

## RESEARCH ARTICLE

## Estimating the dimensionality of the manifold underlying multi-electrode neural recordings

Ege Altan <sup>1,2\*</sup>, Sara A. Solla<sup>1,3</sup>, Lee E. Miller <sup>1,2,4,5</sup>, Eric J. Perreault <sup>2,4,5</sup>

**1** Department of Neuroscience, Northwestern University, Chicago, Illinois, United States of America, **2** Department of Biomedical Engineering, Northwestern University, Evanston, Illinois, United States of America, **3** Department of Physics and Astronomy, Northwestern University, Evanston, Illinois, United States of America, **4** Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, Illinois, United States of America, **5** Shirley Ryan AbilityLab, Chicago, Illinois, United States of America

\* [egealtan@u.northwestern.edu](mailto:egealtan@u.northwestern.edu) OPEN ACCESS

**Citation:** Altan E, Solla SA, Miller LE, Perreault EJ (2021) Estimating the dimensionality of the manifold underlying multi-electrode neural recordings. *PLoS Comput Biol* 17(11): e1008591. <https://doi.org/10.1371/journal.pcbi.1008591>

**Editor:** Alireza Soltani, Dartmouth College, UNITED STATES

**Received:** December 4, 2020

**Accepted:** November 11, 2021

**Published:** November 29, 2021

**Copyright:** © 2021 Altan et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** All data and code are available from the figshare database at <https://doi.org/10.6084/m9.figshare.13335341>.

**Funding:** This work was supported by grant number NS053603 from the National Institute of Neurological Disorder and Stroke (<https://www.ninds.nih.gov/>) to S.A.S, L.E.M, E.J.P. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** The authors have declared that no competing interests exist.

## Abstract

It is generally accepted that the number of neurons in a given brain area far exceeds the number of neurons needed to carry any specific function controlled by that area. For example, motor areas of the human brain contain tens of millions of neurons that control the activation of tens or at most hundreds of muscles. This massive redundancy implies the covariation of many neurons, which constrains the population activity to a low-dimensional manifold within the space of all possible patterns of neural activity. To gain a conceptual understanding of the complexity of the neural activity within a manifold, it is useful to estimate its dimensionality, which quantifies the number of degrees of freedom required to describe the observed population activity without significant information loss. While there are many algorithms for dimensionality estimation, we do not know which are well suited for analyzing neural activity. The objective of this study was to evaluate the efficacy of several representative algorithms for estimating the dimensionality of linearly and nonlinearly embedded data. We generated synthetic neural recordings with known intrinsic dimensionality and used them to test the algorithms' accuracy and robustness. We emulated some of the important challenges associated with experimental data by adding noise, altering the nature of the embedding of the low-dimensional manifold within the high-dimensional recordings, varying the dimensionality of the manifold, and limiting the amount of available data. We demonstrated that linear algorithms overestimate the dimensionality of nonlinear, noise-free data. In cases of high noise, most algorithms overestimated the dimensionality. We thus developed a denoising algorithm based on deep learning, the "Joint Autoencoder", which significantly improved subsequent dimensionality estimation. Critically, we found that all algorithms failed when the intrinsic dimensionality was high (above 20) or when the amount of data used for estimation was low. Based on the challenges we observed, we formulated a pipeline for estimating the dimensionality of experimental neural data.

## Author summary

The number of neurons that we can simultaneously record from has increased exponentially for decades; today that number is in the thousands of neurons. However, the individual firing rates are highly redundant. One approach to identifying important features from redundant data is to estimate the dimensionality of the neural recordings, which represents the number of degrees of freedom required to describe the data without significant information loss. Better understanding of dimensionality may also uncover the mechanisms of computation within a neural circuit. Circuits carrying out complex computations might be higher-dimensional than those carrying out simpler computations. Typically, studies have quantified neural dimensionality using one of several available methods reflecting a lack of consensus on which method would be most appropriate for neural data. In this work, we used several methods to investigate the accuracy of simulated neural data with properties mimicking those of actual neural recordings. Based on these results, we devised an analysis pipeline to estimate the dimensionality of neural recordings. Our work will allow scientists to extract informative features from a large number of highly redundant neurons, as well as quantify the complexity of information encoded by these neurons.

## Introduction

Studies that simultaneously record the activity of many neurons have shown that cortical neural activity is highly redundant [1]. In primary motor cortex (M1), redundancy arises as tens of millions of neurons control tens or at most hundreds of muscles. This redundancy implies significant covariation in the activity of many neurons, which confines the population neural activity to a low-dimensional manifold embedded in the neural space of all possible patterns of neural population activity [2–9]. Low-dimensional manifolds have also been observed in a variety of other cortical regions [10–18]. Reliable algorithms for identifying these manifolds and characterizing their dimensionality are increasingly important as our ability to record from large populations of neurons increases [19]. The dimensionality of the manifold describing the coordinated firing of a set of neurons quantifies the number of degrees of freedom needed to describe population activity without significant information loss [20,21]. Projecting the observed firing patterns onto the manifold yields a low-dimensional set of latent signals that can simplify the interpretation of population neural activity [2,9,22]. Low-dimensional latent signals can facilitate the manipulation or the extraction of signals for brain-computer interfaces, a rehabilitative technology that converts neural signals into control commands to restore movement to paralyzed patients [23,24].

Unfortunately, it is surprisingly difficult to estimate the dimensionality of neural manifolds, particularly in the realistic condition of a noisy, nonlinear embedding. There is evidence of a nonlinear mapping between the recorded neural activity and the associated low-dimensional latent signals [10,25–27]. Noise propagates from the level of sensory transduction and amplification, the opening and closing of voltage-gated ion channels, and builds up at the level of synapses, causing neural firing to be a stochastic process [28]. The two effects, nonlinearity and noise, combine to pose significant challenges to existing dimensionality estimation algorithms. The accuracy of the estimators also depends on the amount of available data [29,30], which is limited in most experimental paradigms. If we wish to identify the manifolds associated with experimentally measured neural activity, we need methods that are robust in the presence of these challenges.

The methods that have been proposed for estimating the dimensionality of neural manifolds can be broadly categorized into linear or nonlinear algorithms, based on assumptions about the nature of the mapping between the low-dimensional representation of the latent signals and the high-dimensional space of neural activity. The most commonly used linear method for dimensionality reduction is Principal Component Analysis (PCA), based on identifying mutually orthogonal directions in the empirical neural space of recorded activity; these directions are monotonically associated with the largest data variance. PCA provides a hierarchical description in which the data projected onto the manifold subtended by the leading principal components become closer and closer to the recorded data as the dimensionality of the linear manifold is increased towards the dimensionality of the empirical neural space. Although PCA provides a useful and systematic tool for variance-based dimensionality reduction, it does not specify how to uniquely identify the dimensionality of the manifold: the typical implementation requires the choice of an arbitrary variance threshold. Other PCA-based algorithms such as Participation Ratio (PR) [5,18] and Parallel Analysis (PA) [31,32] provide more principled prescriptions for linear dimensionality estimation, by incorporating criteria for determining an optimal number of leading principal components to use when constructing the low-dimensional manifold.

Linear dimensionality estimation algorithms may work well for linear datasets, but are likely to overestimate the dimensionality of a manifold arising from a nonlinear mapping between the low- and high-dimensional spaces [20,21,33,34]. In contrast, nonlinear methods (e.g., Correlation Dimension [35–37], Levina-Bickel Maximum Likelihood Estimation [38], Two Nearest Neighbors [39], and Fisher Separability Analysis [40]) may provide accurate dimensionality estimates for both linearly and nonlinearly embedded data.

Most dimensionality estimation methods have been tested in the absence of noise even though it is known that linear and nonlinear methods overestimate dimensionality when the data is noisy [20]. The robustness of dimensionality estimation algorithms to noise remains to be characterized.

The objective of this study was to characterize the accuracy of several dimensionality estimation algorithms when applied to high-dimensional recordings of neural activity. We evaluated previously proposed algorithms on synthetic datasets of known dimensionality to identify conditions under which each method succeeded and/or failed. Specifically, we evaluated how the algorithms handled the nature of the embedding (linear or nonlinear), the amount of noise added to the simulated neural data, and the amount of data available. We found increasing levels of noise to be a challenge for all tested algorithms. We therefore also evaluated different approaches for reducing noise prior to performing dimensionality estimation, including the “Joint Autoencoder”, a method we developed based on deep learning techniques. Together, our results allowed us to propose a methodological pipeline for estimating the intrinsic dimensionality of high-dimensional datasets of recorded neural activity.

## Methods

### Ethics statement

All surgical and experimental procedures that yielded the multi-electrode array recordings from non-human primates [41], which formed the basis of our simulated neural signals, were approved by Institutional Animal Care and Use Committee (IACUC) of Northwestern University. The subject was monitored daily. The subject’s diet consisted of standard laboratory animal diet, fresh fruits, and vegetables, and was provided with access to various types of enrichment.

## Simulation of neural signals

We generated the synthetic data used to evaluate the various dimensionality estimation algorithms as follows. First, we created  $d$  signals by randomly selecting ( $d \times M$ ) samples from an empirical distribution of firing rates that we obtained from multi-electrode array recordings of neural activity in the macaque primary motor cortex (M1) made while the monkey was performing a center-out task [41]. The firing rates were binned at 50 ms. The sampling was done randomly across all recorded neurons and time bins within successful trials. Our goal was to generate  $M$  samples of  $d$ -dimensional latent variables that were uncorrelated with each other and individually uncorrelated over time; we verified that these randomly selected signals were indeed uncorrelated, as intended. These signals provided a set of variables of known dimension  $d$  that preserved the first-order firing statistics of the neural activity recorded in M1. Our procedure aimed at generating simulated data that reproduces possible states of activity of a neural population, without considering the order in which these states might be visited; in other words, we focused on population statistics as opposed to population dynamics.

These signals provided a  $d$ -dimensional latent set used to construct synthetic high-dimensional data sets (Fig 1). We allowed  $d$  to vary from 3 to 40. For the analyses where a fixed value of  $d$  was used we chose  $d = 6$ , to approximate the characteristics of real data collected in our laboratory (S1 Table). The  $d$ -dimensional latent signals were first smoothed using a Gaussian kernel (s.d.: 50 ms), and then multiplied by a  $N \times d$  mixing matrix  $W$  with entries that were randomly selected from a zero-mean Gaussian distribution with unit variance. This resulted in a dataset  $X$  composed of  $M$  samples, each of them  $N$ -dimensional. We chose  $N = 96$  to reproduce the number of signals recorded by the multi-electrode array used to obtain the original experimental data. The activity in each of the  $N = 96$  simulated channels was scaled to the range from zero to one to compensate for variability in firing rates across neurons and across time. The effect of non-uniform firing rate variances across channels was considered separately (see “Effect of non-uniform variances across channels” in Results).

A nonlinear embedding was implemented by processing each simulated neural recording in  $X$  with an exponential activation function:

$$f(X) = \frac{e^{\alpha X} - 1}{e^{\alpha} - 1} \quad (\text{Eq 1})$$

The choice of an exponential nonlinearity was based on results from Generalized Linear Models, for which the statistics of the modeled variable determines the nonlinear link function [42]. In our case, the variables of interest are spike counts. Under the assumption that these variables follow a Poisson-like distribution, the appropriate choice of link function is the logarithm [42]. The inverse of the link function, the exponential, is the appropriate nonlinear function for relating the linear combination of explanatory covariates, the latent signals, to the variables of interest, the simulated firing rates. The exponential activation function used in our simulations allowed us to control the degree of nonlinearity by varying the single parameter  $\alpha$ , and to ensure that the range of the nonlinearly embedded synthetic data remained between zero and one. Finally, we added independent Gaussian noise to each of the channels in  $X$ , to generate signals with known signal-to-noise ratio. This choice of noise model provides a simple and widely used mechanism for simulating stochastic processes [43].

The various steps in this procedure allowed us to generate datasets of known intrinsic dimensionality, embedding type (linear/nonlinear), and signal-to-noise ratio.



**Fig 1. Generation of simulated datasets.** First, latent neural signals were obtained by randomly sampling the firing rates of primary motor cortical recordings. The number of latent signals determined the intrinsic dimensionality of the dataset. Then, the dimensionality of the dataset was increased through linear combinations effected by multiplication with a weight matrix  $W$ . The entries of  $W$  were sampled from a zero-mean Gaussian distribution with unit variance. The resulting signals were then scaled to the  $[0,1]$  range by dividing them by their maximum value. This procedure yielded noise-free, linear datasets. In nonlinear simulations, the signals were activated nonlinearly using the exponential function in Eq 1 (red box in diagram). In noisy simulations, zero-mean Gaussian noise with variance specified by a predetermined signal-to-noise ratio was added to the signals. This procedure yielded linear or nonlinear, noisy datasets with known signal-to-noise ratio.

<https://doi.org/10.1371/journal.pcbi.1008591.g001>

## Dimensionality estimation algorithms

We evaluated two classes of dimensionality estimation algorithms, those that assumed a linear embedding and those that also allowed for a nonlinear embedding.

**Linear algorithms.** Linear algorithms map high-dimensional data to a lower dimensional, linear subspace. Principal Component Analysis (PCA) is often used for linear dimensionality estimation in neuroscience [2,4,7,41,44,45]. All the linear algorithms that we tested (summarized below) are based on PCA but use different criteria for dimensionality estimation.

**Principal Component Analysis with a variance cutoff.** PCA creates a low-dimensional representation of the data by sequentially finding orthogonal directions that explain the most remaining variance. Unit vectors that identify those directions, the PCA eigenvectors  $\{v_i\}$ , provide an orthonormal basis for the  $N$ -dimensional data space. The eigenvectors are labeled in decreasing order of the variance associated with each direction, given by the eigenvalues  $\{\lambda_i\}$ . The simplest way to use PCA for dimensionality estimation is to find the number of principal components required to reach a predetermined threshold of cumulative variance. The selection of a variance threshold can be rather arbitrary, and a range of thresholds have been used in the literature. In this study, we used a threshold of 90%, which yielded accurate estimates of dimensionality for the noise-free linear datasets.

**Participation Ratio (PR).** This approach provides a principled way of finding a variance threshold when the ground truth is not known [5,18]. PR uses a simple formula based on the eigenvalues:

$$PR = \frac{(\sum_{i=1}^N \lambda_i)^2}{\sum_{i=1}^N (\lambda_i)^2} \quad (\text{Eq 2})$$

If the leading eigenvalue carries all the variance ( $\lambda_i \neq 0$  for  $i = 1$  and  $\lambda_i = 0$  for all  $i \geq 2$ ), then  $PR = 1$ . At the other extreme, if all eigenvalues are equal, the variance is spread evenly across all the dimensions, and  $PR = N$ . The actual value of PR interpolates between these two extreme conditions to estimate the intrinsic dimensionality, and thus the number of principal components to be kept [5].

**Parallel Analysis (PA).** Much like the Participation Ratio, Parallel Analysis is a principled approach to finding a variance threshold [31,32]. Parallel Analysis generates null distributions for the eigenvalues by repeatedly shuffling each of the  $N$  dimensions of the data separately. The shuffling step ensures that the remaining correlations across the different dimensions of the data are due to chance. We repeated the shuffling procedure 200 times, resulting in a null distribution for each eigenvalue based on 200 samples. The eigenvalues that exceeded the 95<sup>th</sup> percentile of their null distribution were identified as significant; the number of significant eigenvalues determined the number of dimensions to be kept. Although this method has not

been directly applied to neural data, similar approaches based on finding null distributions of eigenvalues have been used for neural dimensionality estimation [46].

**Nonlinear algorithms.** Nonlinear algorithms can in principle estimate the dimensionality of either linearly or nonlinearly embedded data. Unlike the linear algorithms we tested, the nonlinear algorithms need not rely on a global model for the probability distribution from which the data are assumed to be drawn (in the case of PCA, the model is a multivariate Gaussian distribution). Instead, many nonlinear algorithms estimate intrinsic dimensionality directly from local geometric properties of the data. Commonly used local properties include distance and separability of each data point relative to its neighbors. Although nonlinear algorithms are not yet much used in neuroscience, they have been used to estimate dimensionality in several other fields that produce high-dimensional datasets [47].

**Correlation Dimension (CD).** Correlation Dimension estimates dimensionality by calculating how the number of data samples that fall within a hypersphere change as a function of its radius. This method, originally developed in 1983 [35], has benefitted from recent efforts to improve computational speed and accuracy [36,37]. Although there are only a few applications of Correlation Dimension analysis to neural data [48,49], it is widely used in other disciplines [36].

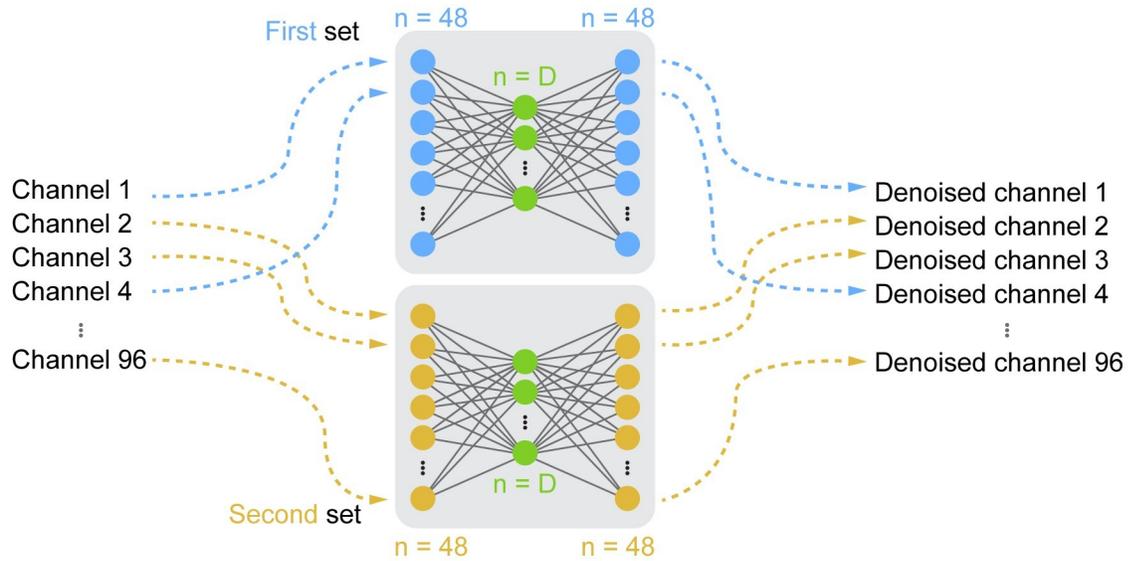
**Levina-Bickel Maximum Likelihood Estimation (LBMLE).** The Levina-Bickel Maximum Likelihood Estimation method [38] is an extension of Correlation Dimension that uses a maximum likelihood approach to estimate distances between data points. This method has been successfully applied to some of the benchmark datasets used in machine learning, such as the Faces [33] and Hands datasets [50].

**Two Nearest Neighbors (TNN).** The Two Nearest Neighbors method also uses the distance between data points to estimate dimensionality [39]. However, unlike Levina-Bickel Maximum Likelihood Estimation, it considers only the first and second neighbors of each point. The ratio of the cumulative distribution of second-neighbor to first-neighbor distances is a function of data dimensionality. By focusing on shorter distances, the method avoids unwanted effects resulting from density changes across the manifold. This method has been successfully applied to synthetic datasets of hyperspheres with known dimensionality [39], and to real-world datasets including molecular simulations [51] and images of hand-written digits [33].

**Fisher Separability Analysis (FSA).** High-dimensional datasets exhibit simple geometric properties such as the likely orthogonality of two randomly picked directions. These properties have recently been characterized as the *blessings of dimensionality* [52], in contrast to the well-known concept of the *curse of dimensionality*. A useful example is the increasing ease with which a hyperplane can separate any given sample in a dataset from all other samples as the dimensionality of the dataset increases. Fisher separability is a computationally efficient, simple, and robust method to assess such separability [53,54]. Dimensionality can be estimated in terms of the probability that a point in the dataset is Fisher separable from the remaining points [40]. The probability distribution of Fisher separability allows the dimensionality of both linear and nonlinear manifolds to be estimated. This method has been applied to study the mutation profiles of the genes resulting in tumors as a means to evaluate therapeutic approaches [55].

## Denoising algorithms

Noise that is uncorrelated across channels will lead to dimensionality estimates that approach the number of channels as the level of noise increases. To mitigate this overestimation problem, we implemented two approaches to denoise neural data. Both rely on an initial estimate



**Fig 2. Architecture of the Joint Autoencoder.** Channels of the 96-dimensional simulated datasets were randomly partitioned into two sets of signals (blue and yellow). Each 48-dimensional set was reconstructed through the corresponding  $D$ -dimensional subspace,  $Z_1$  and  $Z_2$  (green). The reconstructed outputs of the networks were the denoised channels.

<https://doi.org/10.1371/journal.pcbi.1008591.g002>

of an upper bound dimensionality  $D$ , estimated here by using Parallel Analysis. To quantify the performance of the denoising algorithms, we reported Variance Accounted For (VAF) between the denoised signals and the noise-free signals, the latter providing the ground truth.

**PCA denoising.** The linear approach to denoising was based on PCA. Once the value of  $D$  was determined, we used the  $D$  leading principal components to reconstruct the original data. PCA-based denoising is based on the assumption that most of the noise is relegated to the discarded, low-variance principal components.

**Joint Autoencoder denoising.** We also used a neural network for denoising (Fig 2). For this purpose, we divided the 96-dimensional simulated dataset  $X$  into two 48-dimensional partitions:  $X_1$  and  $X_2$ . These partitions were each mapped by the compressive halves of the respective autoencoders to the  $D$ -dimensional subspaces  $Z_1$  and  $Z_2$ . These compressed subspaces were used to obtain reconstructed versions of  $X_1$  and  $X_2$ , respectively denoted  $\hat{X}_1$  and  $\hat{X}_2$ , using the expansive halves of the corresponding autoencoders. The cost function  $C$  used to train the Joint Autoencoder network not only minimized the reconstruction error for  $X_1$  and  $X_2$ , but also the difference between  $Z_1$  and  $Z_2$ :

$$C = MSE(X_1, \hat{X}_1) + MSE(X_2, \hat{X}_2) + MSE(Z_1, Z_2) \tag{Eq 3}$$

This design assumes that each of the partitions  $X_1$  and  $X_2$  contains the information necessary to robustly identify the underlying  $D$ -dimensional signals  $Z_1$  and  $Z_2$ , but not the independent noise components that will differ between the two partitions. We trained the Joint Autoencoder using the ADAM optimizer with a learning rate  $\eta = 0.001$  and dropout regularization on the input layer with  $p = 0.05$ . The use of Rectified Linear Unit (ReLU) activation functions in all layers ensured that the autoencoder network would both operate on and output non-negative signals while allowing for nonlinear embeddings. Our choice of using the ReLU activation function was motivated by its documented success in modeling a wide variety of nonlinearities for deep learning applications [56,57]. In addition, the strict nonnegativity of the ReLU function mimics that of real neural recordings.

## Statistical analyses

We used Monte Carlo simulations to generate up to 10 replications of synthetic data sets, each corresponding to multi-electrode array recording data from an experimental session. We noted the number of replications ( $n$ ) in the figure captions where applicable. Our choice of the number of replications is reasonable compared to the number of experimental sessions that we would expect to see in experiments with monkeys [41,58,59]. The simulations differed by their random number generator seed, which dictated the pseudorandom sampling procedures required for generating the signals. There were three sampling steps in our simulations (Fig 1). First was the creation of the low-dimensional latent signals, which were sampled from an empirical firing rate distribution. The second was the entries of the mixing matrix  $W$ , which were sampled from a zero-mean Gaussian distribution with unit variance. The third was the additive noise, sampled from a zero-mean Gaussian distribution with variance determined by the specified signal-to-noise ratio. We used bootstrapping with 10,000 iterations to compute the statistic of interest and computed its confidence interval using  $\alpha = 0.05$ . We used Bonferroni correction for multiple comparisons.

## Results

Despite the large number of available algorithms for dimensionality estimation, there has been no systematic study of how well-suited they are for the analysis of neural data. Here we test several representative algorithms on synthetic datasets for which the intrinsic dimensionality is known, to assess their ability to estimate the true dimensionality of the data across a range of simulated conditions relevant to neuroscience. These assessments resulted in a recommended procedural pipeline for estimating the intrinsic dimensionality of a set of neural recordings.

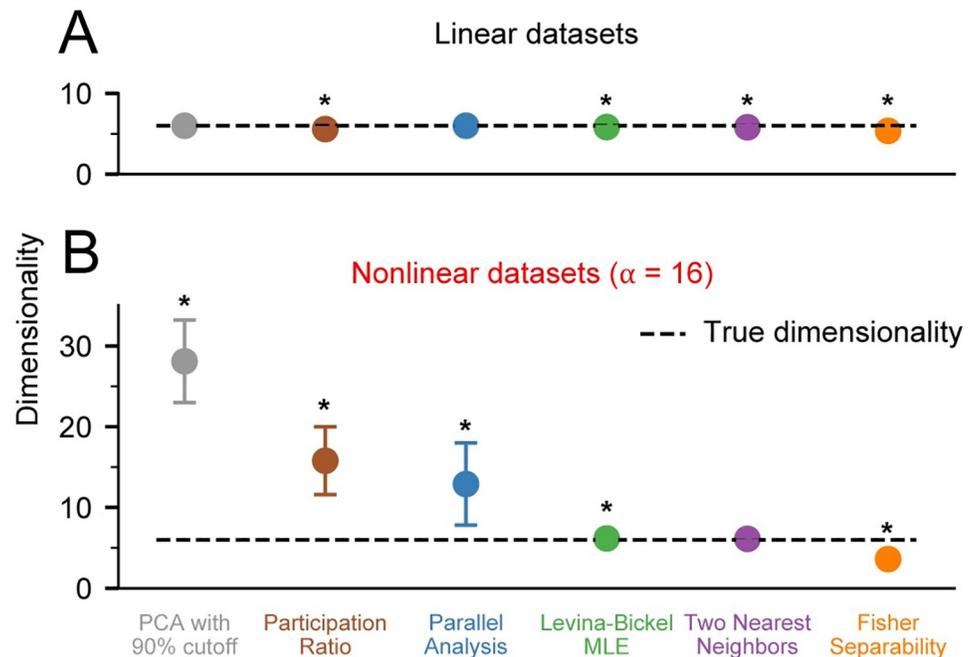
### Dimensionality of noise-free datasets

We first considered the simplest case: how accurately can we determine the dimensionality of linearly embedded, noise-free datasets? To answer this question, we applied the six selected algorithms to datasets with dimensionality  $d = 6$ . We focused on  $d = 6$  as this was the dimensionality estimate of actual multi-electrode array recordings found when using the methods investigated here. In this scenario, all tested linear and nonlinear algorithms estimated the true dimensionality accurately (Fig 3A). Under noise-free conditions, the nonlinear algorithms were as accurate as the linear ones on linearly embedded datasets.

Next, we evaluated all algorithms on nonlinearly embedded noise-free datasets, also for  $d = 6$ . Nonlinearities were introduced as in Eq 1, using  $\alpha = 16$ . In this case, the three linear algorithms dramatically overestimated the true dimensionality, with errors reaching more than 400% of the true value (Fig 3B). In contrast, the nonlinear algorithms performed well; the Levina-Bickel Maximum Likelihood Estimation and the Two Nearest Neighbors methods were more accurate than Fisher Separability Analysis, which slightly underestimated the true dimensionality.

### Effect of non-uniform variances across channels

The normalization of our simulated datasets restricted channel activity to the [0,1] interval, thus imposing a large degree of variance similarity across channels. In contrast, variances of real neural recordings can vary as much as 10-fold from channel to channel. To evaluate the performance of the dimensionality estimation algorithms considered here in the presence of non-uniform variances across channels, we scaled each channel of simulated neural data by a randomly chosen real number between 1 and 10. We found that most algorithms yielded



**Fig 3. Dimensionality of noise free datasets.** A) We applied PCA with 90% variance cutoff (PCA90, gray), Participation Ratio (PR, brown), Parallel Analysis (PA, blue), Levina-Bickel Maximum Likelihood Estimation (LBMLE, green), Two Nearest Neighbors (TNN, purple), and Fisher Separability Analysis (FSA, orange) to linearly embedded,  $d = 6$  datasets ( $n = 10$ ). B) Same as in A, but for nonlinearly embedded datasets. Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant difference of the mean from the true dimensionality of 6 at the significance level of  $\alpha = 0.05$ .

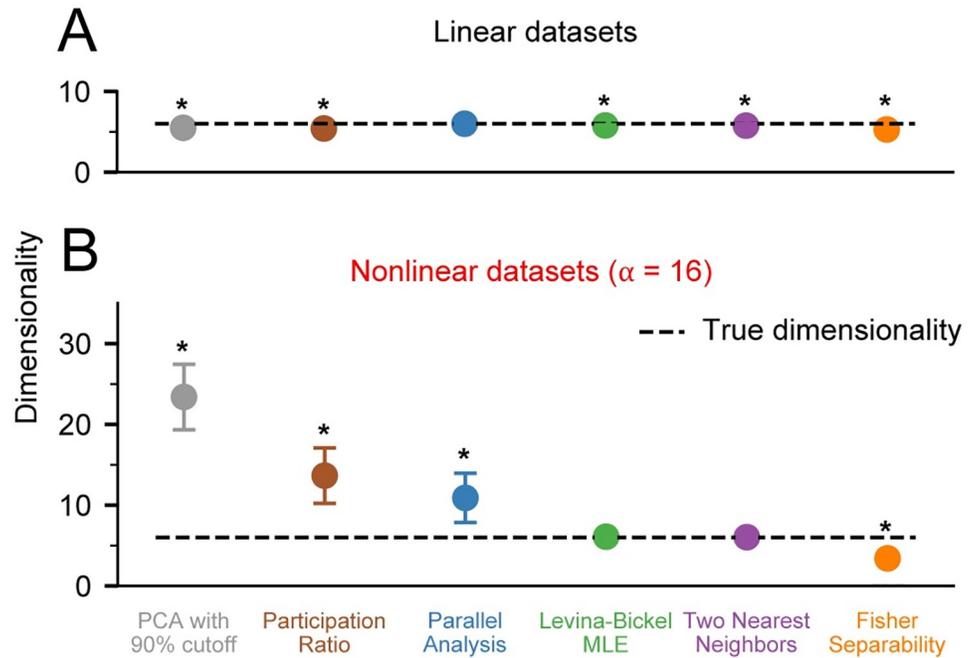
<https://doi.org/10.1371/journal.pcbi.1008591.g003>

lower estimates of dimensionality when applied to the rescaled data in comparison to the estimates obtained when the algorithms were applied to the data before rescaling (Fig 4A and 4B). However, note that both Levina-Bickel Maximum Likelihood Estimation and the Two Nearest Neighbors yielded remarkably accurate dimensionality estimates when applied to rescaled data.

Because of the superior accuracy of Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors, we focused on these two methods for the remainder of the nonlinear analyses. We also retained Parallel Analysis as a benchmark for some of the analyses, as it was the most accurate linear method for estimating the dimensionality of nonlinearly embedded data.

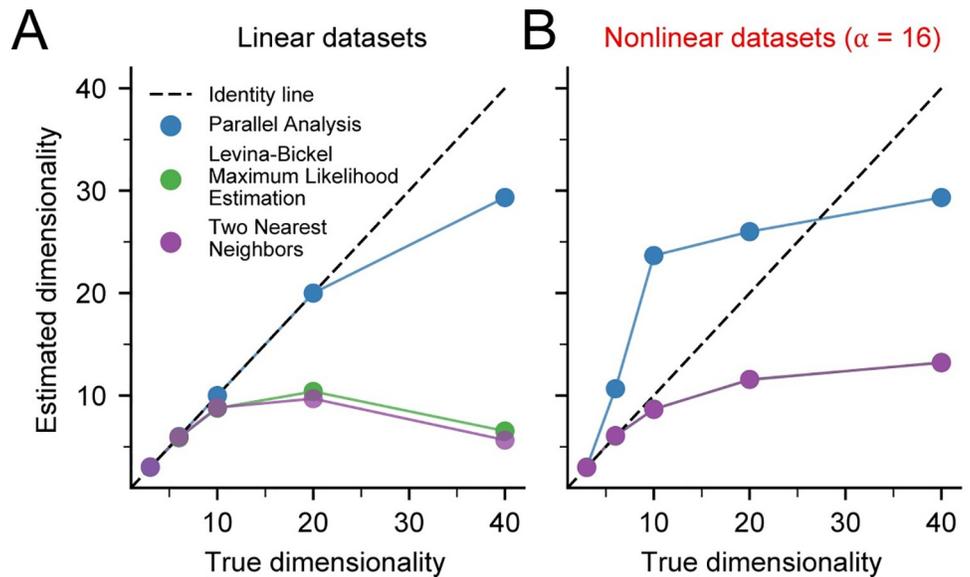
### Effect of true dimensionality on algorithm accuracy

We next evaluated how the true intrinsic dimensionality of the noise-free data influenced algorithm accuracy. Can any intrinsic dimensionality be reliably estimated? We found that the answer is no: the accuracy of all algorithms suffered when the intrinsic dimensionality of the synthetic data was too high. Parallel Analysis was accurate on linear datasets with  $d < 20$ , but inaccurate on nonlinear datasets of all dimensions, as expected (Fig 5). Below about  $d = 6$ , Levina-Bickel Maximum Likelihood Estimation and Two-Nearest Neighbors were accurate on both linear and nonlinear datasets. However, Levina-Bickel Maximum Likelihood Estimation began to underestimate the dimensionality of both linearly embedded (Fig 5A) and nonlinearly embedded (Fig 5B) datasets for  $d > 6$ . This underestimation increased with increasing  $d$ . For nonlinear datasets, the estimate saturated at  $d = 13$ , where underestimation began to get



**Fig 4. Dimensionality of noise free datasets with unequal variance.** A) We applied PCA with 90% variance cutoff (PCA90, gray), Participation Ratio (PR, brown), Parallel Analysis (PA, blue), Levina-Bickel Maximum Likelihood Estimation (LBMLE, green), Two Nearest Neighbors (TNN, purple), and Fisher Separability Analysis (FSA, orange) to linearly embedded,  $d = 6$  datasets ( $n = 10$ ) after randomly scaling each of the  $N = 96$  channels. B) Same as in A, but for nonlinearly embedded datasets. Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of  $\alpha = 0.05$ .

<https://doi.org/10.1371/journal.pcbi.1008591.g004>



**Fig 5. Effect of increasing true dimensionality on dimensionality estimates.** A) The dimensionality of noise free, linear datasets ( $n = 3$ ) was assessed using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Dashed line indicates the identity line. B) Same as A, but for nonlinear datasets. The curve for TNN precisely overlays that of LBMLE, causing it to be obscured.

<https://doi.org/10.1371/journal.pcbi.1008591.g005>

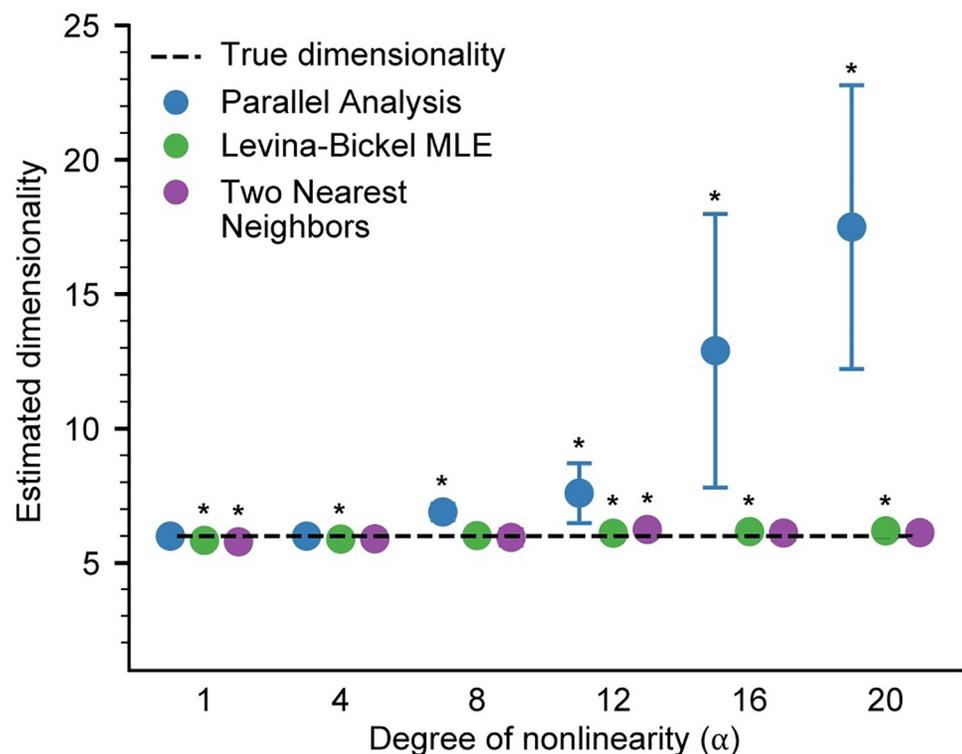
much worse. These results revealed that the intrinsic dimensionality of nonlinearly embedded datasets is hard to estimate reliably when it is large.

### Effect of the level of nonlinearity

We next evaluated how the degree of nonlinearity influenced the accuracy of the dimensionality estimation algorithms. We controlled the degree of nonlinearity by varying the parameter  $\alpha$  in Eq 1; this parameter controls the slope of the exponential activation function used to generate the nonlinearly embedded datasets. We found that both Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors provided accurate dimensionality estimates for all tested levels of nonlinearity (Fig 6). Surprisingly, even Parallel Analysis was accurate up to levels of nonlinearity corresponding to  $\alpha \approx 8$ , where it started to overestimate the intrinsic dimensionality. These results revealed that Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors provide accurate dimensionality estimates for wide levels of nonlinearity, whereas Parallel Analysis is accurate only for low levels of nonlinearity.

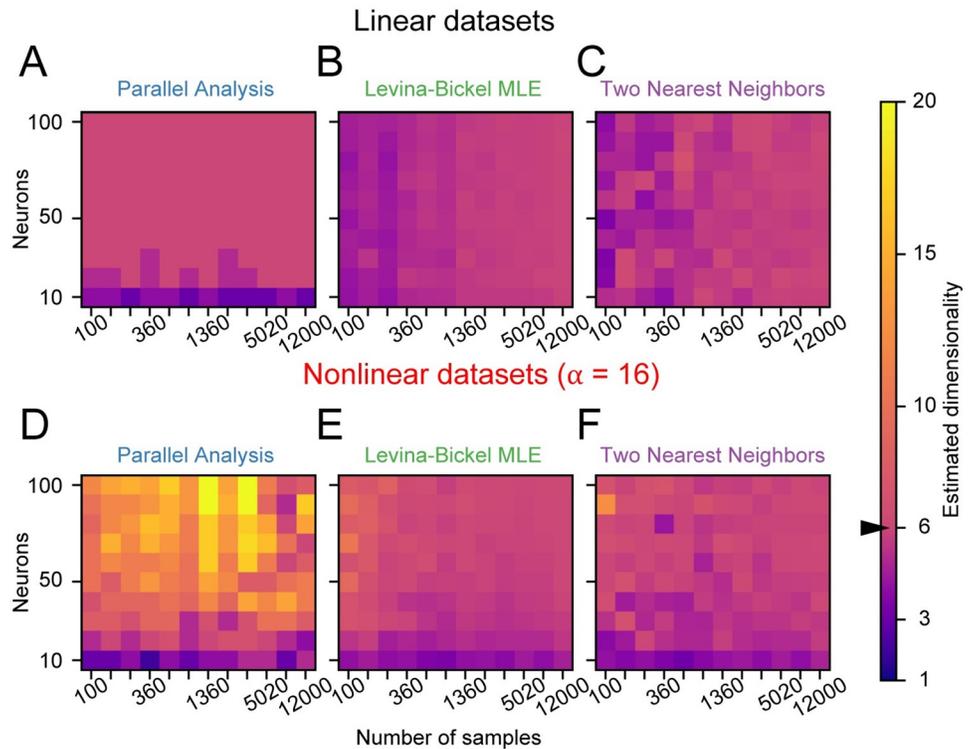
### Amount of data required for estimating dimensionality

Ideally, algorithms would require only small amounts of data, so that the intrinsic dimensionality could be estimated even during transient behaviors and for a small number of recording channels. We thus evaluated the amount of data required to estimate the dimensionality of



**Fig 6. Effect of changing the degree of nonlinearity.** Dimensionality of nonlinear datasets ( $n = 10$ ) with varying levels of nonlinearity controlled by the  $\alpha$  parameter (See Methods), was assessed using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of  $\alpha = 0.05$ .

<https://doi.org/10.1371/journal.pcbi.1008591.g006>



**Fig 7. Amount of data required by dimensionality estimators.** Amount of data required by A) Parallel Analysis (PA), B) Levina-Bickel Maximum Likelihood Estimation (LBMLE), and C) Two Nearest Neighbors (TNN) on linear datasets. Data length was logarithmically scaled between  $M = 100$  and  $M = 12,000$  samples. The correct dimensionality  $d = 6$  is shown in pink. Light colors indicate overestimation and dark colors indicate underestimation of dimensionality. D, E, and F) Same as A, B, and C, respectively, but for nonlinear datasets.

<https://doi.org/10.1371/journal.pcbi.1008591.g007>

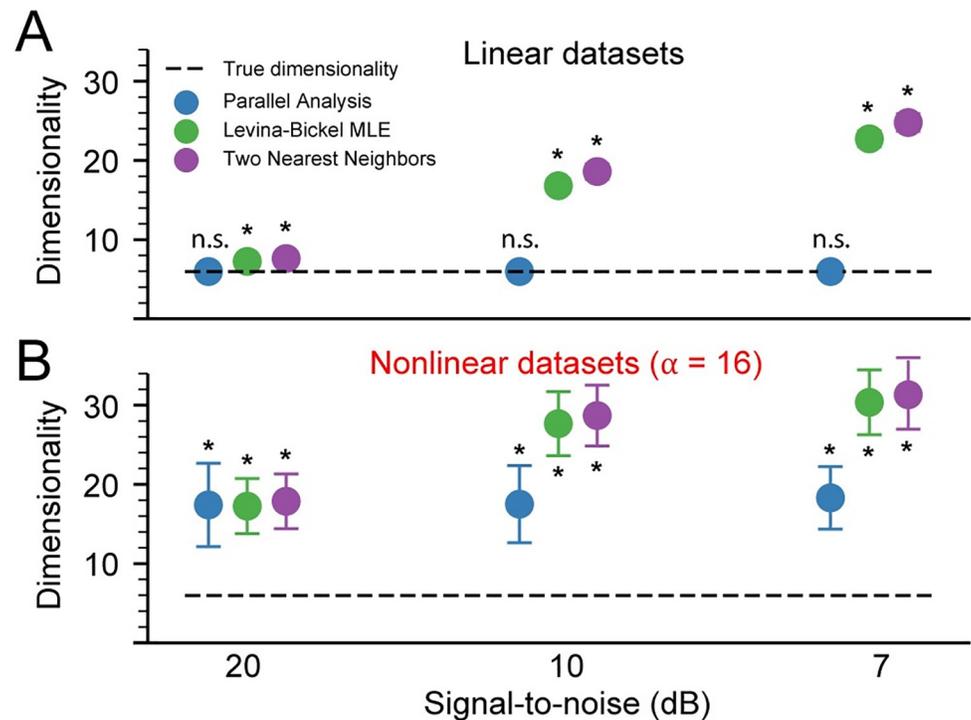
datasets with  $d = 6$ , by varying both the number of samples  $M$  and the number of recording channels  $N$ .

On linear datasets, the accuracy of Parallel Analysis depended only on the number of channels: the algorithm was accurate if 20 or more channels were available (Fig 7A). In contrast, the accuracy of both Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors also depended on the number of samples (Fig 7B and 7C). Around  $M = 600$ , requiring about 30 seconds of data binned at 50 ms, was sufficient for accurate estimates of intrinsic dimensionality using either of these two nonlinear methods.

As expected for highly nonlinear datasets ( $\alpha = 16$ ,  $d = 6$ ), Parallel Analysis was not accurate (Fig 7D) regardless of the amount of data. Both Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors were accurate provided that data from more than 50 channels were available (Fig 7E and 7F). Furthermore, while Levina-Bickel Maximum Likelihood estimation required around 600 samples of data for accurate dimensionality estimates, Two Nearest Neighbors required more than twice as many samples. These results would also depend on the actual dimensionality  $d$ ; here we focused on  $d = 6$ .

### Evaluating and reducing the effects of noise

Any experiment will include some amount of noise in the recorded signals. As expected, all tested algorithms overestimated intrinsic dimensionality in the presence of noise (Fig 8). For any given noise level, estimation errors for the linear datasets (Fig 8A) were a bit smaller than those for the nonlinear datasets (Fig 8B). Adding noise with a power of only 1% of that of the



**Fig 8. Effect of noise on dimensionality estimates.** Estimated dimensionality of linear (A) and nonlinear (B) datasets ( $n = 10$ ) with 20 dB, 10 dB, and 7 dB signal-to-noise ratio was assessed using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of  $\alpha = 0.05$ .

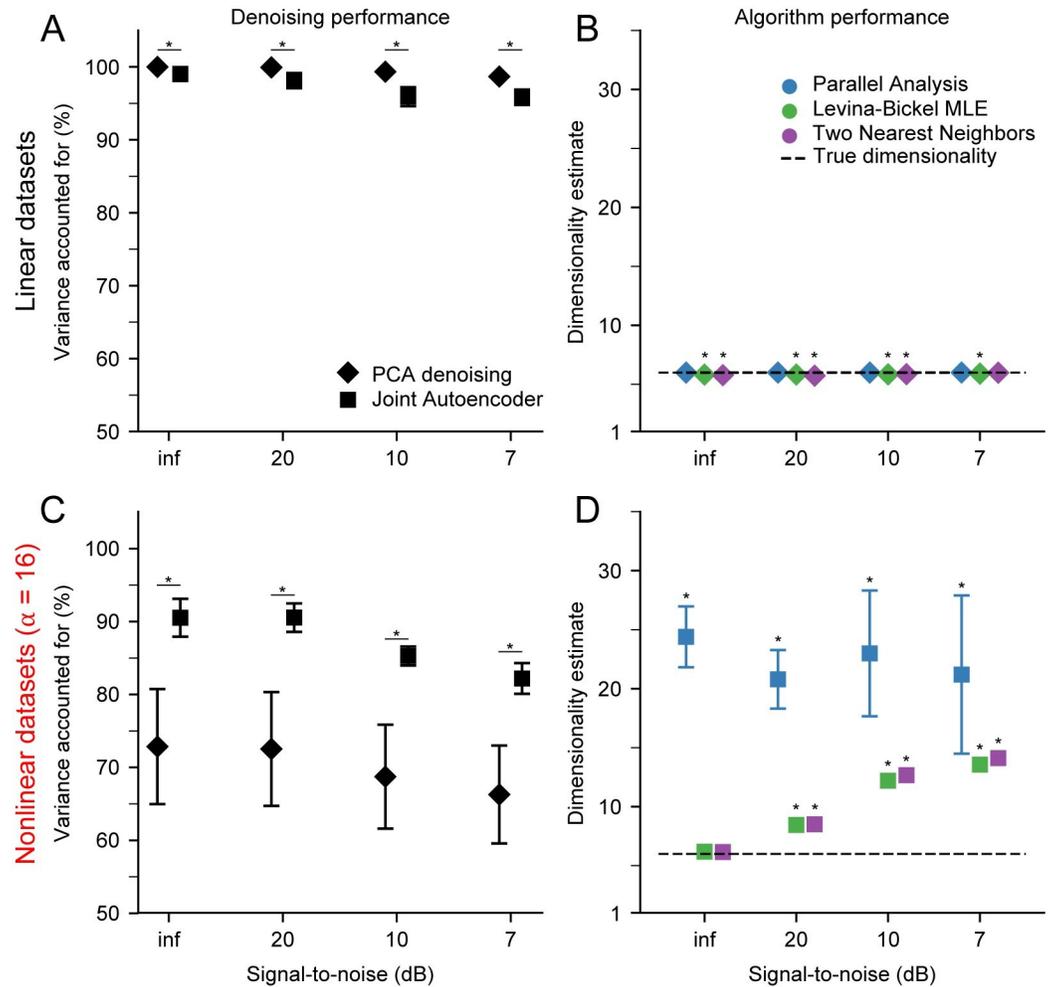
<https://doi.org/10.1371/journal.pcbi.1008591.g008>

signal (SNR = 20 dB) caused Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors to overestimate the dimensionality of the nonlinear data by ~200% (Fig 8B). PA yielded consistent overestimation errors across all nonzero levels of noise for both linear and nonlinear data.

We evaluated two algorithms for mitigating the effects of noise prior to estimating dimensionality: a PCA-based linear method and a Joint Autoencoder nonlinear neural network (see Methods). Both methods were quite effective for denoising the linear datasets (Fig 9A), with the PCA-based approach slightly better than the Joint Autoencoder at the higher noise levels. For linear datasets, dimensionality estimates following PCA-based denoising were highly accurate, yielding correct estimates of the true intrinsic dimension even for high-noise signals (Fig 9B). The Joint Autoencoder was significantly more effective for denoising the nonlinear datasets (Fig 9C). Joint Autoencoder denoising on nonlinear datasets resulted in dimensionality estimates that still increasingly overestimated with increasing noise, but at a much slower rate than without denoising (Fig 9D). The highest noise level we tested (20%; SNR = 7 dB) caused the dimensionality to be overestimated by about 100%. These results were consistent for different degrees of nonlinearity. The more nonlinear the data, the more appropriate it was to use the Joint Autoencoder for denoising.

## Discussion

This study evaluated techniques for estimating the intrinsic dimensionality of high-dimensional neural recordings. We considered representative linear and nonlinear algorithms,



**Fig 9. Performance of PCA and Joint Autoencoder (JAE) denoising algorithms.** A) PCA and JAE denoising applied to linear datasets ( $n = 10$ ) with varying signal-to-noise ratio. Symbols indicate the mean and error bars indicate the standard deviation of the VAF between noise-free and denoised signals. Asterisks indicate significant difference between mean values at the significance level of  $\alpha = 0.05$ . B) Dimensionality estimation on linear datasets after PCA denoising. Dimensionality was estimated using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Symbols indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of  $\alpha = 0.05$ . C) Same as in A, but for nonlinear datasets. D) Same as in B, but for nonlinear datasets after JAE denoising.

<https://doi.org/10.1371/journal.pcbi.1008591.g009>

testing their performance on synthetic datasets that captured properties of neural recordings likely to affect dimensionality estimation. The tested datasets had known intrinsic dimensionality, known levels of noise, and embeddings that were either linear or nonlinear. Our results demonstrated that none of the tested algorithms work for all possible scenarios, but they yielded important insights for when estimates of intrinsic dimensionality are likely to be valid and when they are not. As expected, we found that linear estimation methods are generally not as accurate as nonlinear methods when the mapping between the low-dimensional latent space and the high-dimensional space of neural recordings is nonlinear. Surprisingly, the linear method Parallel Analysis estimated the dimensionality of mildly nonlinear datasets well though it failed for more highly nonlinear embeddings. In contrast, the nonlinear methods

worked well on both linear and highly nonlinear datasets but failed once the intrinsic dimensionality of the data became too high.

Noise was a challenge for all methods, causing dimensionality to be overestimated even for signal-to-noise ratios as low as 20 dB (1% noise variance). We presented two approaches for denoising the data so as to improve the accuracy of the dimensionality estimation. These were a linear PCA-based approach and a novel nonlinear, deep learning approach that we call the Joint Autoencoder. Both denoising approaches attempted to remove signal components that were not shared across the data channels. To achieve this, the PCA-based approach simply removed Principal Components with low variance, whereas the Joint Autoencoder identified an underlying manifold that was common to two randomly sampled sets of channels. Both approaches relied on a linear, upper-bound estimate of the intrinsic dimensionality. Denoising by either method substantially improved subsequent dimensionality estimation, but the Joint Autoencoder was substantially more effective in denoising nonlinear datasets. For linear datasets, dimensionality estimates using Parallel Analysis, Levina-Bickel Maximum Likelihood Estimation, and Two Nearest Neighbors were accurate after PCA-denoising. In the nonlinear case, dimensionality estimates using the two nonlinear methods (LBMLE and TNN) were similarly accurate after JAE-denoising.

### Implications for evaluation of experimental recordings

Due to its computational efficiency and ease of interpretation, most studies have used PCA with an arbitrary variance cutoff to estimate the dimensionality of M1 neural recordings [4,17,41,44,45]. While we have shown that some of the linear methods can be quite effective, simply eliminating non leading PCs based on a cumulative variance cutoff was the least accurate of the algorithms that we tested. Parallel Analysis, the most accurate linear method, performed as well or even better than some of the more advanced and computationally demanding nonlinear methods. Therefore, PA should suffice as a quick and effective approach to estimating dimensionality, even for mildly noisy and nonlinear datasets.

Some of the linear methods used in neuroscience studies rely on a structure of repeated trials in the data [7,46]. These methods use the regularity of repeated trials in a supervised scenario to identify neural dimensions associated with specific experimentally controlled conditions. Such supervised methods cannot be applied to data obtained during non-stereotyped, non-repeating behaviors. All of the methods that we assess in this study are unsupervised and thus applicable to datasets with no repeated trial structure.

Despite the simplicity of linear algorithms, estimating the dimensionality of nonlinear manifolds requires nonlinear algorithms. There is some evidence that neural manifolds may be nonlinear. Recent studies have shown that nonlinear methods for inferring behavioral parameters from M1 neural manifolds are superior to linear methods [60–63]. This suggests that the underlying neural manifold representing motor intent may be nonlinear, and that linear dimensionality estimation methods may be inadequate when estimating the intrinsic dimensionality of primary motor cortical recordings. Studies that investigated the dimensionality of M1 using linear methods most likely overestimated its true intrinsic dimensionality.

Nonlinear algorithms were more accurate than linear algorithms for nonlinear datasets of dimensionality below 10. However, nonlinear methods underestimated dimensionalities above 10. This is a critical concern for experimental recordings, since a low dimensionality estimate from a nonlinear method might be inaccurate if the true dimensionality were large. Multiple studies using linear methods have reported an estimated dimensionality of M1 of around 10 for simple, well-practiced behaviors [5,45,59]. Our results show that linear methods provide an upper bound to the estimate of intrinsic dimensionality as long as the true dimensionality of the data is below 20. If the intrinsic dimensionality of M1 is substantially

higher for more dexterous use of arm and hand than for the scenarios that have typically been studied, the nonlinear methods investigated here may underestimate it.

One method for addressing this concern would be to use nonlinear methods to reduce the dimensionality of a dataset to that of its nonlinear dimensionality estimate, and then to assess the amount of variance that the nonlinear low-dimensional representation captures. If the VAF is high, the data may be truly nonlinearly low dimensional. If, on the other hand, the VAF is low, the true intrinsic dimensionality could be higher than estimated. For the latter case, a practical approach would be to report only the linear dimensionality estimate and emphasize that it only provides an upper bound to the true dimensionality.

We currently lack techniques for reliably assessing datasets with high intrinsic dimensionality, at least when considering practical situations with limited data. There have been some theoretical studies of the amount of data needed for accurate estimation of dimensionality [64,65]. Correlation Dimension, the method on which many nonlinear algorithms are based, requires that the number of data samples  $M$  be on the order of  $10^{d/2}$  [29]. The total amount of data can be increased by either recording from more channels or for a longer duration. Studies that investigated the dimensionality of the primary visual cortex (V1) found that the eigenvalue spectrum of the neural signals obtained from approximately one thousand neurons decayed as a power law [66,67]. These findings would not have been possible if recording from a hundred neurons, which would not have revealed the long, slow-decaying tail of the eigenvalue distribution. One interpretation of these findings is that the linear dimensionality of V1 is arbitrarily large. However, an alternative interpretation is that the neural data are embedded in a very nonlinear manifold, causing the intrinsic dimensionality to be overestimated by the linear methods used in these studies.

The stochastic nature of neural firing and the noise associated with experimental measurements will also cause the intrinsic dimensionality to be overestimated. The two denoising approaches that we presented are simple and effective. Depending on the assumptions about the underlying structure of firing patterns, alternative denoising approaches may be useful. For example, if the temporal relationship between the firing patterns of the population neural activity is of interest, one could use denoising methods that explicitly attempt to model these dynamics, such as Latent Factor Analysis through Dynamical Systems (LFADS), prior to estimating the dimensionality [61].

For the past five decades since the time of Evarts' early experiments [68], assessing the relationship between behaviors and single neuron signals recorded from the brain has been a mainstay of motor systems research. Although the focus of our study was on the dimensionality of the neural manifold to which the population activity is confined, the natural next step in the analysis is to investigate the dynamics of the signals within the neural manifold and their relation to behavior.

### Limitations of the study

While we tried to replicate essential features of experimental data, there are certain characteristics that we did not try to model in our simulations. For example, we only considered additive Gaussian isotropic noise, for simplicity. Experimental recordings might include non-additive, non-isotropic, or non-Gaussian noise. In such cases, PCA may not be an appropriate approach to denoising, even for linearly embedded data. Methods such as Factor Analysis or extensions such as Gaussian-Process Factor Analysis [69], and preprocessing steps such as square-root transforms or pre-whitening could be used instead.

We scaled the firing rates of each channel to be in the [0,1] range. The arbitrary scaling of firing rates provided a simple means for the nonlinear datasets to have the same range as their

linear counterparts, as the nonlinear activation function that we used mapped the  $[0,1]$  range onto itself. However, this modeling restriction does not reflect experimental neural firing data, since the range of neural firing can differ significantly even across neurons of the same type. We have illustrated how heterogeneity in the range of firing rates affects the reliability of dimension estimation algorithms. Soft-normalization approaches that are commonly used in neuroscience (dividing a neuron's firing rate by its range plus a small constant, e.g. [70]), would result in the amplification of signals with low variance and would cause variance-based algorithms to result in higher dimensionality estimates.

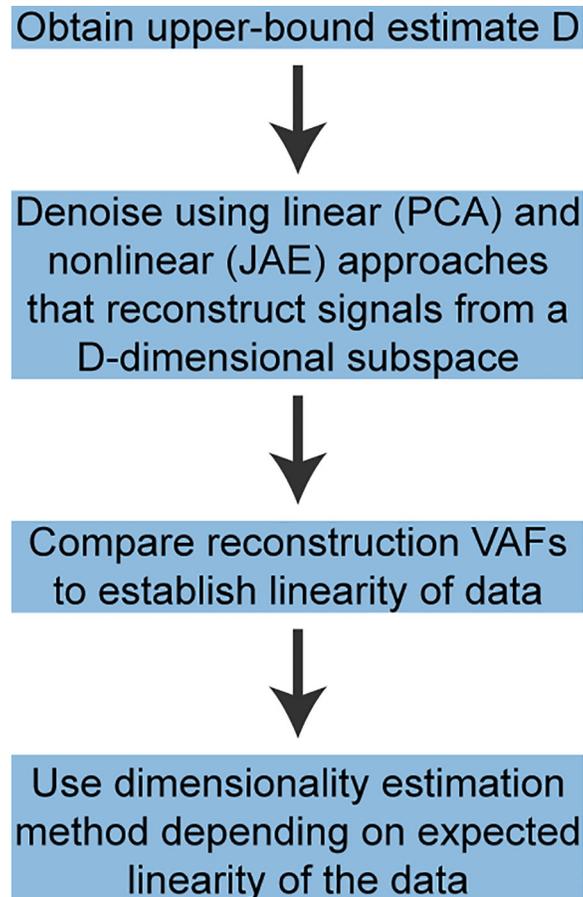
The latent signals used to generate simulated firing rates have the same first-order statistics as the actual data from which they were sampled, but the rescaling of simulated channel activity to a  $[0,1]$  range introduced a departure from realism. This unrealistic scenario is addressed through the random rescaling of individual simulated channels, to reflect the heterogeneity in the range of firing rates observed in actual neural recordings. The latent signals corresponding to the rescaled data no longer share common statistics. This scenario allows us to address an important problem: that low-variance latent signals can be informative [71]. As demonstrated in our study, the use of nonlinear methods for dimensionality estimation ameliorates the problems that arise when neglecting low-variance signals as purely noisy.

### Recommended analysis pipeline

Based on our results, we recommend the following approach for estimating the dimensionality of neural recordings (Fig 10). First, obtain an upper-bound estimate  $D$  of the intrinsic dimensionality of the data. We found that Parallel Analysis works well for this purpose, being both computationally efficient and the most accurate linear method in our tests. Next, the signals should be denoised. Our denoising approach worked by projecting the neural signals into a subspace of dimensionality  $D$  equal to the upper-bound dimensionality estimate, and then reconstructing them based on these projections. A PCA based reconstruction is easy to implement and interpret and may be preferable if computational efficiency is important. A nonlinear denoising algorithm, such as the Joint Autoencoder we proposed, should also be used to assess the degree of nonlinearity of the manifold. The usefulness of the denoising step was quantified through the VAF between the reconstructed signals, assumed to be denoised, and the noise-free synthetic signals before noise was added to them. Our results showed that for nonlinear datasets this VAF was higher for the Joint Autoencoder than it was for PCA. However, this VAF cannot be computed for experimental data, for which we do not have access to the noise-free signals. In this scenario, the reconstruction VAF between noisy inputs and the denoised reconstructed outputs may be useful for detecting nonlinear manifolds: a higher reconstruction VAF for Joint Autoencoder denoising than for PCA denoising would signal a nonlinear manifold. A reconstruction VAF that prefers the Joint Autoencoder indicates that this denoising method yields better denoised signals. Once the signals are denoised, and the linearity or nonlinearity of the manifold is established, either a linear or nonlinear dimensionality estimation method should be used depending on the comparative performance of the corresponding denoising algorithms. The most accurate linear method we tested was Parallel Analysis. Of the nonlinear methods, Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors were the most accurate; Levina-Bickel Maximum Likelihood Estimation required fewer data samples.

### Conclusions

Estimating the dimensionality of neural data is challenging. In this study, we tested several available algorithms and determined the conditions under which estimating dimensionality



**Fig 10. Recommended analysis pipeline for estimating the dimensionality of multi-electrode array recordings.** First, obtain an upper-bound dimensionality estimate  $D$  using a linear algorithm. Parallel Analysis works well for this purpose. Next, denoise the data using both linear (PCA based) and nonlinear (JAE based) denoising approaches and compare their reconstruction VAFs. Higher VAF for the PCA based denoising. Similar VAFs for the linear and nonlinear denoising approaches would signal a linear manifold. In contrast, higher VAF for the JAE based denoising would signal a nonlinear manifold. Finally, once the signals have been denoised using the appropriate denoising method based on the determined linearity of the manifold, estimate the dimensionality of the denoised signals. Parallel Analysis is appropriate for linear manifolds. Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbor are the most accurate nonlinear algorithms that we tested.

<https://doi.org/10.1371/journal.pcbi.1008591.g010>

may be particularly difficult or even impractical. Noise is a confounding factor and must be removed prior to dimensionality estimation. Most existing studies have estimated intrinsic dimensionality using linear methods that are computationally efficient and easy to interpret. We showed that linear methods provide an upper-bound to the intrinsic dimensionality, and in cases of high noise, may even provide better estimates than nonlinear methods, although neither linear nor nonlinear methods will yield accurate estimates in this scenario. Nonlinear algorithms were more accurate for nonlinear datasets when noise was adequately removed. Finally, algorithms failed when the intrinsic dimensionality was high. It may be impractical or impossible to estimate the dimensionality of neural data when it is above  $\sim 20$ . However, estimation of the dimensionality of neural activity in the primary motor cortex may be possible, as many studies have reported its linear dimensionality to be within the practical limits for accurate estimation by the methods we tested.

## Supporting information

**S1 Table. Application of the recommended analysis pipeline to three sets of real neural recordings.** The parallel analysis (PA) estimates of the dimensionality are shown for each of the datasets J1, J2, and J3. These values determined the dimensionality to be used for denoising each dataset. The PCA-based denoising yielded reconstructions with 53%, 51%, and 56% VAF. The JAE-based denoising was slightly better for all datasets, with 61%, 59%, and 62% VAF. The better performance of the JAE-based denoising is indicative of modest nonlinearity in all three datasets. Once each dataset had been denoised using JAE, the corresponding dimensionalities were estimated using MLE and TNN. These results motivated our choice of  $d = 6$  for the intrinsic dimensionality of most of our simulated datasets. (DOCX)

## Acknowledgments

The authors would like to thank Juan Á. Gallego for collecting the experimental data which served as the basis for our simulations.

## Author Contributions

**Conceptualization:** Ege Altan, Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Data curation:** Ege Altan.

**Formal analysis:** Ege Altan.

**Funding acquisition:** Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Investigation:** Ege Altan, Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Methodology:** Ege Altan, Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Project administration:** Eric J. Perreault.

**Resources:** Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Software:** Ege Altan.

**Supervision:** Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Validation:** Ege Altan, Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Visualization:** Ege Altan, Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Writing – original draft:** Ege Altan, Sara A. Solla, Lee E. Miller, Eric J. Perreault.

**Writing – review & editing:** Sara A. Solla, Lee E. Miller, Eric J. Perreault.

## References

1. Saxena S, Cunningham JP. Towards the neural population doctrine. *Current Opinion in Neurobiology*. 2019; 55:103–11. <https://doi.org/10.1016/j.conb.2019.02.002> PMID: 30877963
2. Cunningham JP, Yu BM. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*. 2014; 17(11):1500–9. <https://doi.org/10.1038/nn.3776> PMID: 25151264
3. Elsayed GF, Cunningham JP. Structure in neural population recordings: an expected byproduct of simpler phenomena? *Nature Neuroscience*. 2017; 20(9):1310–8. <https://doi.org/10.1038/nn.4617> PMID: 28783140
4. Gao P, Ganguli S. On Simplicity and Complexity in the Brave New World of Large-Scale Neuroscience. *Current Opinion in Neurobiology*. 2015. <https://doi.org/10.1016/j.conb.2015.04.003> PMID: 25932978

5. Gao P, Trautmann E, Yu B, Santhanam G, Ryu S, Shenoy K, et al. A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*. 2017. <https://doi.org/10.1101/214262>
6. Trautmann EM, Stavisky SD, Lahiri S, Ames KC, Kaufman MT, O'Shea DJ, et al. Accurate Estimation of Neural Population Dynamics without Spike Sorting. *Neuron*. 2019; 103(2):292–308.e4. Epub 2019/06/07. <https://doi.org/10.1016/j.neuron.2019.05.003> PMID: 31171448; PubMed Central PMCID: PMC7002296.
7. Williams AH, Kim TH, Wang F, Vyas S, Ryu SI, Shenoy KV, et al. Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron*. 2018; 98(6):1099–115.e8. <https://doi.org/10.1016/j.neuron.2018.05.015> PMID: 29887338
8. Williamson RC, Doiron B, Smith MA, Yu BM. Bridging large-scale neuronal recordings and large-scale network models using dimensionality reduction. *Current Opinion in Neurobiology*. 2019; 55:40–7. <https://doi.org/10.1016/j.conb.2018.12.009> PMID: 30677702
9. Gallego JA, Perich MG, Miller LE, Solla SA. Neural Manifolds for the Control of Movement. *Neuron*. 2017; 94(5):978–84. Epub 2017/06/09. <https://doi.org/10.1016/j.neuron.2017.05.025> PMID: 28595054; PubMed Central PMCID: PMC6122849.
10. Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, et al. The importance of mixed selectivity in complex cognitive tasks. *Nature*. 2013; 497(7451):585–90. <https://doi.org/10.1038/nature12160> PMID: 23685452
11. Mazor O, Laurent G. Transient Dynamics versus Fixed Points in Odor Representations by Locust Antennal Lobe Projection Neurons. *Neuron*. 2005; 48(4):661–73. <https://doi.org/10.1016/j.neuron.2005.09.032> PMID: 16301181
12. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503(7474):78–84. Epub 2013/11/10. <https://doi.org/10.1038/nature12742> PMID: 24201281; PubMed Central PMCID: PMC4121670.
13. Churchland MM, Abbott LF. Two layers of neural variability. *Nature Neuroscience*. 2012; 15(11):1472–4. <https://doi.org/10.1038/nn.3247> PMID: 23103992
14. Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*. 2012; 484(7392):62–8. <https://doi.org/10.1038/nature10918> PMID: 22419153
15. Warnberg E, Kumar A. Perturbing low dimensional activity manifolds in spiking neuronal networks. *PLoS Comput Biol*. 2019; 15(5):e1007074. Epub 2019/06/01. <https://doi.org/10.1371/journal.pcbi.1007074> PMID: 31150376; PubMed Central PMCID: PMC6586365.
16. Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE. Long-term stability of cortical population dynamics underlying consistent behavior. *Nat Neurosci*. 2020; 23(2):260–70. Epub 2020/01/08. <https://doi.org/10.1038/s41593-019-0555-4> PMID: 31907438; PubMed Central PMCID: PMC7007364.
17. Williamson RC, Cowley BR, Litwin-Kumar A, Doiron B, Kohn A, Smith MA, et al. Scaling Properties of Dimensionality Reduction for Neural Populations and Network Models. *PLoS Comput Biol*. 2016; 12(12):e1005141. Epub 2016/12/08. <https://doi.org/10.1371/journal.pcbi.1005141> PMID: 27926936; PubMed Central PMCID: PMC5142778.
18. Mazzucato L, Fontanini A, La Camera G. Stimuli Reduce the Dimensionality of Cortical Activity. *Front Syst Neurosci*. 2016; 10:11–. <https://doi.org/10.3389/fnsys.2016.00011> PMID: 26924968.
19. Stevenson IH, Kording KP. How advances in neural recording affect data analysis. *Nature Neuroscience*. 2011; 14(2):139–42. <https://doi.org/10.1038/nn.2731> PMID: 21270781
20. Camastra F, Staiano A. Intrinsic dimension estimation: Advances and open problems. *Information Sciences*. 2016; 328:26–41. <https://doi.org/10.1016/j.ins.2015.08.029>
21. Lee JA, Verleysen M. *Nonlinear Dimensionality Reduction*. 1 ed. New York: Springer; 2007. XVII, 309 p.
22. Pang R, Lansdell BJ, Fairhall AL. Dimensionality reduction in neuroscience. *Current Biology*. 2016; 26(14):R656–R60. <https://doi.org/10.1016/j.cub.2016.05.029> PMID: 27458907
23. Pandarinath C, Nuyujukian P, Blabe CH, Sorice BL, Saab J, Willett FR, et al. High performance communication by people with paralysis using an intracortical brain-computer interface. *Elife*. 2017; 6. Epub 2017/02/22. <https://doi.org/10.7554/eLife.18554> PMID: 28220753; PubMed Central PMCID: PMC5319839.
24. Degenhart AD, Bishop WE, Oby ER, Tyler-Kabara EC, Chase SM, Batista AP, et al. Stabilization of a brain-computer interface via the alignment of low-dimensional spaces of neural activity. *Nat Biomed Eng*. 2020; 4(7):672–85. Epub 2020/04/22. <https://doi.org/10.1038/s41551-020-0542-9> PMID: 32313100.
25. Gao Y, Archer E, Paninski L, Cunningham J. Linear dynamical neural population models through nonlinear embeddings. *Advances in Neural Information Processing Systems*; Spain 2016. p. 163–71.

26. Wu A, Roy N, Keeley S, Pillow J. Gaussian process based nonlinear latent structure discovery in multivariate spike train data. *Advances in Neural Information Processing Systems*; California 2017. p. 3496–505. PMID: [31244512](#)
27. Batty E, Merel J, Brackbill N, Heitman A, Sher A, Litke A, et al. Multilayer recurrent network models of primate retinal ganglion cell responses. *ICLR*. 2017.
28. Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci*. 2008; 9(4):292–303. Epub 2008/03/06. <https://doi.org/10.1038/nrn2258> PMID: [18319728](#); PubMed Central PMCID: [PMC2631351](#).
29. Camastra F, Vinciarelli A. Estimating the intrinsic dimension of data with a fractal-based method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2002; 24(10):1404–7. <https://doi.org/10.1109/tpami.2002.1039212>
30. Camastra F. Data dimensionality estimation methods: a survey. *Pattern Recognition*. 2003; 36(12):2945–54. [https://doi.org/10.1016/s0031-3203\(03\)00176-6](https://doi.org/10.1016/s0031-3203(03)00176-6)
31. Horn JL. A rationale and test for the number of factors in factor analysis. *Psychometrika*. 1965; 30(2):179–85. <https://doi.org/10.1007/BF02289447> PMID: [14306381](#)
32. Andreas Buja NE. Remarks on Parallel Analysis. *Multivariate Behavioral Research*. 1992; 27:509–40. [https://doi.org/10.1207/s15327906mbr2704\\_2](https://doi.org/10.1207/s15327906mbr2704_2) PMID: [26811132](#)
33. Tenenbaum JB, Silva Vd, Langford JC. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*. 2000; 290(5500):2319–23. <https://doi.org/10.1126/science.290.5500.2319> PMID: [11125149](#)
34. Hinton GE, Salakhutdinov RR. Reducing the Dimensionality of Data with Neural Networks. *Science*. 2006; 313. <https://doi.org/10.1126/science.1127647> PMID: [16873662](#)
35. Grassberger P, Procaccia I. Measuring the strangeness of strange attractors. *Physica D: Nonlinear Phenomena*. 1983; 9(1):189–208. [https://doi.org/10.1016/0167-2789\(83\)90298-1](https://doi.org/10.1016/0167-2789(83)90298-1).
36. Kalantana Z, Einbeck J. On the computation of the correlation integral for fractal dimension estimation. *International Conference on Statistics in Science, Business and Engineering (ICSSBE)*. 2012.
37. Einbeck J, Kalantana Z. Intrinsic Dimensionality Estimation for High-dimensional Data Sets: New Approaches for the Computation of Correlation Dimension. *Journal of Emerging Technologies in Web Intelligence*. 2013; 5(2). <https://doi.org/10.4304/jetwi.5.2.91-97>
38. Levina E, Bickel PJ. Maximum Likelihood Estimation of Intrinsic Dimension. *Advances in Neural Information Processing Systems*. 2004.
39. Facco E, d'Errico M, Rodriguez A, Laio A. Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Scientific Reports*. 2017; 7(1). <https://doi.org/10.1038/s41598-017-11873-y> PMID: [28939866](#)
40. Zinovyev LAJBA. Estimating the effective dimension of large biological datasets using Fisher separability analysis. *International Joint Conference on Neural Networks (IJCNN)*. 2019.
41. Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE. Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nat Commun*. 2018; 9(1):4233. Epub 2018/10/14. <https://doi.org/10.1038/s41467-018-06560-z> PMID: [30315158](#); PubMed Central PMCID: [PMC6185944](#).
42. Nelder JA, Wedderburn RWM. Generalized Linear Models. *Journal of the Royal Statistical Society: Series A (General)*. 1972; 135(3):370–84. <https://doi.org/10.2307/2344614>.
43. Cunningham JP, Gilja V, Ryu SI, Shenoy KV. Methods for estimating neural firing rates, and their application to brain–machine interfaces. *Neural Networks*. 2009; 22(9):1235–46. <https://doi.org/10.1016/j.neunet.2009.02.004> PMID: [19349143](#)
44. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci*. 2014; 17(3):440–8. Epub 2014/02/04. <https://doi.org/10.1038/nn.3643> PMID: [24487233](#); PubMed Central PMCID: [PMC3955357](#).
45. Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, et al. Neural constraints on learning. *Nature*. 2014; 512(7515):423–6. <https://doi.org/10.1038/nature13665> PMID: [25164754](#)
46. Machens CK, Romo R, Brody CD. Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. *J Neurosci*. 2010; 30(1):350–60. Epub 2010/01/08. <https://doi.org/10.1523/JNEUROSCI.3276-09.2010> PMID: [20053916](#); PubMed Central PMCID: [PMC2947945](#).
47. Campadelli P, Casiraghi E, Ceruti C, Rozza A. Intrinsic Dimension Estimation: Relevant Techniques and a Benchmark Framework. *Mathematical Problems in Engineering*. 2015; 2015:759567. <https://doi.org/10.1155/2015/759567>

48. Kobayashi T, Misaki K, Nakagawa H, Madokoro S, Ota T, Ihara H, et al. Correlation dimension of the human sleep electroencephalogram. *Psychiatry Clin Neurosci*. 2000; 54(1):11–6. Epub 2004/11/24. <https://doi.org/10.1046/j.1440-1819.2000.00629.x> PMID: 15558873.
49. Boon MY, Henry BI, Suttle CM, Dain SJ. The correlation dimension: A useful objective measure of the transient visual evoked potential? *Journal of Vision*. 2008; 8(1):6–. <https://doi.org/10.1167/8.1.6> PMID: 18318609
50. Kégl B. Intrinsic Dimension Estimation Using Packing Numbers. In: Becker S, Thrun S, Obermayer K, editors. *Advances in Neural Information Processing Systems 15*: MIT Press; 2003. p. 697–704. [https://doi.org/10.1016/S0911-6044\(02\)00014-3](https://doi.org/10.1016/S0911-6044(02)00014-3) PMID: 21274410
51. Pinamonti G, Zhao J, Condon DE, Paul F, Noè F, Turner DH, et al. Predicting the Kinetics of RNA Oligonucleotides Using Markov State Models. *Journal of Chemical Theory and Computation*. 2017; 13(2):926–34. <https://doi.org/10.1021/acs.jctc.6b00982> PMID: 28001394
52. Gorban AN, Makarov VA, Tyukin IY. High-Dimensional Brain in a High-Dimensional World: Blessing of Dimensionality. *Entropy*. 2020; 22(1). <https://doi.org/10.3390/e22010082> PMID: 33285855
53. Fisher RA. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*. 1936; 7(2):179–88. <https://doi.org/10.1111/j.1469-1809.1936.tb02137.x>
54. Gorban AN, Golubkov A, Grechuk B, Mirkes EM, Tyukin IY. Correction of AI systems by linear discriminants: Probabilistic foundations. *Information Sciences*. 2018; 466:303–22. <https://doi.org/10.1016/j.ins.2018.07.040>
55. Le Morvan M, Zinovyev A, Vert J-P. NetNorM: Capturing cancer-relevant information in somatic exome mutation data with gene networks for cancer stratification and prognosis. *PLOS Computational Biology*. 2017; 13(6):e1005573. <https://doi.org/10.1371/journal.pcbi.1005573> PMID: 28650955
56. Xavier G, Antoine B, Yoshua B. Deep Sparse Rectifier Neural Networks. 2011/06/14: PMLR. p. 315–23.
57. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015; 521(7553):436–44. Epub 2015/05/29. <https://doi.org/10.1038/nature14539> PMID: 26017442.
58. Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH, et al. Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron*. 2018; 97(4):953–66.e8. <https://doi.org/10.1016/j.neuron.2018.01.004> PMID: 29398358
59. Perich MG, Gallego JA, Miller LE. A Neural Population Mechanism for Rapid Learning. *Neuron*. 2018; 100(4):964–76.e7. <https://doi.org/10.1016/j.neuron.2018.09.030> PMID: 30344047
60. Sussillo D, Stavisky SD, Kao JC, Ryu SI, Shenoy KV. Making brain-machine interfaces robust to future neural variability. *Nat Commun*. 2016; 7:13749. Epub 2016/12/14. <https://doi.org/10.1038/ncomms13749> PMID: 27958268; PubMed Central PMCID: PMC5159828.
61. Pandarinath C, O’Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*. 2018; 15(10):805–15. <https://doi.org/10.1038/s41592-018-0109-9> PMID: 30224673
62. Pandarinath C, Ames KC, Russo AA, Farshchian A, Miller LE, Dyer EL, et al. Latent Factors and Dynamics in Motor Cortex and Their Application to Brain–Machine Interfaces. *The Journal of Neuroscience*. 2018; 38(44):9390–401. <https://doi.org/10.1523/JNEUROSCI.1669-18.2018> PMID: 30381431
63. Farshchian A, Gallego JA, Miller LE, Solla SA, Cohen JP, Bengio Y. Adversarial Domain Adaptation For Stable Brain-Machine Interfaces. *ICLR*. 2019.
64. Smith LA. Intrinsic limits on dimension calculations. *Physics Letters A*. 1988; 133(6):283–8. [https://doi.org/10.1016/0375-9601\(88\)90445-8](https://doi.org/10.1016/0375-9601(88)90445-8).
65. Eckmann JP, Ruelle D. Fundamental limitations for estimating dimensions and Lyapunov exponents in dynamical systems. *Physica D: Nonlinear Phenomena*. 1992; 56(2):185–7. [https://doi.org/10.1016/0167-2789\(92\)90023-G](https://doi.org/10.1016/0167-2789(92)90023-G).
66. Stringer C, Pachitariu M, Steinmetz N, Carandini M, Harris KD. High-dimensional geometry of population responses in visual cortex. *Nature*. 2019; 571(7765):361–5. <https://doi.org/10.1038/s41586-019-1346-5> PMID: 31243367
67. Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*. 2019; 364(6437):eaav7893. <https://doi.org/10.1126/science.aav7893> PMID: 31000656
68. Evarts EV. Pyramidal tract activity associated with a conditioned hand movement in the monkey. *Journal of Neurophysiology*. 1966; 29(6):1011–27. <https://doi.org/10.1152/jn.1966.29.6.1011> PMID: 4961643
69. Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. *Journal of Neurophysiology*. 2009; 102(1):614–35. <https://doi.org/10.1152/jn.90941.2008> PMID: 19357332.

70. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, et al. Neural population dynamics during reaching. *Nature*. 2012; 487(7405):51–6. Epub 2012/06/23. <https://doi.org/10.1038/nature11129> PMID: 22722855; PubMed Central PMCID: PMC3393826.
71. Yan Y, Goodman JM, Moore DD, Solla SA, Bensmaia SJ. Unexpected complexity of everyday manual behaviors. *Nat Commun*. 2020; 11(1):3564. Epub 2020/07/18. <https://doi.org/10.1038/s41467-020-17404-0> PMID: 32678102; PubMed Central PMCID: PMC7367296.