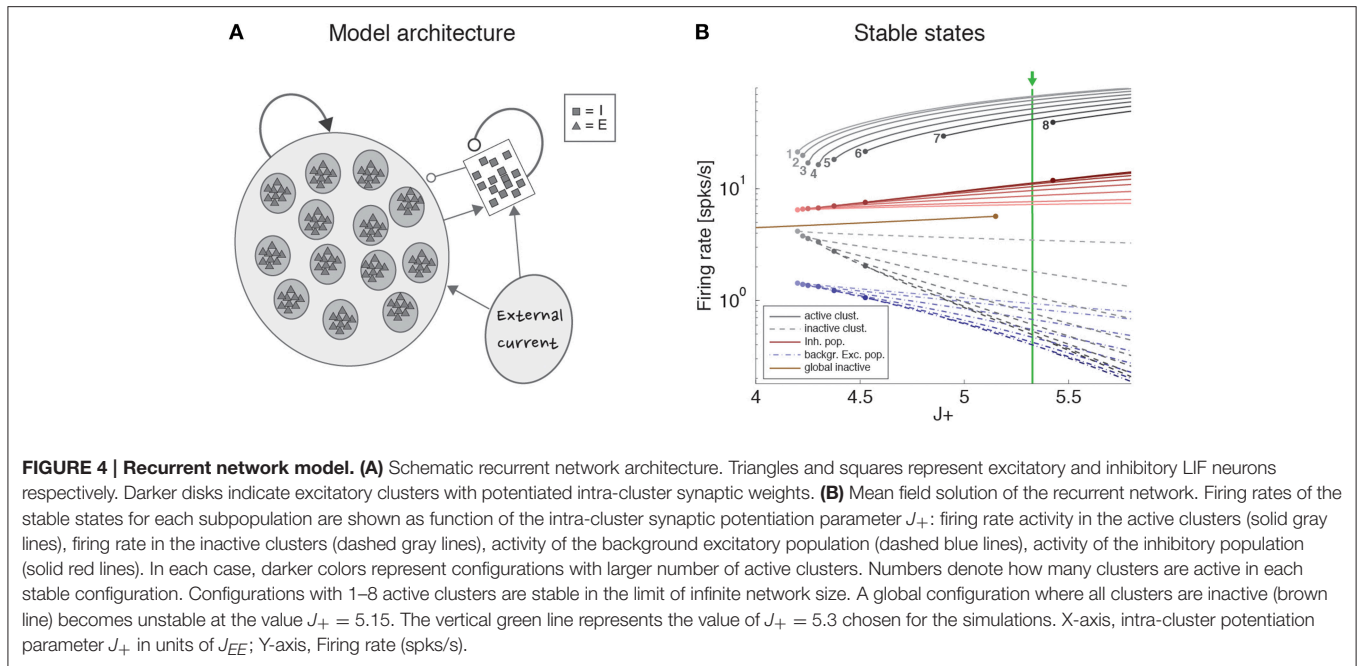
To gain a mechanistic understanding of the different dimensionality of ongoing and evoked activity, we have analyzed a spiking network model with clustered connectivity which has been shown to capture many essential features of the data (Mazzucato et al., 2015). In particular, the model reproduces the transitions among latent states in both ongoing and evoked activity. The network (see Methods for details) comprises Q clusters of excitatory neurons characterized by stronger synaptic connections within each cluster and weaker connections between neurons in different clusters. All neurons receive recurrent input from a pool of inhibitory neurons that keeps the network in a balanced regime of excitation and inhibition in the absence of external stimulation (Figure 4A). In very large networks (technically, in networks with an infinite number of neurons), the stable configurations of the neural activity are characterized by a finite number of active clusters whose firing rates depend on the number clusters active at any given moment, as shown



in Figure 4B (where $Q = 30$). In a finite network, however, finite size effects ignite transitions among these configurations, inducing network states (firing rate vectors) on randomly chosen subsets of neurons that resemble the HMM states found in the data (Figure 5; see Mazzucato et al., 2015 for details).

The dimensionality of the simulated sequences during ongoing and evoked activity was computed as done for the data, finding similar results. For the examples in Figure 5, we found $d_{\text{ongoing}} = 4.0$ for ongoing activity (Figure 6A) between $d_{\text{evoked}} = 2.2$ and $d_{\text{evoked}} = 3.2$ across tastes during evoked activity (Figure 6C). Across all simulated sessions, we found an average $d_{\text{ongoing}} = 2.9 \pm 0.9$ (mean \pm SD) for ongoing activity and $d_{\text{evoked}} = 2.4 \pm 0.7$ for evoked activity. The model captured the essential properties of dimensionality observed in the data: the dimensionality did not differ across different tastes (one-way ANOVA, $p > 0.2$) and depended on ensemble size during both

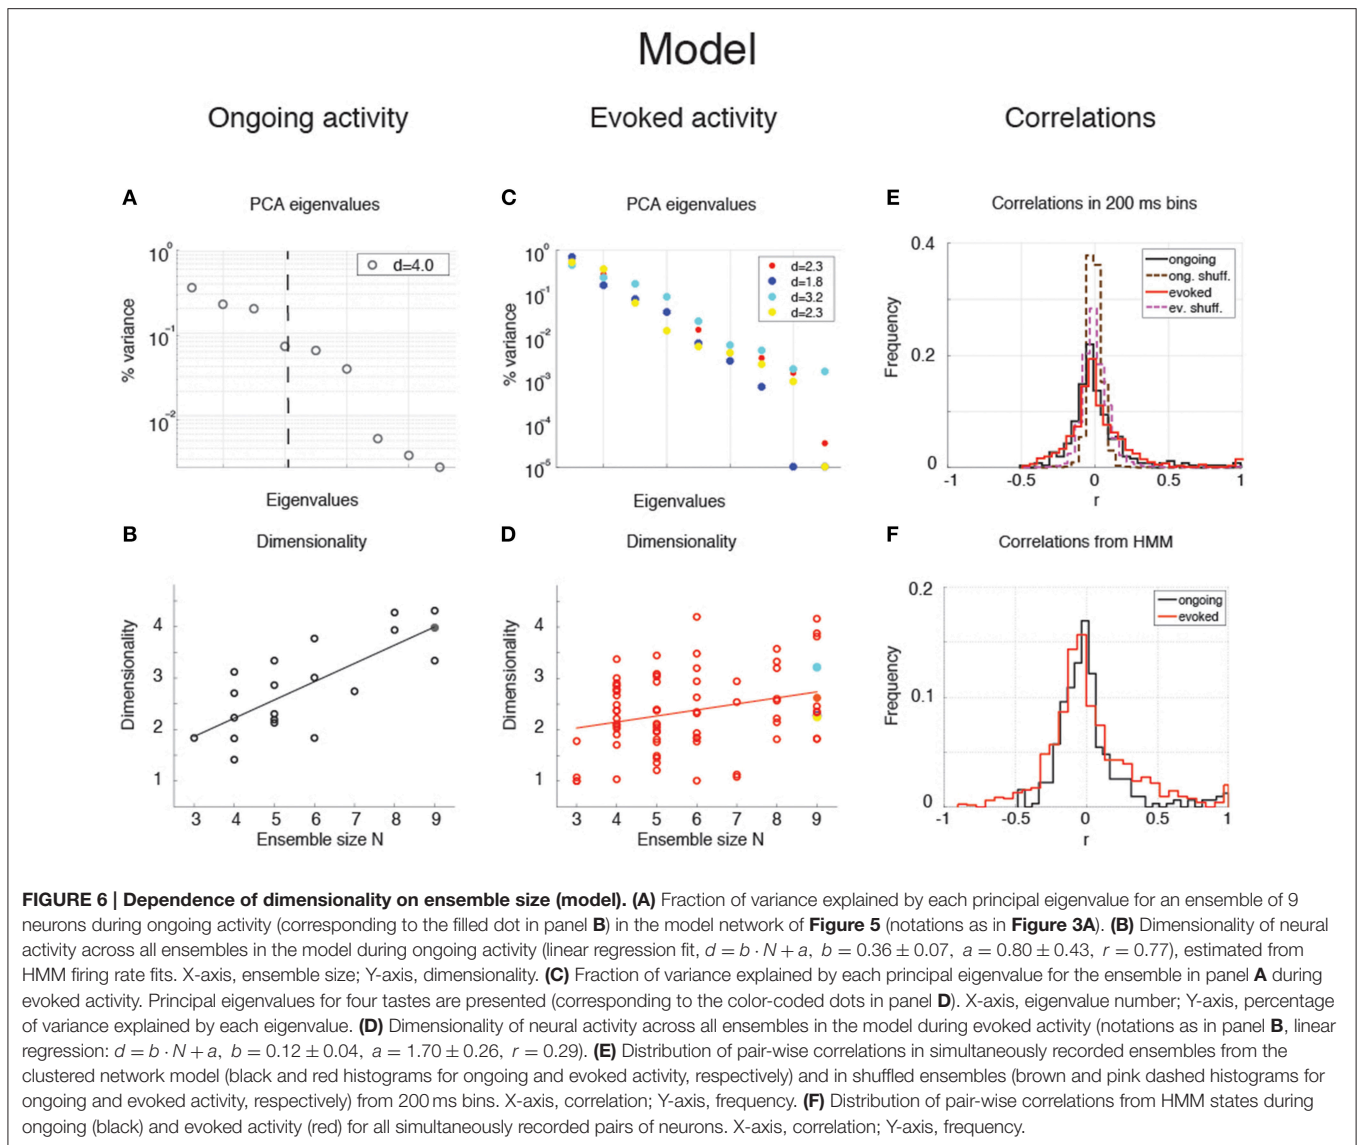
ongoing (Figure 6B; slope = 0.36 ± 0.07 , $r = 0.77$, $p < 10^{-4}$) and evoked periods (Figure 6D; slope = 0.12 ± 0.04 , $r = 0.29$, $p = 0.01$). As for the data, the dependency on ensemble size was smaller for evoked compared to ongoing activity. We performed a trial-matched two-way ANOVA as done on the data and found, also in the model, a main effect of condition [ongoing vs. evoked: $F_{(1, 146)} = 22.1$, $p < 10^{-5}$], a main effect of ensemble size [$F_{(6, 146)} = 14.1$, $p < 10^{-11}$], and a significant interaction [$F_{(6, 146)} = 3.8$, $p = 0.001$]. These results were accompanied by patterns of correlations among the model neurons (Figures 6E,F) very similar to those found in the data (Figures 3G,H; see section “Dimensionality is larger in the presence of clusters” for statistics of correlation values). As in the data, narrower distributions of correlations were found for binned firing rates (Figure 6E) compared to firing rates in hidden states (Figure 6F; compare with Figures 3G,H, respectively). Moreover, shuffling neurons



across datasets reduced the correlations (Figure 6E, dashed), resulting in a larger slope of d vs. N (not shown). Finally, d during ongoing activity was always larger than during evoked activity also when computed on binned firing rates (not shown), as found

in the data (see section “Dependence of dimensionality on bin size”).

Since the model was not fine-tuned to find these results, the different dimensionalities of ongoing and evoked activity, and



their associated patterns of pair-wise correlations, are likely the consequence of the organization in clusters and of the ensuing dynamics during ongoing and evoked activity.

Scaling of Dimensionality with Ensemble Size and Pair-Wise Correlations

The dependence of dimensionality on ensemble size observed in the data (Figure 3B) and in the model (Figure 6B) raises the question of whether or not the dimensionality would converge to an upper bound as one increases the number of simultaneously recorded neurons. In general, this question is important in a number of settings, related e.g., to coding in motor cortex (Ganguli et al., 2008; Gao and Ganguli, 2015), performance in a discrimination task (Rigotti et al., 2013), or coding of visual stimuli (Cadieu et al., 2013). We can attack this question aided by the model of Figure 4, where we can study the effect of large numbers of neurons, but also the impact on dimensionality of a

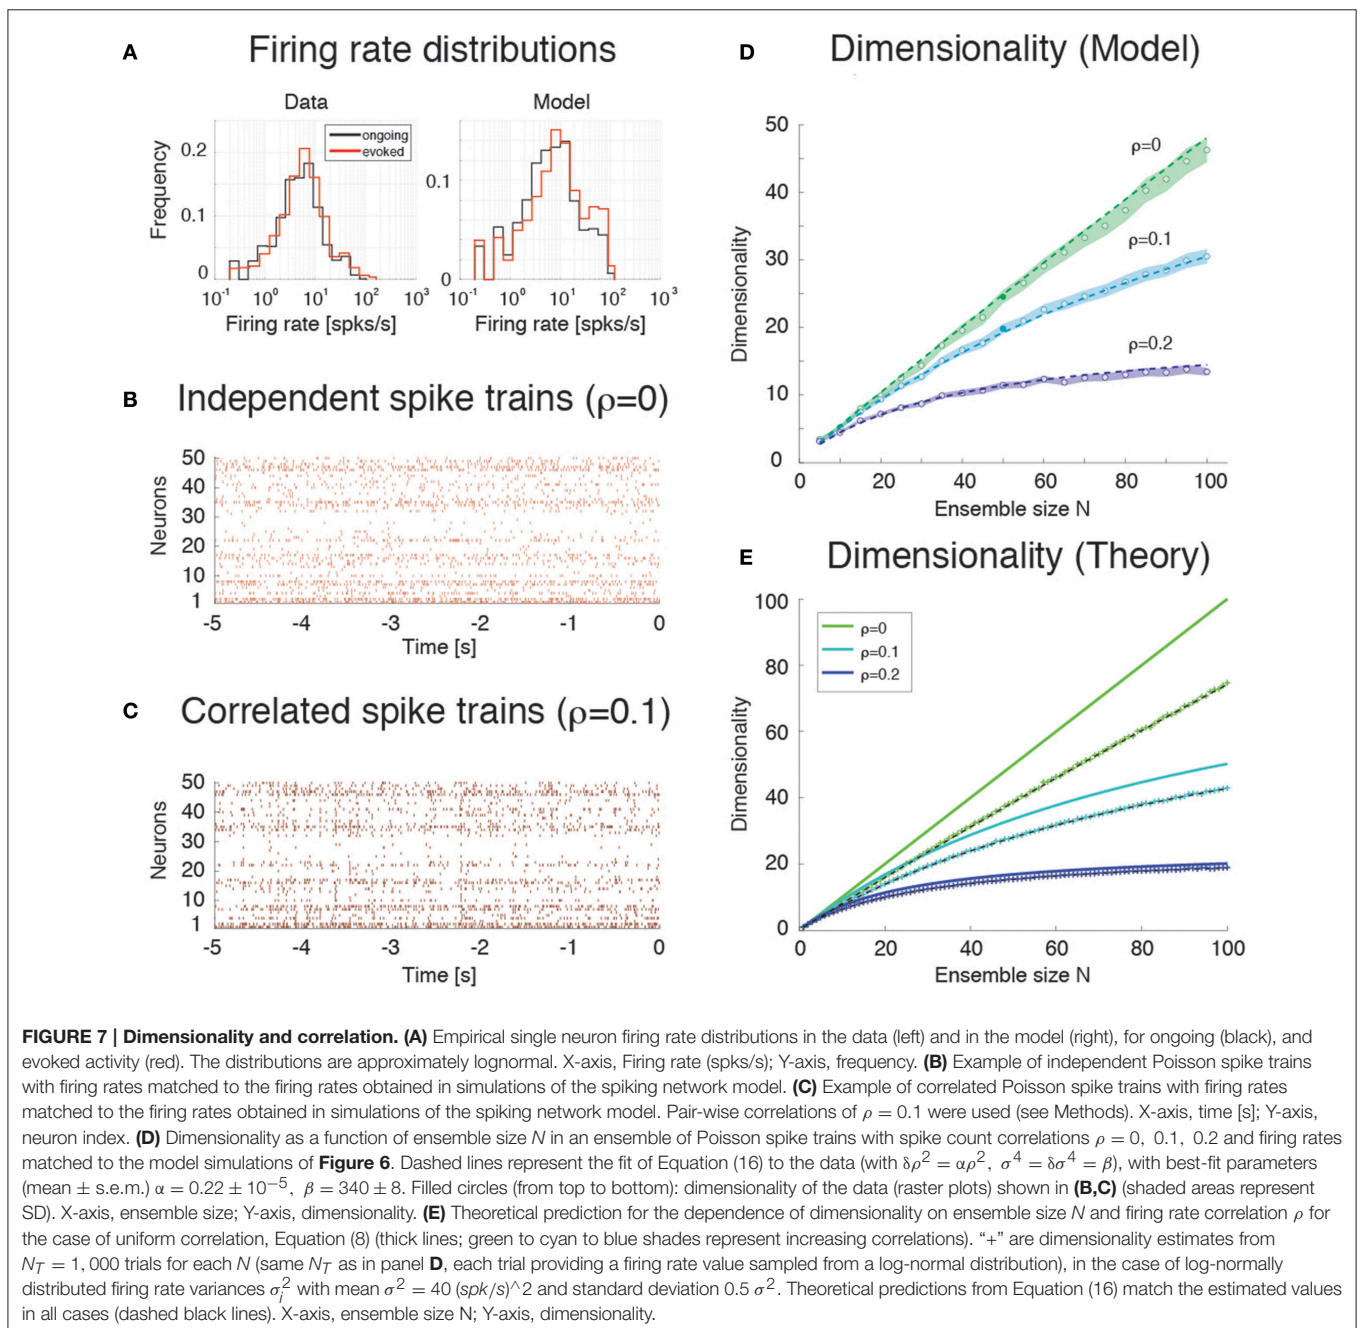
clustered network architecture compared to a homogeneous one, at parity of correlations and ensemble size.

We consider first the case of a homogeneous network of neurons having no clusters and low pair-wise correlations, but having the same firing rates distributions (which were approximately log-normal, Figure 7A) and the same mean pair-wise correlations as found in the data ($\rho \sim 0.01 - 0.2$). This would require solving a homogeneous recurrent network self-consistently for the desired firing rates and correlations. As a proxy for this scenario, we generated 20 sessions of 40 Poisson spike trains having exactly the desired properties (including the case of independent neurons for which $\rho = 0$). Two examples with $\rho = 0$ and $\rho = 0.1$, respectively, are shown in Figures 7B,C. Since in the asynchronous homogeneous network there are no transitions and hence no hidden states, the dimensionality was estimated based on the rate vectors in bins of 200 ms duration (using bin widths of 50–500 ms did not change the results; see Methods for details).

We found that the dimensionality grows linearly with ensemble size in the absence of correlations, but is a concave function of N in the presence of pair-wise correlations (circles in **Figure 7D**). Thus, as expected, the presence of correlations reduces the dimensionality and suggests the possibility of an upper bound. A simple theoretical calculation mimicking this scenario shows that d in this case converges indeed to an upper bound that depends on the inverse of the square of the pair-wise correlations. For example, in the case of uniform correlations (ρ) and equal variances of the spike counts, Equation (8) of Methods, $d(N, \rho) = \frac{1}{\rho^2 + (1 - \rho^2)/N}$, shows that $d = N$ in the absence of

correlations, but $d < 1/\rho^2$ in the presence of correlations. These properties remain approximately true if the variances σ_i^2 of the firing rates are drawn from a distribution with mean $E[\sigma_i^2] = \sigma^2$ and variance $Var[\sigma_i^2] = \delta\sigma^4$. As Equation (9) shows, in such a case dimensionality is reduced compared to the case of equal variances, for example $d \approx \frac{\sigma^4}{\sigma^4 + \delta\sigma^4} N < N$ for large N when $\rho = 0$, $\delta\rho = 0$.

The analytical results are shown in **Figure 7E** (full lines correspond to Equation 8), together with their estimates (“+”) based on 1000 data points (same number as trials in **Figure 7D**; see Methods). The estimates are based on surrogate datasets



with lognormal-distributed variances σ_i^2 to mimic the empirical distribution of variances found in GC (not shown).

Estimation Bias

Comparison of **Figures 7D,E** shows that the dimensionality of the homogeneous network is underestimated compared to the theoretical value given by Equation (8). This is due to a finite number of trials and the presence of unequal variances with spread $\delta\sigma^4$ (“+” in **Figure 7E**). As **Figure 7E** shows, taking this into account will reduce the dimensionality to values comparable to those of the homogeneous network of **Figure 7D**. The dimensionality in that case is well predicted by Equation (16) (broken lines in **Figure 7E**). The same Equation (16) was fitted successfully to the data in **Figure 7D** (dashed) by tuning 2 parameters to account for the unknown variance and correlation width of the firing rates (see Methods for details).

Empirically, estimates of the dimensionality Equation (2) based on a finite number N_T of trials tend to underestimate d (**Figure 3E**). The approximate estimator Equation (16) confirms that, for any ensemble size N , d is a monotonically increasing function of the number of trials (**Figure 8A**). Note that this holds for the mean value of the estimator (Equation 16) over many datasets, not for single estimates, which could overestimate the true d (not shown). Equation (16) also provides an excellent description of dimensionality as a function of firing rates’ variance $\delta\sigma^4$ (**Figure 8B**) and pair-wise correlations width $\delta\rho^2$ (**Figure 8C**). In particular, the mean and the variance of the pair-wise correlations have an interchangeable effect on d (see Equation 16); they both decrease the dimensionality and so does the firing rate variance $\delta\sigma^4$ (**Figure 8B**).

Scaling of Dimensionality in the Presence of Clusters

We next compared the dimensionality of the homogeneous network’s activity to that predicted by the clustered network model of **Figure 4**. To allow comparison with the homogeneous network, dimensionality was computed based on the spike counts in 200 ms bins rather than the HMM’s firing rate vectors as in **Figure 6** (see Methods for details).

We found that the dependence of d on N in the clustered network depends on how the neurons are sampled. If the sampling is completely random, so that any neuron has the same probability of being added to the ensemble regardless of cluster membership, a concave dependence on N will appear, much like the case of the homogeneous network (**Figure 9A**, dashed lines). However, if neurons are selected one from each cluster until all clusters have been sampled once, then one neuron from each cluster until all clusters have been sampled twice, and so on, until all the neurons in the network have been sampled, then the dependence of d on N shows an abrupt transition when $N = Q$, i.e., when the number of sampled neurons reaches the number of clusters in the network (**Figure 9A**, full lines; see **Figure 9B** for raster plots with $Q = 30$ and $N = 50$). In the following, we refer to this sampling procedure as “ordered sampling,” as a reminder that neurons are selected randomly from each cluster, but the clusters are selected in serial order. For $N \leq Q$, the dimensionality grows linearly with ensemble size in both ongoing (slope 0.24 ± 0.01 , $r = 0.79$, $p < 10^{-10}$, black line) and evoked periods (slope 0.19 ± 0.01 , $r = 0.84$, $p < 10^{-10}$; red line), and was larger during ongoing than evoked activity [trial-matched two-way ANOVA, main effect:

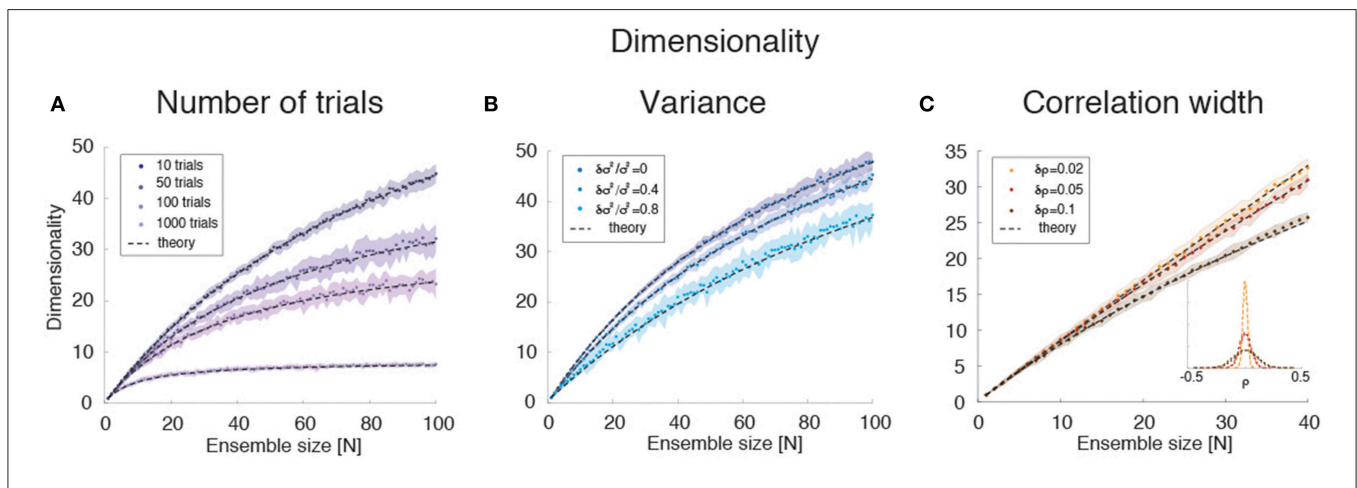
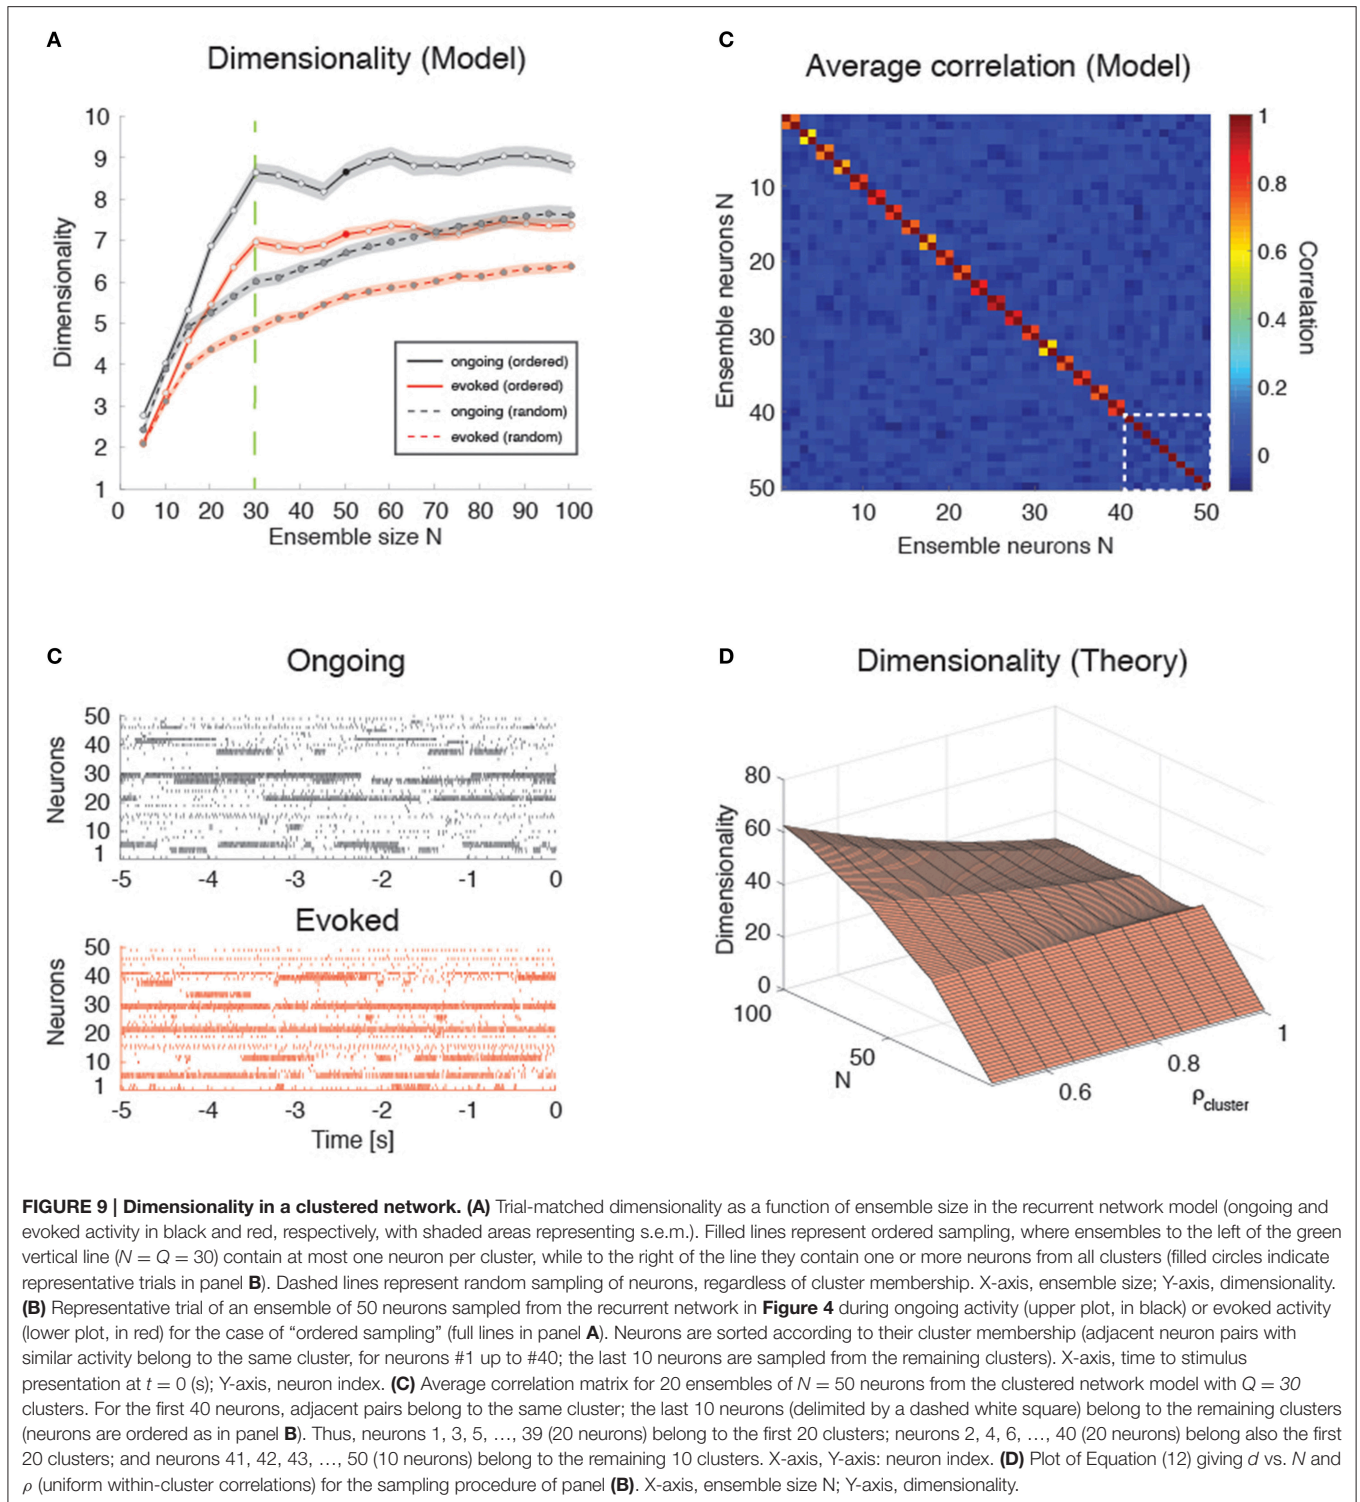


FIGURE 8 | Dimensionality estimation. (A) Dependence of dimensionality on the number of trials for variable ensemble size N , for fixed correlations $\rho = 0.1$ and firing rates variances σ_i^2 with mean σ^2 and standard deviation $\delta\sigma^2 = 0.4\sigma^2$. Dashed lines: theoretical prediction, Equation (16); dots: mean values from simulations of 20 surrogate datasets containing 10–1000 trials each (shaded areas: SD), with darker shades representing increasing number of trials. X-axis: ensemble size; Y-axis, dimensionality. (B) Dependence of dimensionality on the spread $\delta\sigma^2$ of the firing rates variances for fixed correlations $\rho = 0.1$ and firing rate variance with mean σ^2 . Dashed lines: theoretical prediction, Equation (16); dots: mean values from simulations of 20 surrogate datasets containing 1000 trial each (shaded areas: SD), with lighter shades representing increasing values of $\delta\sigma^2/\sigma^2$. X-axis, ensemble size; Y-axis, dimensionality. (C) Dependence of dimensionality on the width $\delta\rho = \sqrt{\text{Var}(\rho)}$ of pair-wise firing rate correlations (with zero mean, $\rho = 0$), for firing rates variances σ_i^2 with mean σ^2 and standard deviation $\delta\sigma^2 = 0.4\sigma^2$. Dashed lines: theoretical prediction, Equation (16); dots: mean values from simulations of 20 surrogate datasets containing 1000 trials each (shaded areas: SD), with darker shades representing increasing values of $\delta\rho$. Inset: distribution of correlation coefficients used in the main figure. X-axis, ensemble size; Y-axis, dimensionality. In all panels, $\sigma^2 = 40 (\text{spk/s})^2$.



$F_{(1, 948)} = 168$, $p < 10^{-30}$; interaction: $F_{(5, 948)} = 4.1$, $p < 0.001$].

These results are in keeping with the empirical and model results based on the HMM analysis (**Figures 3, 6**). However, in the case of ordered sampling, the dependence of dimensionality on ensemble size tends to disappear for $N \geq Q$ both during

ongoing (slope 0.010 ± 0.003 , $r = 0.1$, $p < 0.001$) and evoked periods (slope 0.009 ± 0.002 , $r = 0.13$, $p < 10^{-4}$; **Figure 9A**, full lines). The average dimensionality over the range $30 \leq N \leq 100$ was significantly larger for ongoing, $d_{\text{ongoing}} = 8.74 \pm 0.06$, than for evoked activity, $d_{\text{evoked}} = 7.15 \pm 0.04$ [trial-matched two-way ANOVA, main effect: $F_{(1, 2212)} = 488$, $p < 10^{-30}$], confirming

that dimensionality during ongoing is larger than during evoked activity also in this case. The difference in dimensionality between ongoing and evoked activity also holds in the case of random sampling on the entire range of N values (Figure 9A, dashed lines), confirming the generality of this finding.

Dimensionality is Larger in the Presence of Clusters

Intuitively, the dimensionality saturates at $N = Q$ in the clustered network because additional neurons will be highly correlated with already sampled ones. For $N \leq Q$, each new neuron's activity adds an independent degree of freedom to the neural dynamics and thus increases its dimensionality. For $Q > N$, additional neurons are highly correlated with an existing neuron, adding little or no additional contribution to d . Indeed, compared to the low overall correlations found across all neuron pairs in the data (and used as *desiderata* for the homogeneous network), neurons belonging to the same model cluster had a much higher correlation of $\rho = 0.92$ [0.56, 0.96] (median and [25, 75]-percentile), while neurons belonging to different clusters had negligible correlation ($\rho \approx 0$, [-0.10, 0.06]). A negligible median correlation was typical: for example, negligible was the overall median correlation regardless of cluster membership ($\rho \approx 0$ [-0.109, 0.083]); and the empirical correlation both during ongoing ([-0.047, 0.051], with rare maximal values of $\rho \sim 0.5$), and evoked activity ([-0.085, 0.113], with rare maximal values of $\rho \sim 0.9$). While we note the qualitative agreement of model and empirical correlations, we emphasize that these numbers were obtained using 200 ms bins and that they were quite sensitive to bin duration. In particular, the maximal correlations (regardless of sign) were substantially reduced for smaller bin durations (not shown).

Plugging these values into a correlation matrix reflecting the clustered architecture and the “ordered” sampling procedure used in Figure 9B, we obtained the matrix shown in Figure 9C, where pairwise correlations depend on whether or not the neurons belong to the same cluster (for the first 40 neurons, adjacent pairs belong to the same cluster; the last 10 neurons belong to the remaining clusters). It is natural to interpret such correlation matrix as the noisy observation of a block-diagonal matrix such that neurons in the same cluster have uniform correlation while neurons from different clusters are uncorrelated. For such a correlation matrix the dimensionality can be evaluated exactly (see Equation 12 of Methods). In the approximation where all neurons have the same variance, this reduces to Equation (13), i.e.,

$$d(N, \rho) = \begin{cases} N & , N \leq Q \\ \frac{N}{1+m\rho^2[1-(Q-p)/N]} & , N > Q \end{cases}$$

where $N = mQ + p$. This formula is plotted in Figure 9D for relevant values of ρ and N and it explains the origin of the abrupt transition in dimensionality at $Q = N$. (The reasons for a dimensionality lower than N for $N \leq Q$ in the data—see Figure 9A—are, also in this case, the finite number of data points (250) used for its estimation and the non-uniform distributions of firing rate variances and correlations).

Note that the formula also predicts cusps in dimensionality (which become local maxima for large ρ) whenever the ensemble size is an exact multiple of the number of clusters. This is also visible in the simulated data of Figure 9A, where local maxima seem to appear at $N = 30, 60, 90$ with $Q = 30$ clusters. It is also worth mentioning that, for low intra-cluster correlations, the dependence on N predicted by Equation (13) becomes smoother and the cusps harder to detect (not shown), suggesting that the behavior of a clustered network with weak clusters tends to converge to the behavior of a homogeneous asynchronous network—therefore lacking sequences of hidden states. Thus, the complexity of the network dynamics is reflected in how its dimensionality scales with N , assuming that one may sample one neuron per cluster (i.e., via “ordered sampling”).

Even though it is not clear how to perform ordered sampling empirically (see Discussion), this result is nevertheless useful since it represents an upper bound also in the case of random sampling (see Figure 9A, dashed lines). Equation (13) predicts that $d \leq Q/\rho^2$, with this value reached asymptotically for large N . In the case of random sampling, growth to this bound is even slower (Figure 9A). For comparison, in a homogeneous network $d \leq 1/\rho^2$ from Equation (8), a bound that is smaller by a factor of Q . Finally, homogeneous dimensionality is dominated by clustered dimensionality also in the more realistic case of non-uniform variances and correlations, where similar bounds are found in both cases (see Methods for details).

DISCUSSION

In this paper we have investigated the dimensionality of the neural activity in the gustatory cortex of alert rats. Dimensionality was defined as a collective property of ensembles of simultaneously recorded neurons that reflects the effective space occupied by the ensemble activity during either ongoing or evoked activity. If one represents ensemble activity in terms of firing rate vectors, whose dimension is the number of ensemble neurons N , then the collection of rate vectors across trials takes the form of a set of points in the N -dimensional space of firing rates. Roughly, dimensionality is the minimal number of dimensions necessary to provide an accurate description of such set of points, which may be localized on a lower-dimensional subspace inside the whole firing rate space.

One of the main results of this paper is that the dimensionality of evoked activity is smaller than that of ongoing activity, i.e., stimulus presentation quenches dimensionality. More specifically, the dimensionality is linearly related to the ensemble size, with a significantly larger slope during ongoing activity compared to evoked activity (compare Figures 3B,D). We explained this phenomenon using a biologically plausible, mechanistic spiking network model based on recurrent connectivity with clustered architecture. The model was recently introduced in Mazzucato et al. (2015) to account for the observed dynamics of ensembles of GC neurons as sequences of metastable states, where each state is defined as a vector of firing rates across simultaneously recorded neurons. The model captures the reduction in trial-to-trial variability and the multiple firing rates attained by single neurons across different

states observed in GC upon stimulus presentation. Here, the same model was found to capture also the stimulus-induced reduction of dimensionality. While the set of active clusters during ongoing activity varies randomly, allowing the ensemble dynamics to explore a large portion of firing rate space, the evoked set of active clusters is limited mostly to the stimulus-selective clusters only (see Mazzucato et al., 2015 for a detailed analysis). The dynamics of cluster activation in the model thus explains the more pronounced dependence of dimensionality on ensemble size found during ongoing compared to evoked activity.

We presented a simple theory of how dimensionality depends on the number of simultaneously recorded neurons N , their firing rate correlations, their variance, and the number and duration of recording trials. We found that dimensionality increases with N and decreases with the amount of pair-wise correlations among the neurons (e.g., **Figure 8C**). At parity of correlations, dimensionality is maximal when all neurons have the same firing rate variance, and it decreases as the distribution of count variances becomes more heterogeneous (e.g., **Figure 8B**). The estimation of dimensionality based on a finite dataset is an increasing function of the number of trials (**Figure 8A**). Finally, introducing clustered correlations in the theory, and sampling one neuron per cluster as in **Figure 9B**, results in cusps at values of N that are multiples of the number of clusters (**Figure 9D**), in agreement with the predictions of the spiking network model (**Figure 9A**, full lines).

Dimensionality Scaling with Ensemble Size

The increased dimensionality with sample size, especially during ongoing activity, was found empirically in datasets with 3–9 neurons per ensemble, but could be extrapolated for larger N in a spiking network model with homogeneous or clustered architecture. In homogeneous networks with finite correlations the dimensionality is predicted to increase sub-linearly with N (Equation 8), whereas in the clustered network it may exhibit cusps at multiple values of the number of clusters (**Figure 9A**), and would saturate quickly to a value that depends on the ratio of the number of clusters Q and the amount of pair-wise correlations, $d \leq Q/\rho^2$. Testing this prediction requires the ability to sample neurons one from each cluster, until all clusters are sampled, and seems beyond the current recording techniques. However, looking for natural groupings of neurons based on response similarities could uncover spatial segregation of clusters (Kiani et al., 2015) and could perhaps allow sampling neurons according to this procedure. Moreover, the model predicts a slower approach to a similar bound also in the case of random sampling.

Dimensionality in a homogeneous network is instead bounded by $1/\rho^2$, and hence it is a factor Q smaller than in the clustered network. Dimensionality is maximal in a population of independent neurons ($\rho = 0$), where it grows linearly with N ; however, neurons of recurrent networks have wide-ranging correlations (see e.g., **Figures 6E,F** and its empirical counterpart, **Figures 3G,H**). Since the presence of even low correlations can dramatically reduce the dimensionality (see **Figure 7D**), the neural activity in a clustered architecture can reach much higher

values at parity of correlations, representing an intermediate case between a homogeneous network and a population of independent neurons.

Evidence for the presence of spatial clusters has been recently reported in the prefrontal cortex based on correlations analyses (Kiani et al., 2015). An alternative possibility is that neural clusters are not spatially but functionally arranged, and cluster memberships vary with time and task complexity (Rickert et al., 2009). Can our model provide indirect tools to help uncover the presence of clusters? A closer look at **Figures 6E,F** reveal a small peak at large correlations due to the contribution of highly correlated neurons belonging to the same cluster. This peak would be absent in a homogenous network and thus is the signature of a clustered architecture. However, such peak is populated by only small fraction ($1/Q$) of the total number of neuron pairs, which hinders its empirical detection (no peak at large correlations is clearly visible in our data, see **Figures 3G,H**).

Dimensionality and Trial-to-Trial Variability

Cortical recordings from alert animals show that neurons produce irregular spike trains with variable spike counts across trials (Shadlen and Newsome, 1994; Fontanini and Katz, 2008; Moreno-Bote, 2014). Despite many efforts, it remains a key issue to establish whether variability is detrimental (Gur et al., 1997; White et al., 2012) or useful (McDonnell and Ward, 2011) for neural computation.

Trial-to-trial variability is reduced during preparatory activity (Churchland et al., 2006), during the presentation of a stimulus (Churchland, 2010b), or when stimuli are expected (Samuelsen et al., 2012), a phenomenon that would not occur in a population of independent or homogeneously connected neurons (Litwin-Kumar and Doiron, 2012). Recent work has shown that the stimulus-induced reduction of trial-to-trial variability can be due to spike-frequency adaptation in balanced networks (Farkhooi et al., 2013) or to slow dynamic fluctuations generated in a recurrent spiking networks with clustered connectivity (Deco and Hugues, 2012; Litwin-Kumar and Doiron, 2012; Mazzucato et al., 2015). In clustered network models, slow fluctuations in firing rates across neurons can ignite metastable sequences of neural activity, closely resembling metastable sequences observed experimentally (Abeles et al., 1995; Seidemann et al., 1996; Jones et al., 2007; Kemere et al., 2008; Durstewitz et al., 2010; Ponce-Alvarez et al., 2012; Mazzucato et al., 2015). The slow, metastable dynamics of cluster activation produces high variability in the spike count during ongoing activity. While cluster activations occur at random times during ongoing activity periods, stimulus presentation locks cluster activation at its onset, leading to a decrease in trial-to-trial variability.

Similarly, a stimulus-induced reduction of dimensionality is obtained in the same model. In this case, preferred cluster activation due to stimulus onset generates an increase in pair-wise correlations that reduce dimensionality. Note that the two properties (trial-to-trial variability and dimensionality) are conceptually distinct. An ensemble of Poisson spike trains can be highly correlated (hence have low dimensionality), yet the Fano Factor of each spike train will still be 1 (hence high),

independently of the correlations among neurons. In a recurrent network, however, dimensionality and trial-to-trial variability may become intertwined and exhibit similar properties, such as the stimulus-induced reduction observed in a model with clustered connectivity. A deeper investigation of the link between dimensionality and trial-to-trial variability in recurrent networks is left for future studies.

Alternative Definitions of Dimensionality

Following (Abbott et al., 2011) we have defined dimensionality (Equation 2) as the dimension of an effective linear subspace of firing rate vectors containing the most variance of the neural activity. It differs somewhat from the typical dimensionality reduction based on PCA that retains only the number of eigenvectors explaining a predefined amount of variance (see Broome et al., 2006; Geffen et al., 2009), because Equation (2) includes contribution from all eigenvalues. Moreover, we have computed the firing rate correlations in bins of variable width that match the duration of the HMM states. Although, our main results do not depend on bin size (see Results' section "Time course of dimensionality as a function of ensemble size"), the actual value of dimensionality decreases with increasing bin duration. Thus, any choice of bin size (e.g., 200 ms in **Figures 3E,G**) remains somewhat arbitrary. A better method is to use a variable bin size as dictated by the HMM analysis, as done in **Figures 3B–D**. This method also prevents diluting correlations among firing rates that would occur if one neuron were to change state inside the current bin, because during a hidden state the firing rates of the neurons are constant (by definition). Thus, this provides a principled adaptive procedure for selecting the bin size and eliminates the dependence of dimensionality on the bin width used for the analysis.

Other definitions of neural dimensionality have been proposed in the literature, which aim at capturing different properties of the neural activity, typically during stimulus-evoked activity. A measure of dimensionality related to ours, and referred to as "complexity," was introduced in Cadieu et al. (2013). According to their definition, population firing rate vectors from all evoked conditions were first decomposed along their kernel Principal Components (Montavon et al., 2011). A linear classifier was then trained on an increasing number of leading PCs in order to perform a discrimination task, where the number of PCs used was defined as the complexity of the representation. In general, the classification accuracy improves with increasing complexity, and it may saturate when all PCs containing relevant features are used—with the remaining PCs representing noise or information irrelevant to the task. Reaching high accuracy at low complexity implies good generalization performance, i.e., the ability to classify novel variations of a stimulus in the correct category. Neural representations in monkey inferotemporal cortex (IT) were found to require lower complexity than in area V4, confirming IT's premier role in classifying visual objects despite large variations in shape, orientation and background (Cadieu et al., 2013). Complexity relies on a supervised algorithm and is an efficient tool to capture the generalization properties of evoked representations (see DiCarlo et al., 2012) for its relevance to visual object recognition).

A second definition of dimensionality, sometimes referred to as "shattering dimensionality" in the Machine Learning literature, has been used to assess the discrimination properties of the neural representation (Rigotti et al., 2013). Given a set of p firing rate vectors, one can split them into two classes (e.g., white and black colorings) in 2^p different ways, and train a classifier to learn as many of those binary classification labels as possible. The shattering dimensionality is then defined as (the logarithm of) the largest number of binary classifications that can be implemented. This measure of dimensionality was found to drop significantly in monkey prefrontal cortex during the error trials of a recall task, and thus predicts the ability of the monkey to correctly perform the task (Rigotti et al., 2013).

A flexible and informative neural representation is one that achieves a large shattering dimensionality (good discrimination) while keeping a low complexity (good generalization). Note that both complexity and shattering dimensionality represent measures of classification performance in task-related paradigms, and their definition requires a set of evoked conditions to be classified via a supervised learning algorithm. While both definitions could be applied to neural activity in our stimulus-evoked data, their interpretation is not readily extended to periods of ongoing activity, as the latter is not associated to desired targets in a way that can be learned by a classification algorithm. Since our main aim was to compare the dimensionality of ongoing and evoked activity, the unsupervised approach of Abbott et al. (2011) and their notion of "effective" dimensionality was better suited for our analysis. A related definition of dimensionality has been used by Gao and Ganguli (2015) to investigate neural representations of movements in motor cortex.

Many measures of dimensionality used in the literature (including ours and some of those discussed above) are based on pair-wise correlations. However, neural activity is known to give rise also to higher-order correlations (Martignon et al., 2000). Given that the extent and relevance of higher-order correlations is actively debated (Schneidman et al., 2006; Staude et al., 2010), it would be useful to include them in measures of dimensionality. This is left for a future study.

Ongoing Activity and Task Complexity

The relationship between ongoing and stimulus-evoked activity has been linked to the functional connectivity of local cortical circuits, and their mutual relationship has been the object of both theoretical and experimental investigations, often with contrasting conclusions (e.g., Arieli et al., 1996; Tsodyks et al., 1999; Kenet et al., 2003; Luczak et al., 2009; Tkacik et al., 2010; Berkes et al., 2011; Mazzucato et al., 2015). Here, we have focused on the dimensionality of ongoing and evoked activity and have shown that neural activity during ongoing periods occupies a space of larger dimensionality compared to evoked activity. Although, based on a different measure of dimensionality, recent results on the relation between the dimensionality of evoked activity and task complexity suggest that evoked dimensionality is roughly equal to the number of task conditions (Rigotti et al., 2013). It is natural to ask whether the dimensionality of ongoing activity provides an estimate of the complexity of the hardest task

that can be supported by the neural activity. Moreover, based on the clustered network model, the presence of clusters imposes an upper value $d \leq Q/\rho^2$ during ongoing activity, suggesting that a discrimination task with up to $\propto Q$ different conditions may be supported. The experience of taste consumption is by itself multidimensional, including chemo- and oro-sensory aspects (i.e., taste identity Jezzini et al., 2013, and concentration Sadacca et al., 2012, texture, temperature, Yamamoto et al., 1981, 1988) as well as psychological aspects (hedonic value Katz et al., 2001; Grossman et al., 2008, anticipation Samuelsen et al., 2012; Gardner and Fontanini, 2014, novelty Inberg et al., 2013; Bermudez-Rattoni, 2014, and satiety effects de Araujo et al., 2006). It is tempting to speculate that neural activity during ongoing periods explores all these different dimensions, while evoked activity is confined to the features of the particular taste being delivered or attended in a specific context.

REFERENCES

- Abbott, L. F., Rajan, K., and Sompolinsky, H. (2011). "Interactions between intrinsic and stimulus-evoked activity in recurrent neural networks," in *The Dynamic Brain: An Exploration of Neuronal Variability and its Functional Significance*, eds D. L. Glanzman and M. Ding (New York, NY: Oxford University Press), 65–82.
- Abeles, M., Bergman, H., Gat, I., Meilijson, I., Seidemann, E., Tishby, N., et al. (1995). Cortical activity flips among quasi-stationary states. *Proc. Natl. Acad. Sci. U.S.A.* 92, 8616–8620. doi: 10.1073/pnas.92.19.8616
- Amit, D. J., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* 7, 237–252. doi: 10.1093/cercor/7.3.237
- Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273, 1868–1871. doi: 10.1126/science.273.5283.1868
- Berkes, P., Orbán, G., Lengyel, M., and Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331, 83–87. doi: 10.1126/science.1195870
- Bermudez-Rattoni, F. (2014). The forgotten insular cortex: its role on recognition memory formation. *Neurobiol. Learn. Mem.* 109, 207–216. doi: 10.1016/j.nlm.2014.01.001
- Broome, B. M., Jayaraman, V., and Laurent, G. (2006). Encoding and decoding of overlapping odor sequences. *Neuron* 51, 467–482. doi: 10.1016/j.neuron.2006.07.018
- Brunel, N., and Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.* 11, 1621–1671. doi: 10.1162/089976699300016179
- Brunel, N., and Sergi, S. (1998). Firing frequency of leaky integrate-and-fire neurons with synaptic current dynamics. *J. Theor. Biol.* 195, 87–95. doi: 10.1006/jtbi.1998.0782
- Cadiou, C. F., Hong, H., Yamins, D., Pinto, N., Majaj, N. J., and DiCarlo, J. J. (2013). *The Neural Representation Benchmark and its Evaluation on Brain and Machine. arXiv:13013530 [csNE]*.
- Chapin, J. K., and Nicolelis, M. A. (1999). Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. *J. Neurosci. Methods* 94, 121–140. doi: 10.1016/S0165-0270(99)00130-2
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., and Shenoy, K. V. (2010a). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* 68, 387–400. doi: 10.1016/j.neuron.2010.09.015
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., et al. (2010b). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378. doi: 10.1038/nn.2501
- Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G., and Shenoy, K. V. (2006). Neural variability in premotor cortex provides a signature of motor preparation. *J. Neurosci.* 26, 3697–3712. doi: 10.1523/JNEUROSCI.3762-05.2006
- Cohen, M. R., and Maunsell, J. H. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 12, 1594–1600. doi: 10.1038/nn.2439
- Curti, E., Mongillo, G., La Camera, G., and Amit, D. J. (2004). Mean field and capacity in realistic networks of spiking neurons storing sparsely coded random memories. *Neural Comput.* 16, 2597–2637. doi: 10.1162/089976604.2321805
- de Araujo, I. E., Gutierrez, R., Oliveira-Maia, A. J., Pereira, A., Nicolelis, M. A., and Simon, S. A. (2006). Neural ensemble coding of satiety states. *Neuron* 51, 483–494. doi: 10.1016/j.neuron.2006.07.009
- Deco, G., and Hugues, E. (2012). Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Comput. Biol.* 8:e1002395. doi: 10.1371/journal.pcbi.1002395
- DiCarlo, J. J., Zoccolan, D., and Rust, N. C. (2012). How does the brain solve visual object recognition? *Neuron* 73, 415–434. doi: 10.1016/j.neuron.2012.01.010
- Durstewitz, D., Vitz, N. M., Floresco, S. B., and Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron* 66, 438–448. doi: 10.1016/j.neuron.2010.03.029
- Escola, S., Fontanini, A., Katz, D., and Paninski, L. (2011). Hidden markov models for the stimulus-response relationships of multistate neural systems. *Neural Comput.* 23, 1071–1132. doi: 10.1162/NECO_a_00118
- Farkhooi, F., Froese, A., Muller, E., Menzel, R., and Nawrot, M. P. (2013). Cellular adaptation facilitates sparse and reliable coding in sensory pathways. *PLoS Comput. Biol.* 9:e1003251. doi: 10.1371/journal.pcbi.1003251
- Fontanini, A., and Katz, D. B. (2005). 7 to 12 Hz activity in rat gustatory cortex reflects disengagement from a fluid self-administration task. *J. Neurophysiol.* 93, 2832–2840. doi: 10.1152/jn.01035.2004
- Fontanini, A., and Katz, D. B. (2006). State-dependent modulation of time-varying gustatory responses. *J. Neurophysiol.* 96, 3183–3193. doi: 10.1152/jn.00804.2006
- Fontanini, A., and Katz, D. B. (2008). Behavioral states, network states, and sensory response variability. *J. Neurophysiol.* 100, 1160–1168. doi: 10.1152/jn.90592.2008
- Fourcaud, N., and Brunel, N. (2002). Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Comput.* 14, 2057–2110. doi: 10.1162/089976602320264015
- Fusi, S., and Mattia, M. (1999). Collective behavior of networks with linear (VLSI) integrate-and-fire neurons. *Neural Comput.* 11, 633–652. doi: 10.1162/089976699300016601

- Ganguli, S., Bisley, J. W., Roitman, J. D., Shadlen, M. N., Goldberg, M. E., and Miller, K. D. (2008). One-dimensional dynamics of attention and decision making in LIP. *Neuron* 58, 15–25. doi: 10.1016/j.neuron.2008.01.038
- Ganguli, S., and Sompolinsky, H. (2012). Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Ann. Rev. Neurosci.* 35, 485–508. doi: 10.1146/annurev-neuro-062111-150410
- Gao, P., and Ganguli, S. (2015). On simplicity and complexity in the brave new world of large-scale neuroscience. *Curr. Opin. Neurobiol.* 32, 148–155. doi: 10.1016/j.conb.2015.04.003
- Gardner, M. P., and Fontanini, A. (2014). Encoding and tracking of outcome-specific expectancy in the gustatory cortex of alert rats. *J. Neurosci.* 34, 13000–13017. doi: 10.1523/JNEUROSCI.1820-14.2014
- Geffen, M. N., Broome, B. M., Laurent, G., and Meister, M. (2009). Neural encoding of rapidly fluctuating odors. *Neuron* 61, 570–586. doi: 10.1016/j.neuron.2009.01.021
- Giugliano, M., Darbon, P., Arsiero, M., Luscher, H. R., and Streit, J. (2004). Single-neuron discharge properties and network activity in dissociated cultures of neocortex. *J. Neurophysiol.* 92, 977–996. doi: 10.1152/jn.00067.2004
- Giugliano, M., La Camera, G., Fusi, S., and Senn, W. (2008). The response of cortical neurons to *in vivo*-like input current: theory and experiment. II. Time-varying and spatially distributed inputs. *Biol. Cybern.* 99, 303–318. doi: 10.1007/s00422-008-0270-9
- Grossman, S. E., Fontanini, A., Wieskopf, J. S., and Katz, D. B. (2008). Learning-related plasticity of temporal coding in simultaneously recorded amygdala-cortical ensembles. *J. Neurosci.* 28, 2864–2873. doi: 10.1523/JNEUROSCI.4063-07.2008
- Gur, M., Beylin, A., and Snodderly, D. M. (1997). Response variability of neurons in primary visual cortex (V1) of alert monkeys. *J. Neurosci.* 17, 2914–2920.
- Holland, P. W., and Welsch, R. E. (1977). Robust regression using iteratively reweighted least-squares. *Commun. Statist. Theory Methods A* 6, 813–827. doi: 10.1080/03610927708827533
- Horst, N. K., and Laubach, M. (2013). Reward-related activity in the medial prefrontal cortex is driven by consumption. *Front. Neurosci.* 7:56. doi: 10.3389/fnins.2013.00056
- Inberg, S., Elkobi, A., Edri, E., and Rosenblum, K. (2013). Taste familiarity is inversely correlated with Arc/Arg3.1 hemispheric lateralization. *J. Neurosci.* 33, 11734–11743. doi: 10.1523/JNEUROSCI.0801-13.2013
- Jezzini, A., Mazzucato, L., La Camera, G., and Fontanini, A. (2013). Processing of hedonic and chemosensory features of taste in medial prefrontal and insular networks. *J. Neurosci.* 33, 18966–18978. doi: 10.1523/JNEUROSCI.2974-13.2013
- Jones, L. M., Fontanini, A., Sadacca, B. F., Miller, P., and Katz, D. B. (2007). Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc. Natl. Acad. Sci. U.S.A.* 104, 18772–18777. doi: 10.1073/pnas.0705546104
- Katz, D. B., Simon, S. A., and Nicolelis, M. A. (2001). Dynamic and multimodal responses of gustatory cortical neurons in awake rats. *J. Neurosci.* 21, 4478–4489.
- Kemere, C., Santhanam, G., Yu, B. M., Afshar, A., Ryu, S. I., Meng, T. H., et al. (2008). Detecting neural-state transitions using hidden Markov models for motor cortical prostheses. *J. Neurophysiol.* 100, 2441–2452. doi: 10.1152/jn.00924.2007
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., and Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956. doi: 10.1038/nature02078
- Kiani, R., Cueva, C. J., Reppas, J. B., Peixoto, D., Ryu, S. I., and Newsome, W. T. (2015). Natural grouping of neural responses reveals spatially segregated clusters in primate cortex. *Neuron* 85, 1359–1373. doi: 10.1016/j.neuron.2015.02.014
- La Camera, G., Giugliano, M., Senn, W., and Fusi, S. (2008). The response of cortical neurons to *in vivo*-like input current: theory and experiment. I. Noisy inputs with stationary statistics. *Biol. Cybern.* 99, 279–301. doi: 10.1007/s00422-008-0272-7
- La Camera, G., Rauch, A., Lüscher, H. R., Senn, W., and Fusi, S. (2004). Minimal models of adapted neuronal response to *in vivo*-like input currents. *Neural Comput.* 16, 2101–2124. doi: 10.1162/0899766041732468
- La Camera, G., Rauch, A., Thurbon, D., Lüscher, H. R., Senn, W., and Fusi, S. (2006). Multiple time scales of temporal response in pyramidal and fast spiking cortical neurons. *J. Neurophysiol.* 96, 3448–3464. doi: 10.1152/jn.00453.2006
- Lánský, P., and Sato, S. (1999). The stochastic diffusion models of nerve membrane depolarization and interspike interval generation. *J. Peripher. Nerv. Syst.* 4, 27–42.
- Litwin-Kumar, A., and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* 15, 1498–1505. doi: 10.1038/nn.3220
- Luczak, A., Bartho, P., and Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62, 413–425. doi: 10.1016/j.neuron.2009.03.014
- Macke, J. H., Berens, P., Ecker, A. S., Tolias, A. S., and Bethge, M. (2009). Generating spike trains with specified correlation coefficients. *Neural Comput.* 21, 397–423. doi: 10.1162/neco.2008.02-08-713
- Mardia, K. V., Kent, J. T., and Bibby, J. M. (1979). *Multivariate Analysis*. London: Academic Press.
- Martignon, L., Deco, G., Laskey, K., Diamond, M., Freiwald, W., and Vaadia, E. (2000). Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput.* 12, 2621–2653. doi: 10.1162/089976600300014872
- Mascaro, M., and Amit, D. J. (1999). Effective neural response function for collective population states. *Network* 10, 351–373. doi: 10.1088/0954-898X/10_4_305
- Mazzucato, L., Fontanini, A., and La Camera, G. (2015). Dynamics of multistable states during ongoing and evoked cortical activity. *J. Neurosci.* 35, 8214–8231. doi: 10.1523/JNEUROSCI.4819-14.2015
- McDonnell, M. D., and Ward, L. M. (2011). The benefits of noise in neural systems: bridging theory and experiment. *Nat. Rev. Neurosci.* 12, 415–426. doi: 10.1038/nrn3061
- Miller, P., and Katz, D. B. (2010). Stochastic transitions between neural states in taste processing and decision-making. *J. Neurosci.* 30, 2559–2570. doi: 10.1523/JNEUROSCI.3047-09.2010
- Montavon, G., Braun, M. L., and Müller, K.-R. (2011). Kernel analysis of deep networks. *J. Mach. Learn. Res.* 12, 2563–2581.
- Moreno-Bote, R. (2014). Poisson-like spiking in circuits with probabilistic synapses. *PLoS Comput. Biol.* 10:e1003522. doi: 10.1371/journal.pcbi.1003522
- Nienborg, H., Cohen, M. R., and Cumming, B. G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Ann. Rev. Neurosci.* 35, 463–483. doi: 10.1146/annurev-neuro-062111-150403
- Phillips, M. I., and Norgren, R. (1970). A rapid method for permanent implantation of an intraoral fistula in rats. *Behav. Res. Methods Instrum.* 2:124. doi: 10.3758/BF03211020
- Ponce-Alvarez, A., Nácher, V., Luna, R., Riehle, A., and Romo, R. (2012). Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. *J. Neurosci.* 32, 11956–11969. doi: 10.1523/JNEUROSCI.6176-11.2012
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical Recipes the Art of Scientific Computing, 3rd Edn*. Cambridge: New York, NY: Cambridge University Press.
- Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* 77, 257–286. doi: 10.1109/5.18626
- Rauch, A., La Camera, G., Lüscher, H. R., Senn, W., and Fusi, S. (2003). Neocortical pyramidal cells respond as integrate-and-fire neurons to *in vivo*-like input currents. *J. Neurophysiol.* 90, 1598–1612. doi: 10.1152/jn.00293.2003
- Renart, A., Brunel, N., and Wang, X.-J. (2004). “Mean-field theory of recurrent cortical networks: from irregularly spiking neurons to working memory,” in *Computational Neuroscience: A Comprehensive Approach*, ed J. Feng (Boca Raton, FL: CRC), 431–490.
- Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., et al. (2010). The asynchronous state in cortical circuits. *Science* 327, 587–590. doi: 10.1126/science.1179850
- Richardson, M. J. (2004). The effects of synaptic conductance on the voltage distribution and firing rate of spiking neurons. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 69:051918. doi: 10.1103/PhysRevE.69.051918
- Rickert, J., Riehle, A., Aertsen, A., Rotter, S., and Nawrot, M. P. (2009). Dynamic encoding of movement direction in motor cortical neurons. *J. Neurosci.* 29, 13870–13882. doi: 10.1523/JNEUROSCI.5441-08.2009

- Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., et al. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585–590. doi: 10.1038/nature12160
- Sadacca, B. F., Rothwax, J. T., and Katz, D. B. (2012). Sodium concentration coding gives way to evaluative coding in cortex and amygdala. *J. Neurosci.* 32, 9999–10011. doi: 10.1523/JNEUROSCI.6059-11.2012
- Samuelsen, C. L., Gardner, M. P., and Fontanini, A. (2012). Effects of cue-triggered expectation on cortical processing of taste. *Neuron* 74, 410–422. doi: 10.1016/j.neuron.2012.02.031
- Schneidman, E., Berry, M. J. II., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440, 1007–1012. doi: 10.1038/nature04701
- Seidemann, E., Meilijson, I., Abeles, M., Bergman, H., and Vaadia, E. (1996). Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J. Neurosci.* 16, 752–768.
- Shadlen, M. N., and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4, 569–579. doi: 10.1016/0959-4388(94)90059-0
- Stauder, B., Rotter, S., and Grün, S. (2010). CuBIC: cumulant based inference of higher-order correlations in massively parallel spike trains. *J. Comput. Neurosci.* 29, 327–350. doi: 10.1007/s10827-009-0195-x
- Tkacik, G., Prentice, J. S., Balasubramanian, V., and Schneidman, E. (2010). Optimal population coding by noisy spiking neurons. *Proc. Natl. Acad. Sci. U.S.A.* 107, 14419–14424. doi: 10.1073/pnas.1004906107
- Tsodyks, M., Kenet, T., Grinvald, A., and Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286, 1943–1946. doi: 10.1126/science.286.5446.1943
- Tuckwell, H. C. (1988). *Introduction to Theoretical Neurobiology*. Cambridge, UK: Cambridge University Press.
- van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726. doi: 10.1126/science.274.5293.1724
- van Vreeswijk, C., and Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput.* 10, 1321–1371. doi: 10.1162/089976698300017214
- White, B., Abbott, L. F., and Fiser, J. (2012). Suppression of cortical neural variability is stimulus- and state-dependent. *J. Neurophysiol.* 108, 2383–2392. doi: 10.1152/jn.00723.2011
- Yamamoto, T., Matsuo, R., Kiyomitsu, Y., and Kitamura, R. (1988). Sensory inputs from the oral region to the cerebral cortex in behaving rats: an analysis of unit responses in cortical somatosensory and taste areas during ingestive behavior. *J. Neurophysiol.* 60, 1303–1321.
- Yamamoto, T., Yuyama, N., and Kawamura, Y. (1981). Cortical neurons responding to tactile, thermal and taste stimulations of the rat's tongue. *Brain Res.* 221, 202–206. doi: 10.1016/0006-8993(81)91075-1
- Zucchini, W., and MacDonald, I. L. (2009). *Hidden Markov Models for Time Series: An Introduction Using R*. Boca Raton, FL: CRC Press. doi: 10.1201/9781420010893

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

Copyright © 2016 Mazzucato, Fontanini and La Camera. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.