Full title

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

Short title

Intrinsic dimensionality of neural signals

Authors

Ege Altan[1,2,*], Sara A. Solla[1,3], Lee E. Miller[1,2,4,5], Eric J. Perreault[2,4,5]

[1] Department of Physiology, Northwestern University, Chicago, IL, United States of America
[2] Department of Biomedical Engineering, Northwestern University, Evanston, IL, United States of America
[3] Department of Physics and Astronomy, Northwestern University, Evanston, IL, United States of America
[4] Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL, United States of America
[5] Shirley Ryan AbilityLab, Chicago, IL, United States of America
[*] Corresponding author. Email: egealtan@u.northwestern.edu

Abstract

It is generally accepted that the number of neurons in a given brain area far exceeds the information that area encodes. 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 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 from the low-dimensional manifold to 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 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 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 record from has increased exponentially for decades; today we can simultaneously record from 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 despite 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 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

Simulation of neural signals

We generated the synthetic data used to evaluate dimensionality estimation algorithms as follows. First, we created d signals by randomly sampling 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 subject was performing a center-out reaching task [41]. We verified that these randomly selected signals were uncorrelated. These signals provided a d-dimensional set used to construct synthetic high-dimensional data sets **(Fig 1)**. We allowed d to vary from 3 to 40. These signals were multiplied by a d x 96 mixing matrix W with entries that were randomly selected from a zero-mean Gaussian distribution with unit variance. This resulted in a 96-dimensional data set X. 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. A nonlinear embedding was implemented by processing each simulated channel X with an exponential activation function:

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

We chose this exponential activating function to control the degree of nonlinearity by varying the parameter α, 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 procedure generated datasets of known dimensionality, embedding type (linear/nonlinear), and signal-to-noise ratio.

Fig 1: **Generation of simulated datasets.** First, representative neural signals were obtained by randomly sampling the firing rates of primary motor cortical recordings. The number of sampled signals determined the intrinsic dimensionality of the dataset. Then, the dimensionality of the

sampled signals was increased through linear combinations by multiplying the signals with a weight matrix W. The entries of W were sampled from a zero-mean Gaussian distribution with unit variance. Then, the 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 only, the signals were then activated nonlinearly through the exponential function in Equation 1 (red box in diagram). In noisy simulations, zero-mean Gaussian noise with variance specified by the predetermined signal-to-noise ratio was added to the signals. This procedure yielded linear or nonlinear, noisy datasets with known signal-to-noise ratio.

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, 41-43]. 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, 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

148 literature. In this study, we used a threshold of 90%, which yielded accurate estimates of
149 dimensionality for the noise-free linear datasets.

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

153 $$PR = \frac{(\sum_{i=1}^{N} \lambda_i)^2}{\sum_{i=1}^{N} (\lambda_i)^2}$$ (Equation 2)

154 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
155 = 1. At the other extreme, if all eigenvalues are equal, the variance is spread evenly across all
156 the dimensions, and PR=N. The actual value of PR interpolates between these two extreme
157 conditions to estimate the intrinsic dimensionality, and thus the number of principal components
158 to be kept [5].

159 **Parallel Analysis (PA).** Much like the Participation Ratio, Parallel Analysis is a principled
160 approach to finding a variance threshold [31, 32]. Parallel Analysis generates null distributions
161 for the eigenvalues by repeatedly shuffling each dimension of the data separately. The shuffling
162 step ensures that the correlations remaining across the different dimensions of the data are due
163 to chance. The eigenvalues that exceed the 95th percentile of the null distribution are identified
164 as significant, and their number is the number of dimensions to be kept. Although this method
165 has not been directly applied to neural data, similar approaches based on finding null
166 distributions of eigenvalues have been used for neural dimensionality estimation [44].

167 **Nonlinear algorithms.** Nonlinear algorithms can in principle estimate the dimensionality of
168 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. Common local properties include distance and separability of each data point relative to its neighbors. Although nonlinear algorithms are not yet commonly used in neuroscience, they have been used to estimate dimensionality in several other fields that produce high-dimensional datasets [45].

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 [46, 47], 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 [48].

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 [49] 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* [50], 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 [51, 52]. 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 [53].

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 of an upper bound dimensionality D, for which we used 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.

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, under the assumption that most of the noise was relegated to the discarded, low-variance principal components.

Joint Autoencoder denoising. We also used a neural network for denoising **(Fig 2).** 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 half of an autoencoder to compressed subspaces Z_1 and Z_2 respectively, each of dimension $D < 48$. These compressed subspaces were used to obtain reconstructed versions of X_1 and X_2, denoted $\hat{X}_1$ and $\hat{X}_2$, using the expansive halves of the corresponding autoencoders. The cost function C for 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:

$$C = MSE\left(X_1, \hat{X}_1\right) + MSE\left(X_2, \hat{X}_2\right) + MSE\left(Z_1, Z_2\right) \quad \text{(Equation 3)}$$

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 a D-dimensional subspace (green). The reconstructed outputs of the networks were the denoised channels.

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 activation functions in all layers ensured that the

236 autoencoder network would both operate on and output non-negative signals while allowing for
237 nonlinear embeddings.

238 ## Ethics statement

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

244 ## Statistical analyses

245 We used Monte Carlo simulations to generate up to 10 replications of synthetic data sets, each
246 corresponding to microelectrode array recording data from an experimental session. We noted
247 the number of replications (n) in the figure captions where applicable. Our choice of the number
248 of replications is reasonable compared to the number of experimental sessions that we would
249 expect to see in experiments with monkeys [41, 54, 55]. The simulations differed by their
250 random number generator seed, which dictated the pseudorandom sampling procedures
251 required for generating the signals. There were three sampling steps in our simulations **(Fig 1)**.
252 First was the creation of the low-dimensional basis signals, which were sampled from an
253 empirical firing rate distribution. The second was the entries of the mixing matrix W, which were
254 sampled from a zero-mean Gaussian distribution with unit variance. The third was the additive
255 noise, sampled from a zero-mean Gaussian distribution with variance determined by the signal-
256 to-noise ratio. We used bootstrapping with 10,000 iterations to compute the statistic of interest
257 and computed its confidence interval using $\alpha = 0.05$. We used Bonferroni correction for multiple
258 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 3)**. Under noise-free conditions, the nonlinear algorithms were as accurate as the linear ones on linearly embedded datasets.

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 (bootstrapped confidence intervals do not overlap 6 at the significance level of α=0.05.

Next, we evaluated all algorithms on nonlinearly embedded noise-free datasets, also for $d = 6$. Nonlinearities were introduced as in Equation 1, using α = 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.

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 4)**. 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 4A)** and nonlinearly embedded **(Fig 4B)** datasets for $d > 6$. This underestimation increased with increasing d. For nonlinear datasets, the estimate saturated at $d = 13$, where underestimation began to get much

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

Fig 4. 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.

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 α in Equation 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 5**). Surprisingly, even Parallel Analysis was accurate at levels of nonlinearity around $\alpha \approx 8$, 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.

Fig 5. Effect of changing the degree of nonlinearity. Dimensionality of nonlinear datasets (n=10) with varying levels of nonlinearity, controlled by the α 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 difference between mean

values (bootstrapped confidence intervals do not overlap 0 at the significance level of α=0.05/3, Bonferroni corrected for multiple comparisons).

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 dimensionality of datasets with d = 6, by varying both the duration of the recordings and the number of recording channels.

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 6A**). In contrast, the accuracy of both Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors also depended on the duration of the data (**Fig 6B and C**). Around 30 seconds of data was sufficient for accurate estimates of intrinsic dimensionality using either of these two nonlinear methods.

Fig 6. 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 is logarithmically scaled between 5 seconds and 600 seconds. Correct dimensionality d = 6 is shown in gray. Warm colors indicate overestimation and cold colors indicate underestimation of dimensionality. D, E, and F) Same as A, B, and C, respectively, but for nonlinear datasets.

As expected for highly nonlinear datasets (α = 16, d = 6), Parallel Analysis was not accurate (**Fig 6D**) 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 6E and F**). Furthermore, while Levina-Bickel Maximum Likelihood estimation required around 30 seconds of data for accurate dimensionality estimates, Two Nearest Neighbors required more than one minute. These results would also depend on the actual dimensionality d of the tested signals; 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 7**). For any given noise level, estimation errors for the linear datasets **(Fig 7A)** were a bit smaller than those for the nonlinear datasets **(Fig 7B)**. Adding noise with a power of only 1% of that of the 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 7B)**. PA yielded consistent overestimation errors across all nonzero levels of noise for both linear and nonlinear data.

Fig 7. 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 difference between mean values (bootstrapped confidence intervals do not overlap 0 at the significance level of α=0.05/3, Bonferroni corrected for multiple comparisons).

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 8A)**,

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 8B)**. The Joint Autoencoder was significantly more effective for denoising the nonlinear datasets **(Fig 8C)**. 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 8D)**. The highest noise level we tested (20%; SNR = 7 dB) caused the dimensionality to be overestimated by about 100%**.**

Fig 8. 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 Variance accounted for between noise-free and denoised signals. Asterisks indicate significant difference between mean values (bootstrapped confidence intervals do not overlap 0 at the significance level of α=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 difference between mean values (bootstrapped confidence intervals do not overlap 0 at the significance level of α=0.05/3, Bonferroni corrected for multiple comparisons). C) Same as in A, but for nonlinear datasets. D) Same as in B, but for nonlinear datasets after JAE denoising.

Discussion

This study evaluated techniques for estimating the intrinsic dimensionality of high-dimensional neural recordings. We considered representative linear and nonlinear algorithms, 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. In the linear case, 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 same three methods 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, 41-43]. 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.

Despite the simplicity of linear algorithms, estimating 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 “decoding” behavioral parameters from M1 neural manifolds are superior to linear methods [56-59]. 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 methods 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, 43, 55]. 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 resulting variance accounted for (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 [60, 61]. Correlation Dimension, the method on which many nonlinear algorithms are based, requires that the number of data samples be on the order of $10^{d/2}$ [29]. The amount of data can be increased by either recording from more channels or for a longer duration. One study that investigated the

dimensionality of the primary visual cortex (V1) found the eigenvalue spectrum of the neural signals obtained from approximately thousand neurons to decay as a power law [62]. Their finding would not have been possible had they recorded from fewer neurons, which would have prevented them from observing the long tail of the eigenvalue distribution. One interpretation of this finding is that the linear dimensionality is arbitrarily large. However, an alternative interpretation is the existence of an extremely nonlinear manifold embedded within the neural space investigated in that study.

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 [57].

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 [63], 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. This procedure does not reflect experimental neural firing data, since the range of neural firing can differ significantly even across neurons of the same type. 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 activation function that we used mapped the [0,1] range onto itself.

Recommended analysis pipeline

Based on our results, we recommend the following approach for estimating the dimensionality of neural recordings. 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 nonlinearity of the manifold. The usefulness of the denoising step was quantified through the variance accounted for (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. If the reconstruction VAF results prefer the Joint Autoencoder, this denoising method yields better denoised signals. Once the signals are denoised, and the linearity of the

manifold is established, either a linear or nonlinear dimensionality estimation method should be used depending on the expected linearity of the manifold as determined by the comparative performance of the 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 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 may be particularly difficult or even impractical. Noise is a confounding factor and must be eliminated prior to dimensionality estimation. Most existing studies have estimated intrinsic dimensionality using linear methods, as they 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 work better 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 ~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.

Acknowledgments

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

References

1. Saxena S, Cunningham JP. Towards the neural population doctrine. Current Opinion in Neurobiology. 2019;55:103-11. doi: 10.1016/j.conb.2019.02.002.
2. Cunningham JP, Yu BM. Dimensionality reduction for large-scale neural recordings. Nature Neuroscience. 2014;17(11):1500-9. doi: 10.1038/nn.3776.
3. Elsayed GF, Cunningham JP. Structure in neural population recordings: an expected byproduct of simpler phenomena? Nature Neuroscience. 2017;20(9):1310-8. doi: 10.1038/nn.4617.
4. Gao P, Ganguli S. On Simplicity and Complexity in the Brave New World of Large-Scale Neuroscience. Current Opinion in Neurobiology. 2015.
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. doi: 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. doi: 10.1016/j.neuron.2019.05.003. PubMed PMID: 31171448; PubMed Central PMCID: PMCPMC7002296.
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. doi: 10.1016/j.neuron.2018.05.015.
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. doi: 10.1016/j.conb.2018.12.009.
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. doi: 10.1016/j.neuron.2017.05.025. PubMed PMID: 28595054; PubMed Central PMCID: PMCPMC6122849.
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. doi: 10.1038/nature12160.
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. doi: https://doi.org/10.1016/j.neuron.2005.09.032.
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. doi: 10.1038/nature12742. PubMed PMID: 24201281; PubMed Central PMCID: PMCPMC4121670.
13. Churchland MM, Abbott LF. Two layers of neural variability. Nature Neuroscience. 2012;15(11):1472-4. doi: 10.1038/nn.3247.
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. doi: 10.1038/nature10918.
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. doi: 10.1371/journal.pcbi.1007074. PubMed PMID: 31150376; PubMed Central PMCID: PMCPMC6586365.
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. doi: 10.1038/s41593-019-0555-4. PubMed PMID: 31907438; PubMed Central PMCID: PMCPMC7007364.
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

581 Comput Biol. 2016;12(12):e1005141. Epub 2016/12/08. doi: 10.1371/journal.pcbi.1005141.
582 PubMed PMID: 27926936; PubMed Central PMCID: PMCPMC5142778.
583 18. Mazzucato L, Fontanini A, La Camera G. Stimuli Reduce the Dimensionality of Cortical
584 Activity. Front Syst Neurosci. 2016;10:11-. doi: 10.3389/fnsys.2016.00011. PubMed PMID:
585 26924968.
586 19. Stevenson IH, Kording KP. How advances in neural recording affect data analysis.
587 Nature Neuroscience. 2011;14(2):139-42. doi: 10.1038/nn.2731.
588 20. Camastra F, Staiano A. Intrinsic dimension estimation: Advances and open problems.
589 Information Sciences. 2016;328:26-41. doi: 10.1016/j.ins.2015.08.029.
590 21. Lee JA, Verleysen M. Nonlinear Dimensionality Reduction. 1 ed. New York: Springer;
591 2007. XVII, 309 p.
592 22. Pang R, Lansdell BJ, Fairhall AL. Dimensionality reduction in neuroscience. Current
593 Biology. 2016;26(14):R656-R60. doi: 10.1016/j.cub.2016.05.029.
594 23. Pandarinath C, Nuyujukian P, Blabe CH, Sorice BL, Saab J, Willett FR, et al. High
595 performance communication by people with paralysis using an intracortical brain-computer
596 interface. Elife. 2017;6. Epub 2017/02/22. doi: 10.7554/eLife.18554. PubMed PMID: 28220753;
597 PubMed Central PMCID: PMCPMC5319839.
598 24. Degenhart AD, Bishop WE, Oby ER, Tyler-Kabara EC, Chase SM, Batista AP, et al.
599 Stabilization of a brain-computer interface via the alignment of low-dimensional spaces of neural
600 activity. Nat Biomed Eng. 2020;4(7):672-85. Epub 2020/04/22. doi: 10.1038/s41551-020-0542-
601 9. PubMed PMID: 32313100.
602 25. Gao Y, Archer E, Paninski L, Cunningham J. Linear dynamical neural population models
603 through nonlinear embeddings. Advances in Neural Information Processing Systems;
604 Spain2016. p. 163-71.
605 26. Wu A, Roy N, Keeley S, Pillow J. Gaussian process based nonlinear latent structure
606 discovery in multivariate spike train data. Advances in Neural Information Processing Systems;
607 California2017. p. 3496-505.
608 27. Batty E, Merel J, Brackbill N, Heitman A, Sher A, Litke A, et al. Multilayer recurrent
609 network models of primate retinal ganglion cell responses. ICLR. 2017.
610 28. Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. Nat Rev Neurosci.
611 2008;9(4):292-303. Epub 2008/03/06. doi: 10.1038/nrn2258. PubMed PMID: 18319728;
612 PubMed Central PMCID: PMCPMC2631351.
613 29. Camastra F, Vinciarelli A. Estimating the intrinsic dimension of data with a fractal-based
614 method. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2002;24(10):1404-7.
615 doi: 10.1109/tpami.2002.1039212.
616 30. Camastra F. Data dimensionality estimation methods: a survey. Pattern Recognition.
617 2003;36(12):2945-54. doi: 10.1016/s0031-3203(03)00176-6.
618 31. Horn JL. A rationale and test for the number of factors in factor analysis. Psychometrika.
619 1965;30(2):179-85. doi: 10.1007/BF02289447.
620 32. Andreas Buja NE. Remarks on Parallel Analysis. MULTIVARIATE BEHAVIORAL
621 RESEARCH. 1992;27:509-40.
622 33. Tenenbaum JB, Silva Vd, Langford JC. A Global Geometric Framework for Nonlinear
623 Dimensionality Reduction. Science. 2000;290(5500):2319-23. doi:
624 10.1126/science.290.5500.2319.
625 34. Hinton GE, Salakhutdinov RR. Reducing the Dimensionality of Data with Neural
626 Networks. Science. 2006;313.
627 35. Grassberger P, Procaccia I. Measuring the strangeness of strange attractors. Physica D:
628 Nonlinear Phenomena. 1983;9(1):189-208. doi: https://doi.org/10.1016/0167-2789(83)90298-1.
629 36. Kalantana Z, Einbeck J. On the computation of the correlation integral for fractal
630 dimension estimation. International Conference on Statistics in Science, Business and
631 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). doi: 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). doi: 10.1038/s41598-017-11873-y.
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. doi: 10.1038/s41467-018-06560-z. PubMed PMID: 30315158; PubMed Central PMCID: PMCPMC6185944.
42. 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. doi: 10.1038/nn.3643. PubMed PMID: 24487233; PubMed Central PMCID: PMCPMC3955357.
43. 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. doi: 10.1038/nature13665.
44. 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. doi: 10.1523/JNEUROSCI.3276-09.2010. PubMed PMID: 20053916; PubMed Central PMCID: PMCPMC2947945.
45. Campadelli P, Casiraghi E, Ceruti C, Rozza A. Intrinsic Dimension Estimation: Relevant Techniques and a Benchmark Framework. Mathematical Problems in Engineering. 2015;2015:759567. doi: 10.1155/2015/759567.
46. 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. doi: 10.1046/j.1440-1819.2000.00629.x. PubMed PMID: 15558873.
47. 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-. doi: 10.1167/8.1.6.
48. 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.
49. 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. doi: 10.1021/acs.jctc.6b00982.
50. Gorban AN, Makarov VA, Tyukin IY. High-Dimensional Brain in a High-Dimensional World: Blessing of Dimensionality. Entropy. 2020;22(1). doi: 10.3390/e22010082.
51. Fisher RA. The use of multiple measurements in taxonomic problems. Annals of Eugenics. 1936;7(2):179-88. doi: 10.1111/j.1469-1809.1936.tb02137.x.
52. 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. doi: 10.1016/j.ins.2018.07.040.
53. 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. doi: 10.1371/journal.pcbi.1005573.
54. 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. doi: 10.1016/j.neuron.2018.01.004.

683 55. Perich MG, Gallego JA, Miller LE. A Neural Population Mechanism for Rapid Learning.
684 Neuron. 2018;100(4):964-76.e7. doi: 10.1016/j.neuron.2018.09.030.
685 56. Sussillo D, Stavisky SD, Kao JC, Ryu SI, Shenoy KV. Making brain-machine interfaces
686 robust to future neural variability. Nat Commun. 2016;7:13749. Epub 2016/12/14. doi:
687 10.1038/ncomms13749. PubMed PMID: 27958268; PubMed Central PMCID:
688 PMCPMC5159828.
689 57. Pandarinath C, O'Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, et al. Inferring
690 single-trial neural population dynamics using sequential auto-encoders. Nature Methods.
691 2018;15(10):805-15. doi: 10.1038/s41592-018-0109-9.
692 58. Pandarinath C, Ames KC, Russo AA, Farshchian A, Miller LE, Dyer EL, et al. Latent
693 Factors and Dynamics in Motor Cortex and Their Application to Brain–Machine Interfaces. The
694 Journal of Neuroscience. 2018;38(44):9390-401. doi: 10.1523/jneurosci.1669-18.2018.
695 59. Farshchian A, Gallego JA, Miller LE, Solla SA, Cohen JP, Bengio Y. Adversarial Domain
696 Adaptation For Stable Brain-Machine Interfaces. ICLR. 2019.
697 60. Smith LA. Intrinsic limits on dimension calculations. Physics Letters A. 1988;133(6):283-
698 8. doi: https://doi.org/10.1016/0375-9601(88)90445-8.
699 61. Eckmann JP, Ruelle D. Fundamental limitations for estimating dimensions and
700 Lyapunov exponents in dynamical systems. Physica D: Nonlinear Phenomena. 1992;56(2):185-
701 7. doi: https://doi.org/10.1016/0167-2789(92)90023-G.
702 62. Stringer C, Pachitariu M, Steinmetz N, Carandini M, Harris KD. High-dimensional
703 geometry of population responses in visual cortex. Nature. 2019;571(7765):361-5. doi:
704 10.1038/s41586-019-1346-5.
705 63. Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-
706 Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity.
707 Journal of Neurophysiology. 2009;102(1):614-35. doi: 10.1152/jn.90941.2008. PubMed PMID:
708 19357332.
709


Sampled Signals
W
Scale to [0,1]
Nonlinear Activation
High-dimensional signals
Add noise
Simulated Dataset

Fig1

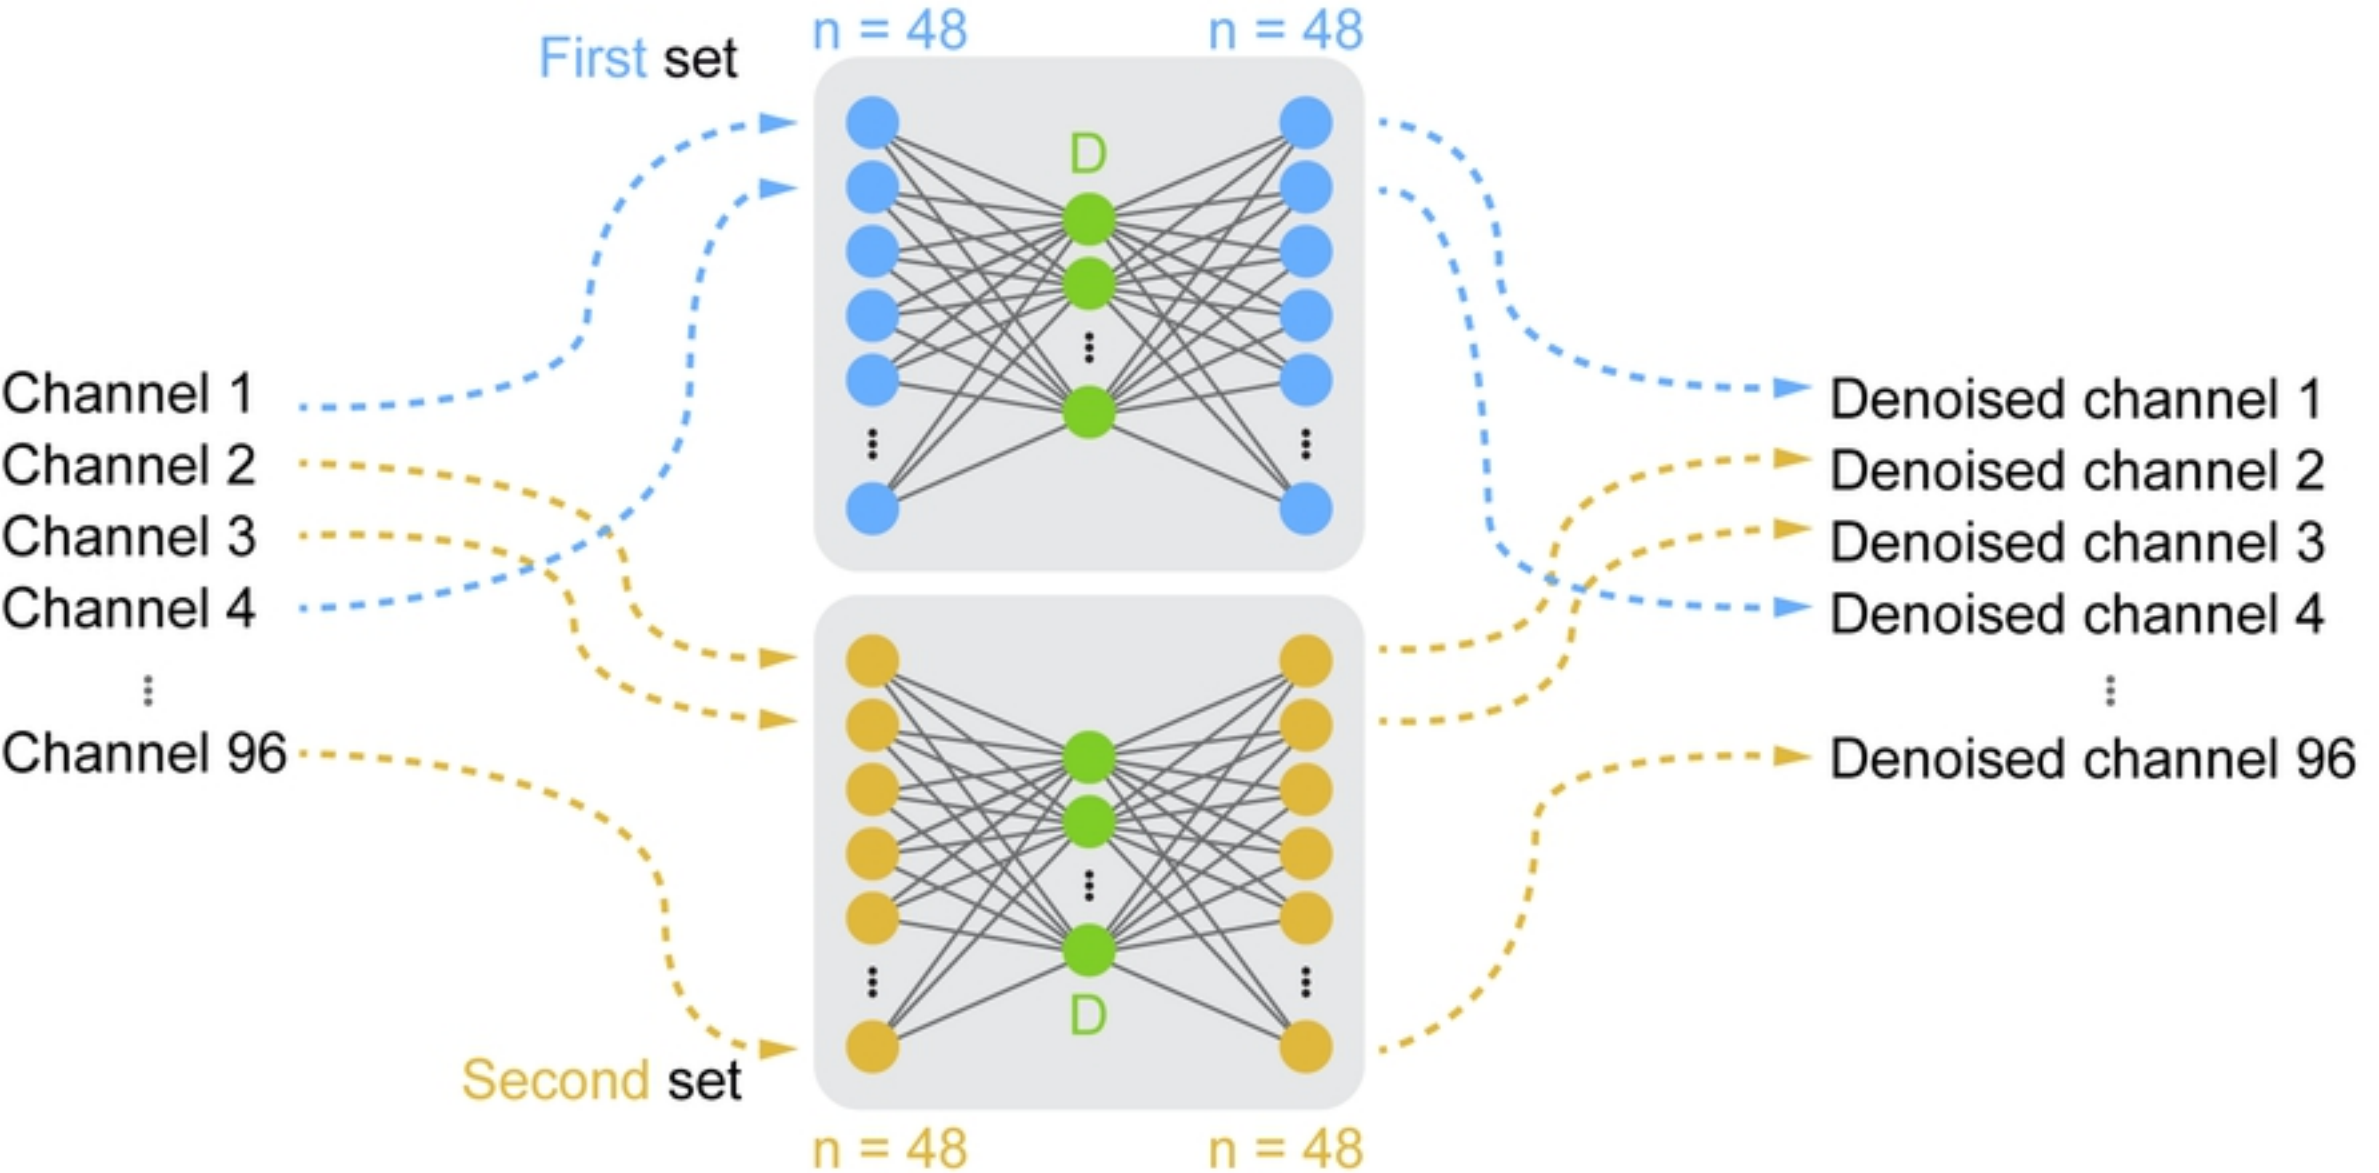
First set
n = 48
n = 48
D
Channel 1
Channel 2
Channel 3
Channel 4
Channel 96
Denoised channel 1
Denoised channel 2
Denoised channel 3
Denoised channel 4
Denoised channel 96
D
Second set
n = 48
n = 48

Fig2

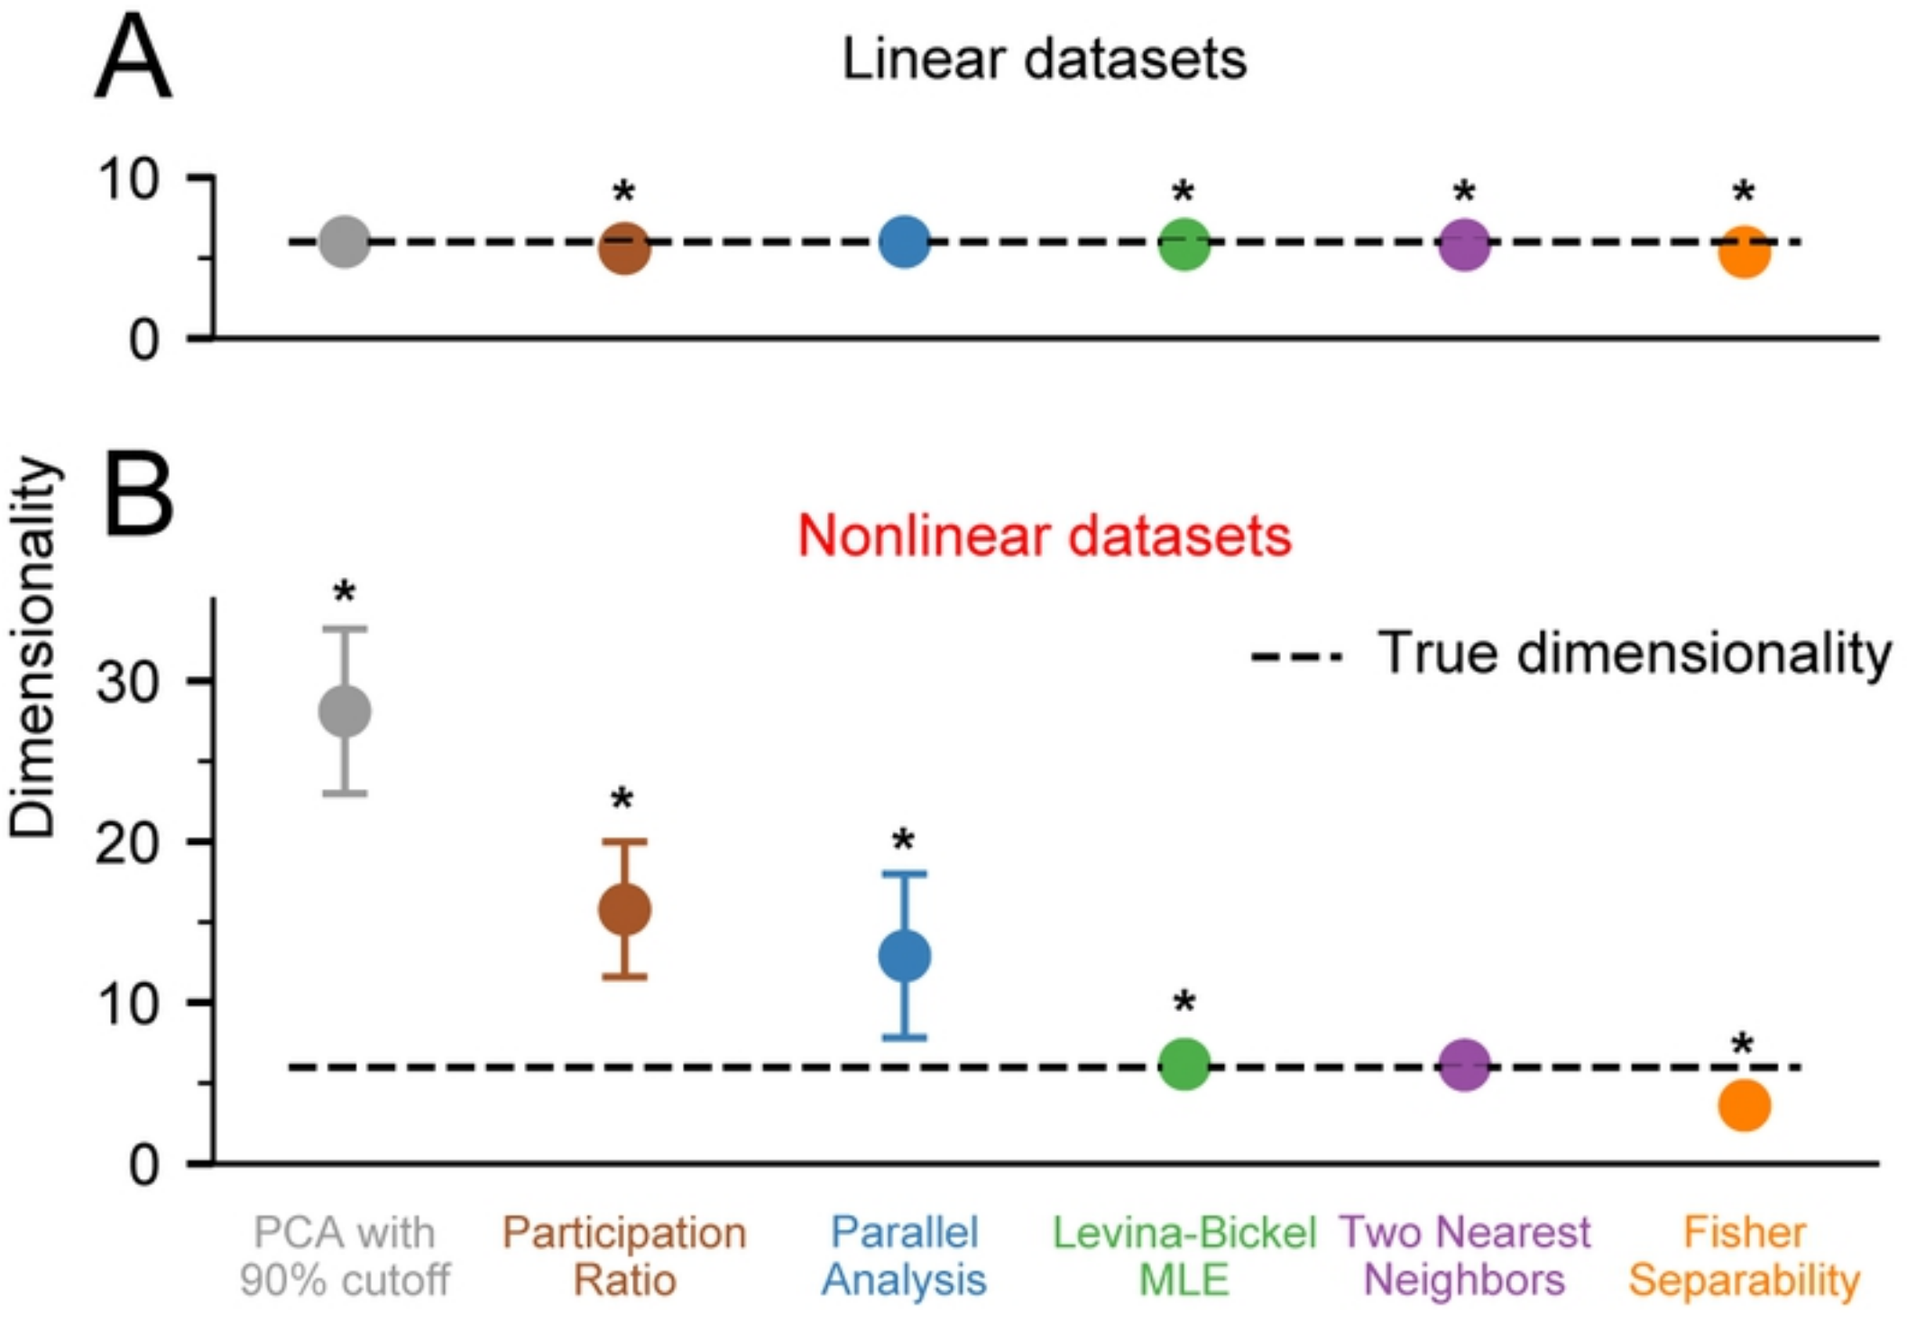
A
Linear datasets
10
0
*
*
*
*
B
Nonlinear datasets
Dimensionality
True dimensionality
30
20
10
0
*
*
*
*
*
PCA with 90% cutoff
Participation Ratio
Parallel Analysis
Levina-Bickel MLE
Two Nearest Neighbors
Fisher Separability

Fig3

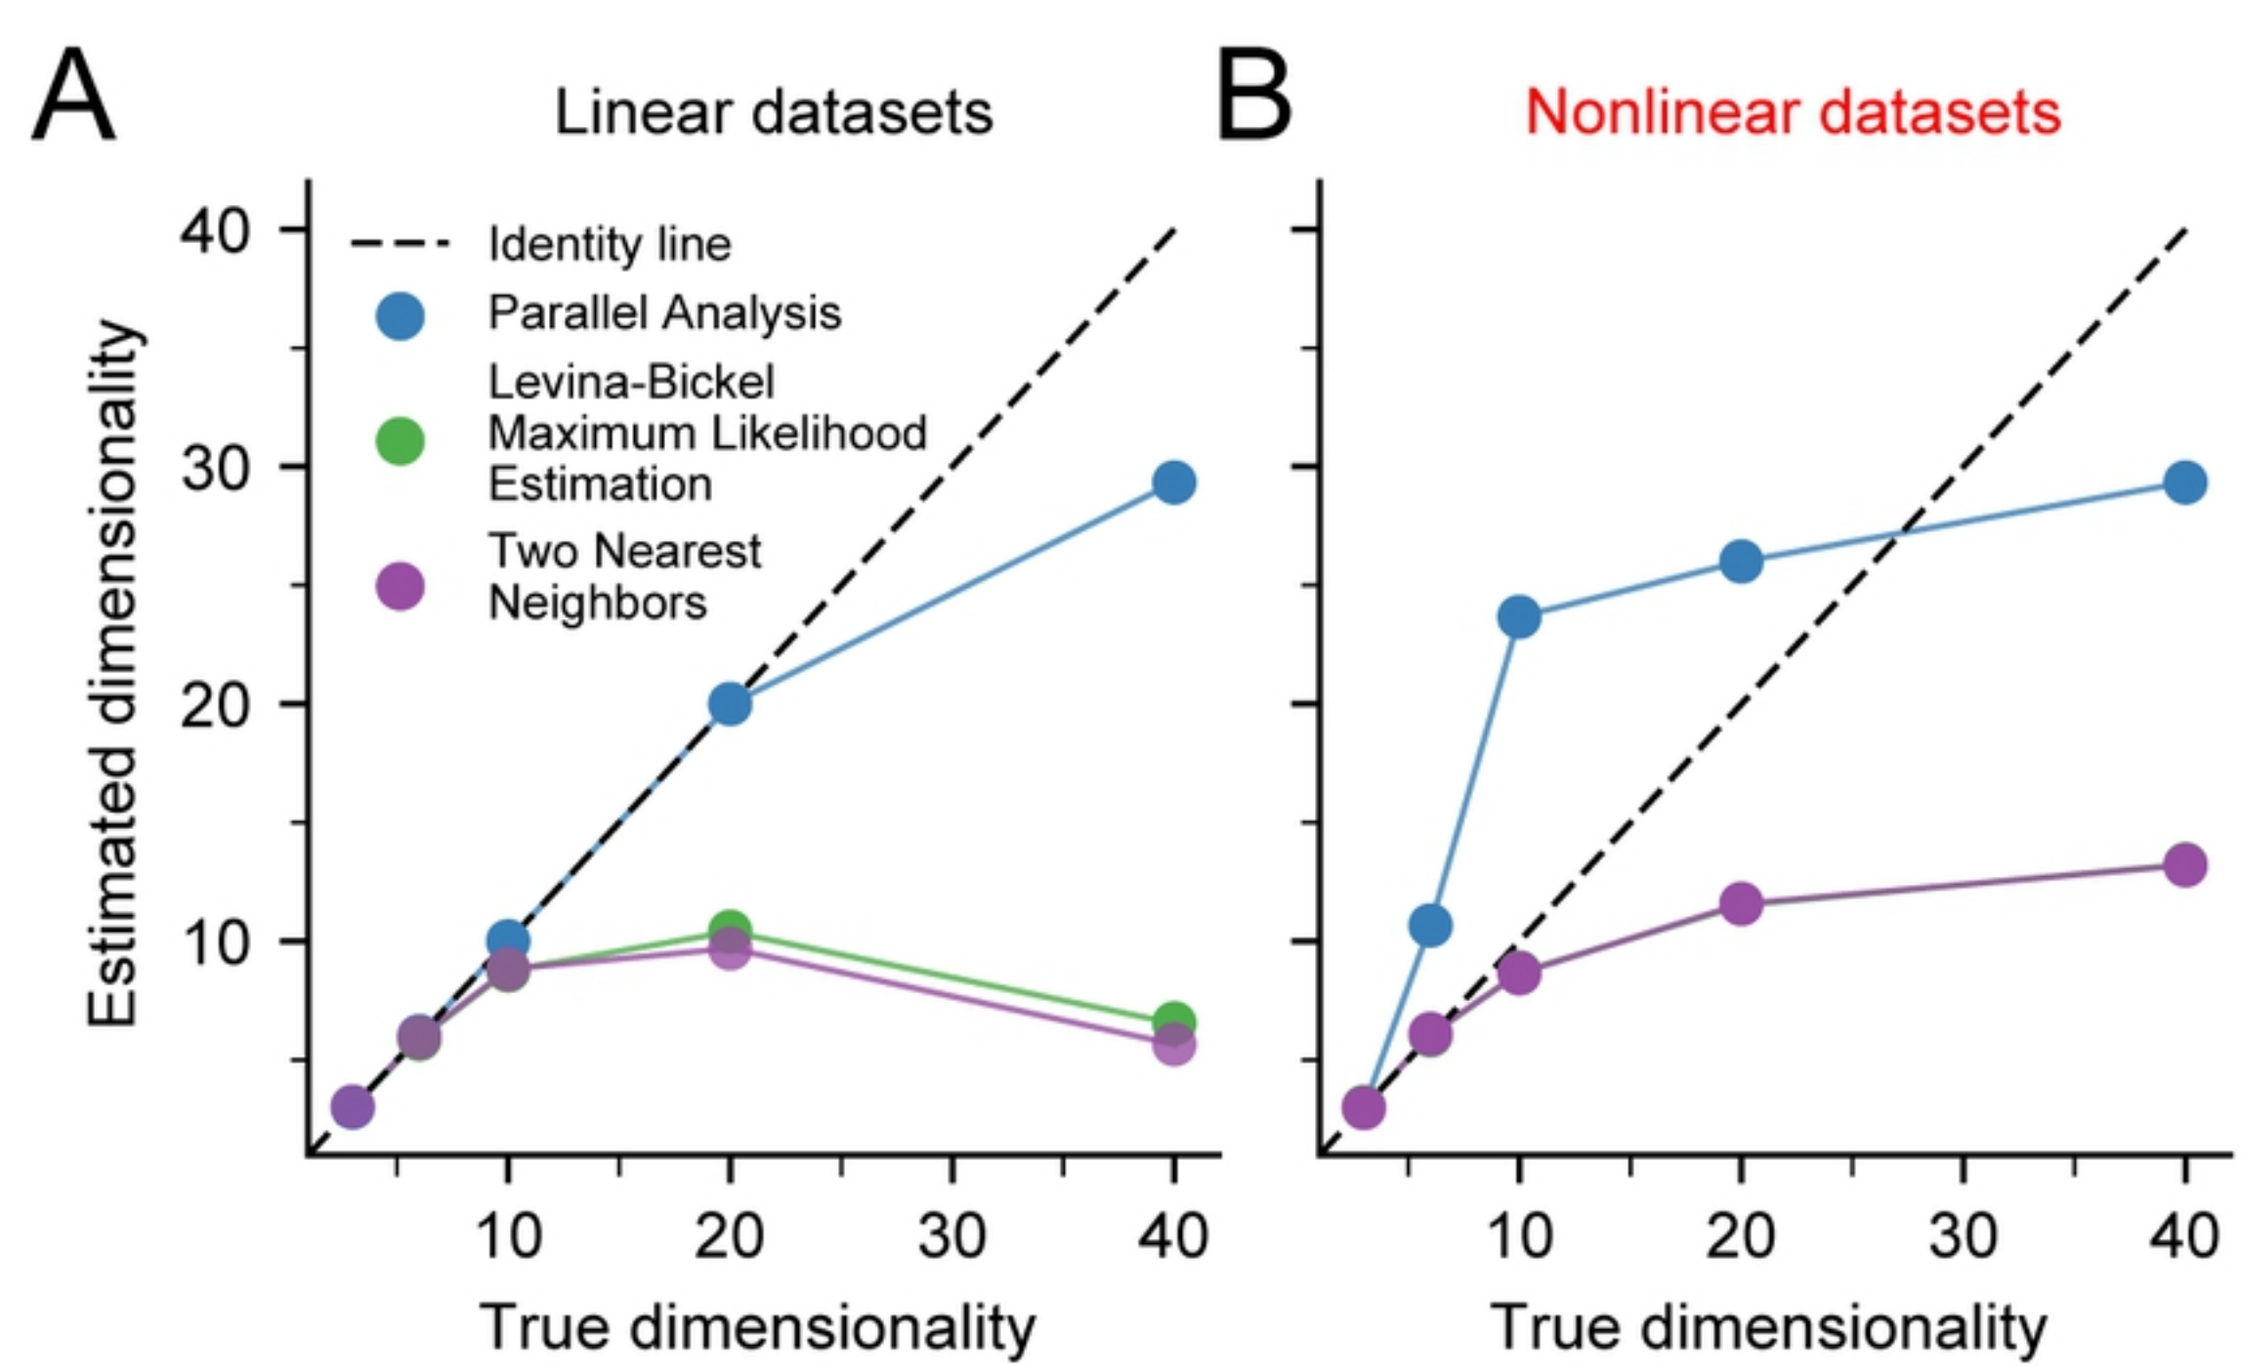
A
Linear datasets
B
Nonlinear datasets
Identity line
Parallel Analysis
Levina-Bickel
Maximum Likelihood
Estimation
Two Nearest
Neighbors
Estimated dimensionality
True dimensionality
True dimensionality
10
20
30
40

Fig4

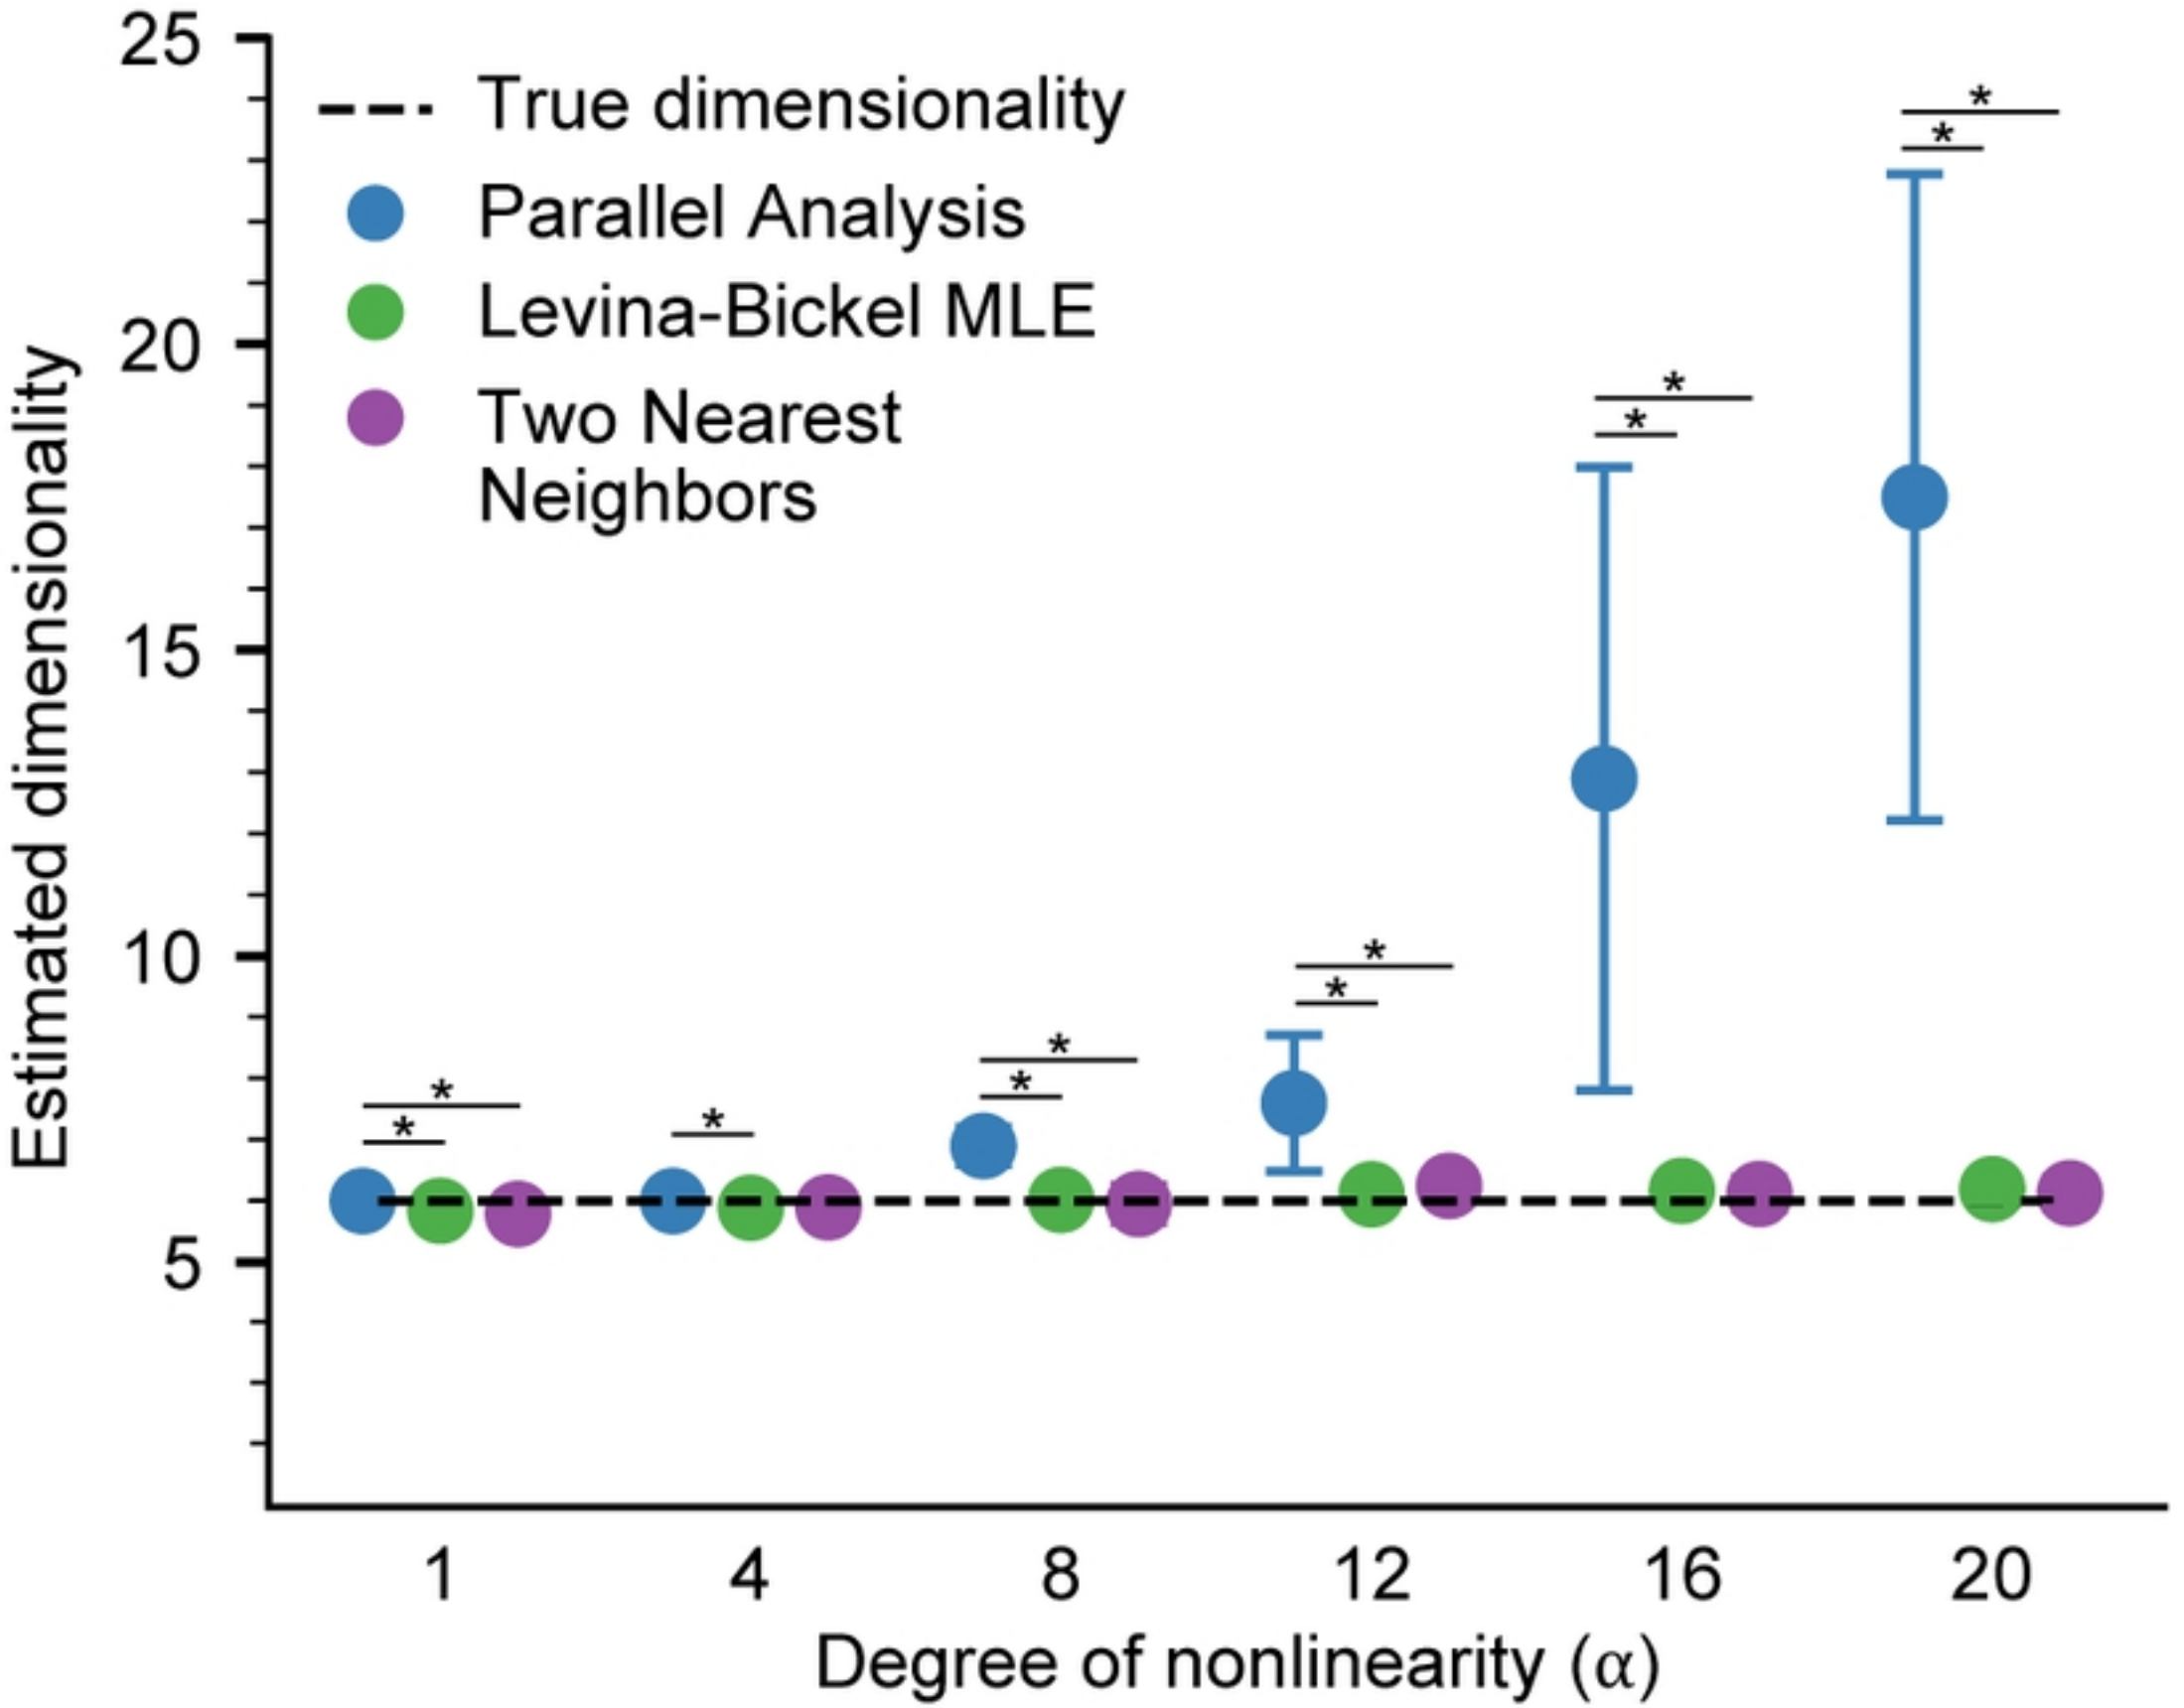

Fig5

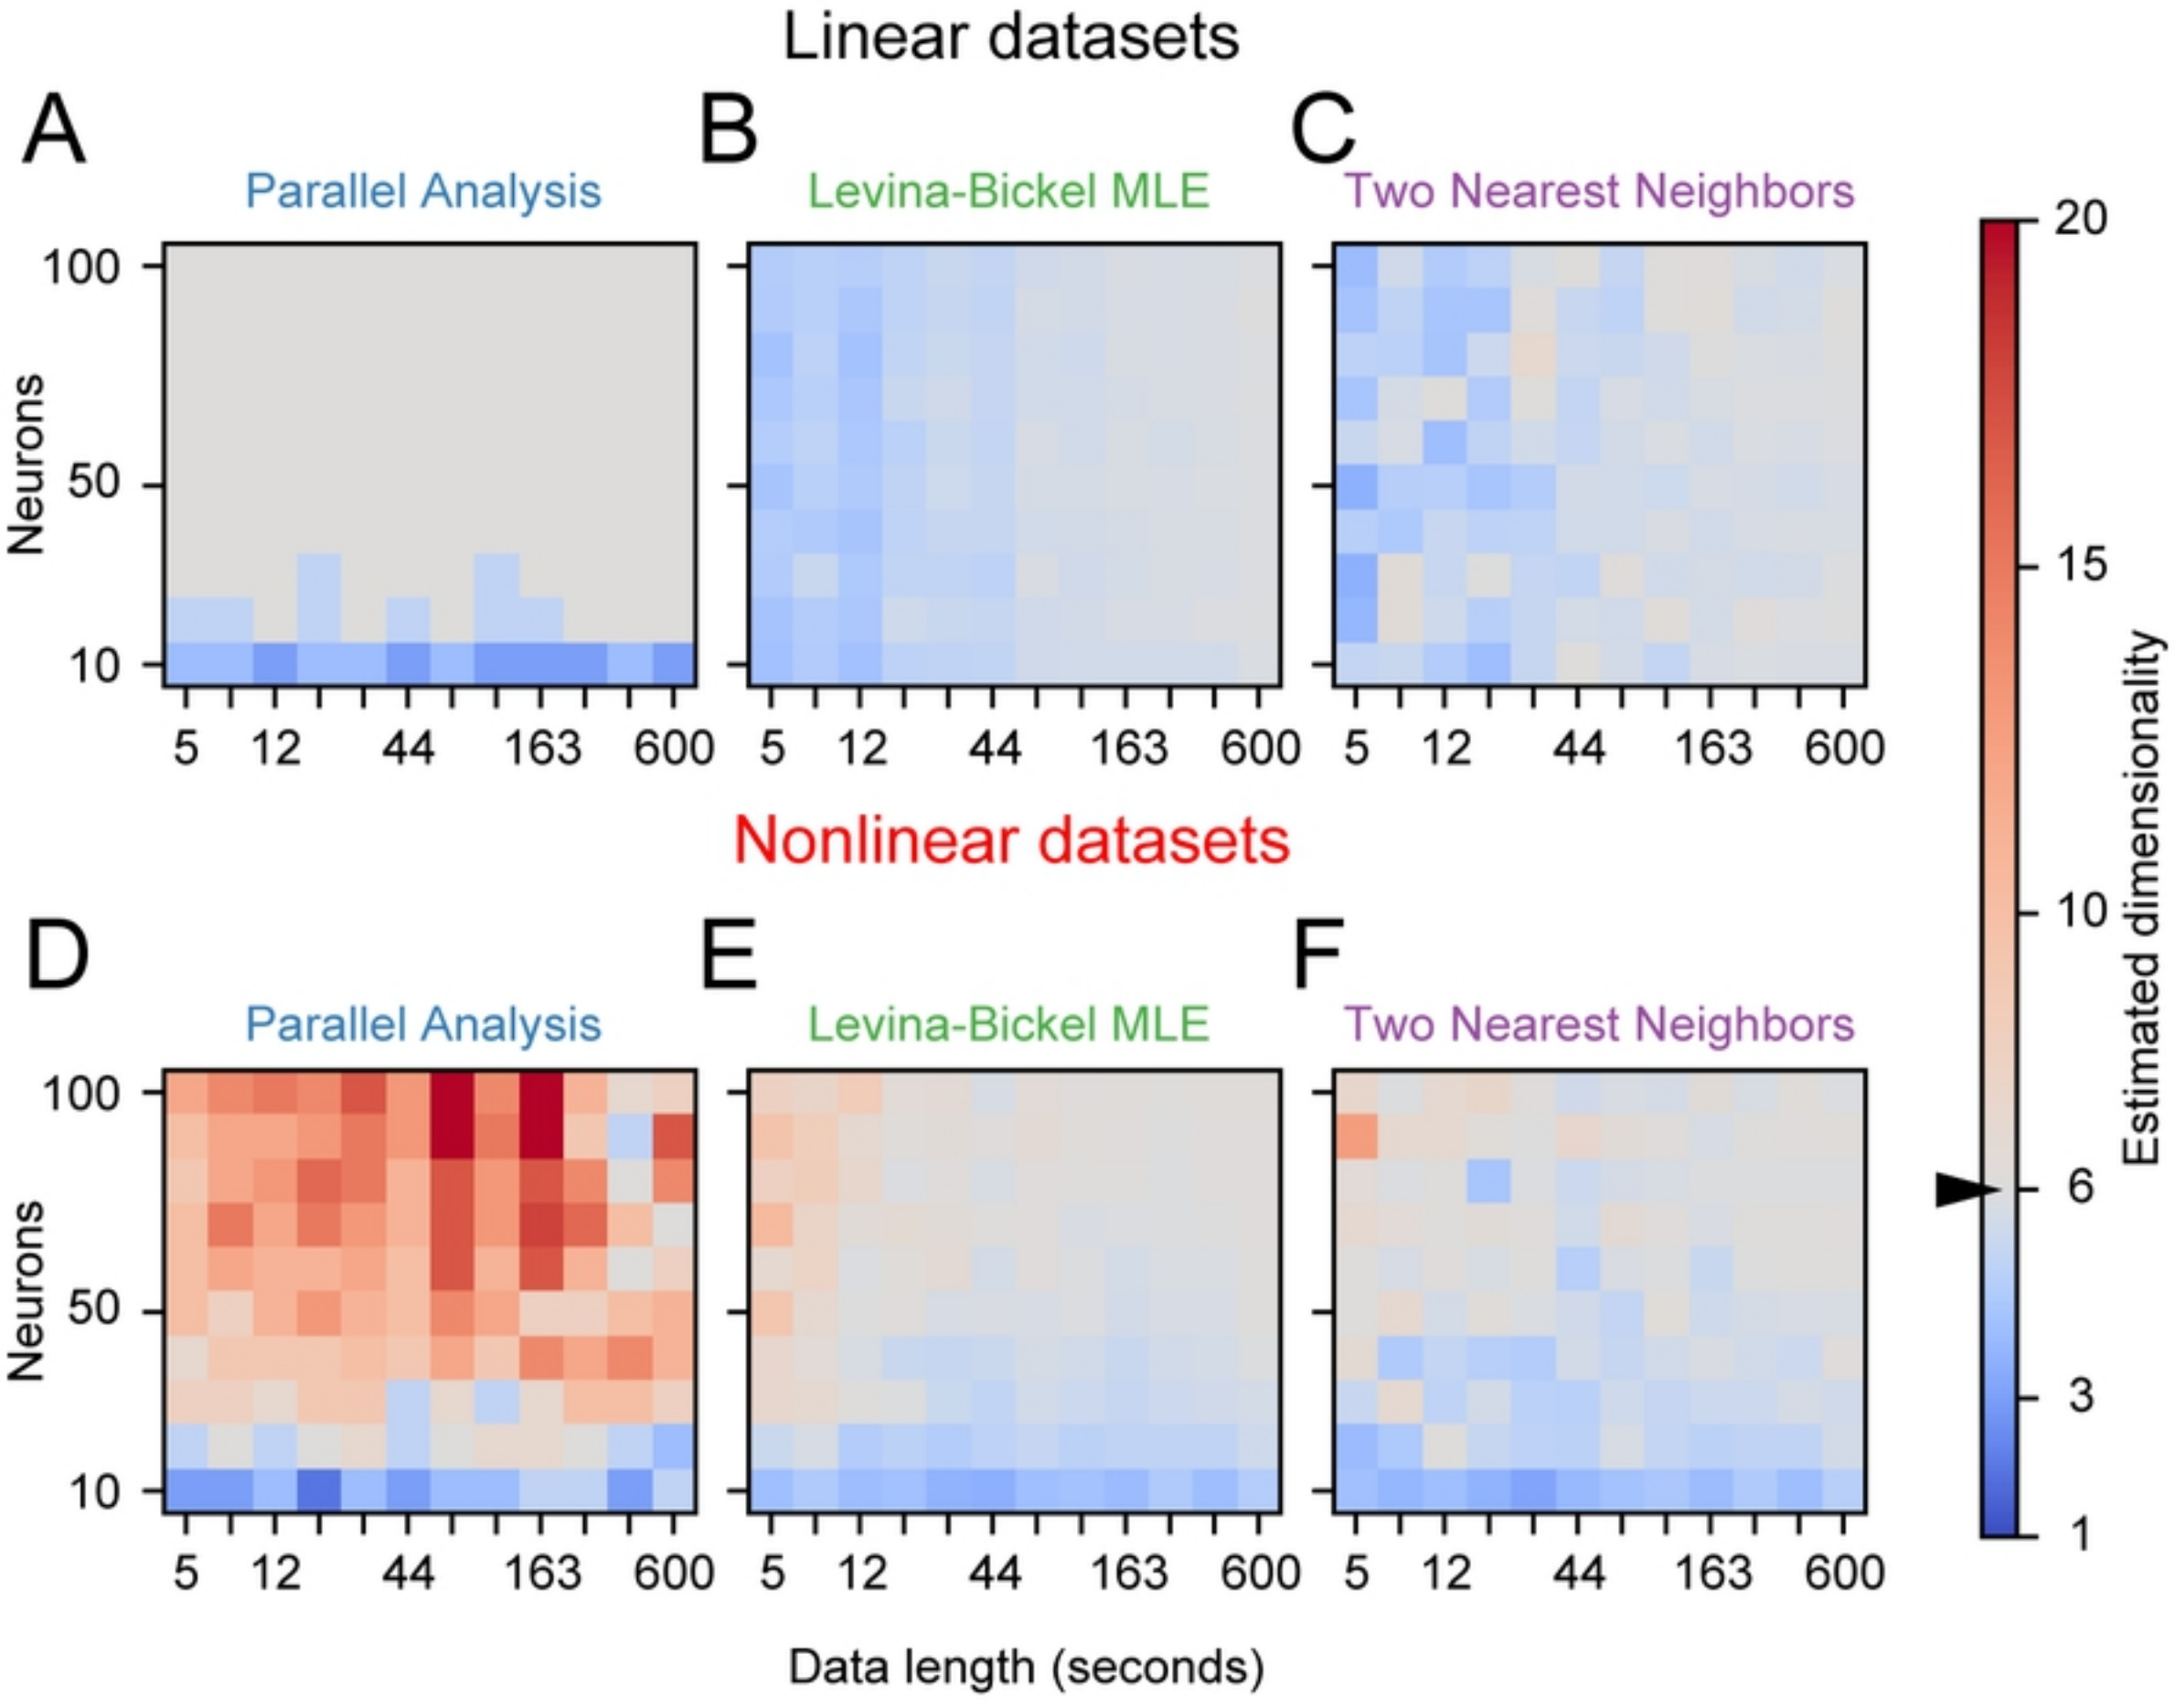
Linear datasets
A
Parallel Analysis
B
Levina-Bickel MLE
C
Two Nearest Neighbors
Neurons
100
50
10
5
12
44
163
600
Nonlinear datasets
D
Parallel Analysis
E
Levina-Bickel MLE
F
Two Nearest Neighbors
Data length (seconds)
Estimated dimensionality
20
15
10
6
3
1

Fig6

A

Linear datasets

True dimensionality

Parallel Analysis

Levina-Bickel MLE

Two Nearest Neighbors

Dimensionality

30

20

10

*

*

*

*

*

*

*

*

B

Nonlinear datasets

n.s.

Dimensionality

30

20

10

*

*

*

*

20

10

7

Signal-to-noise (dB)

Fig7

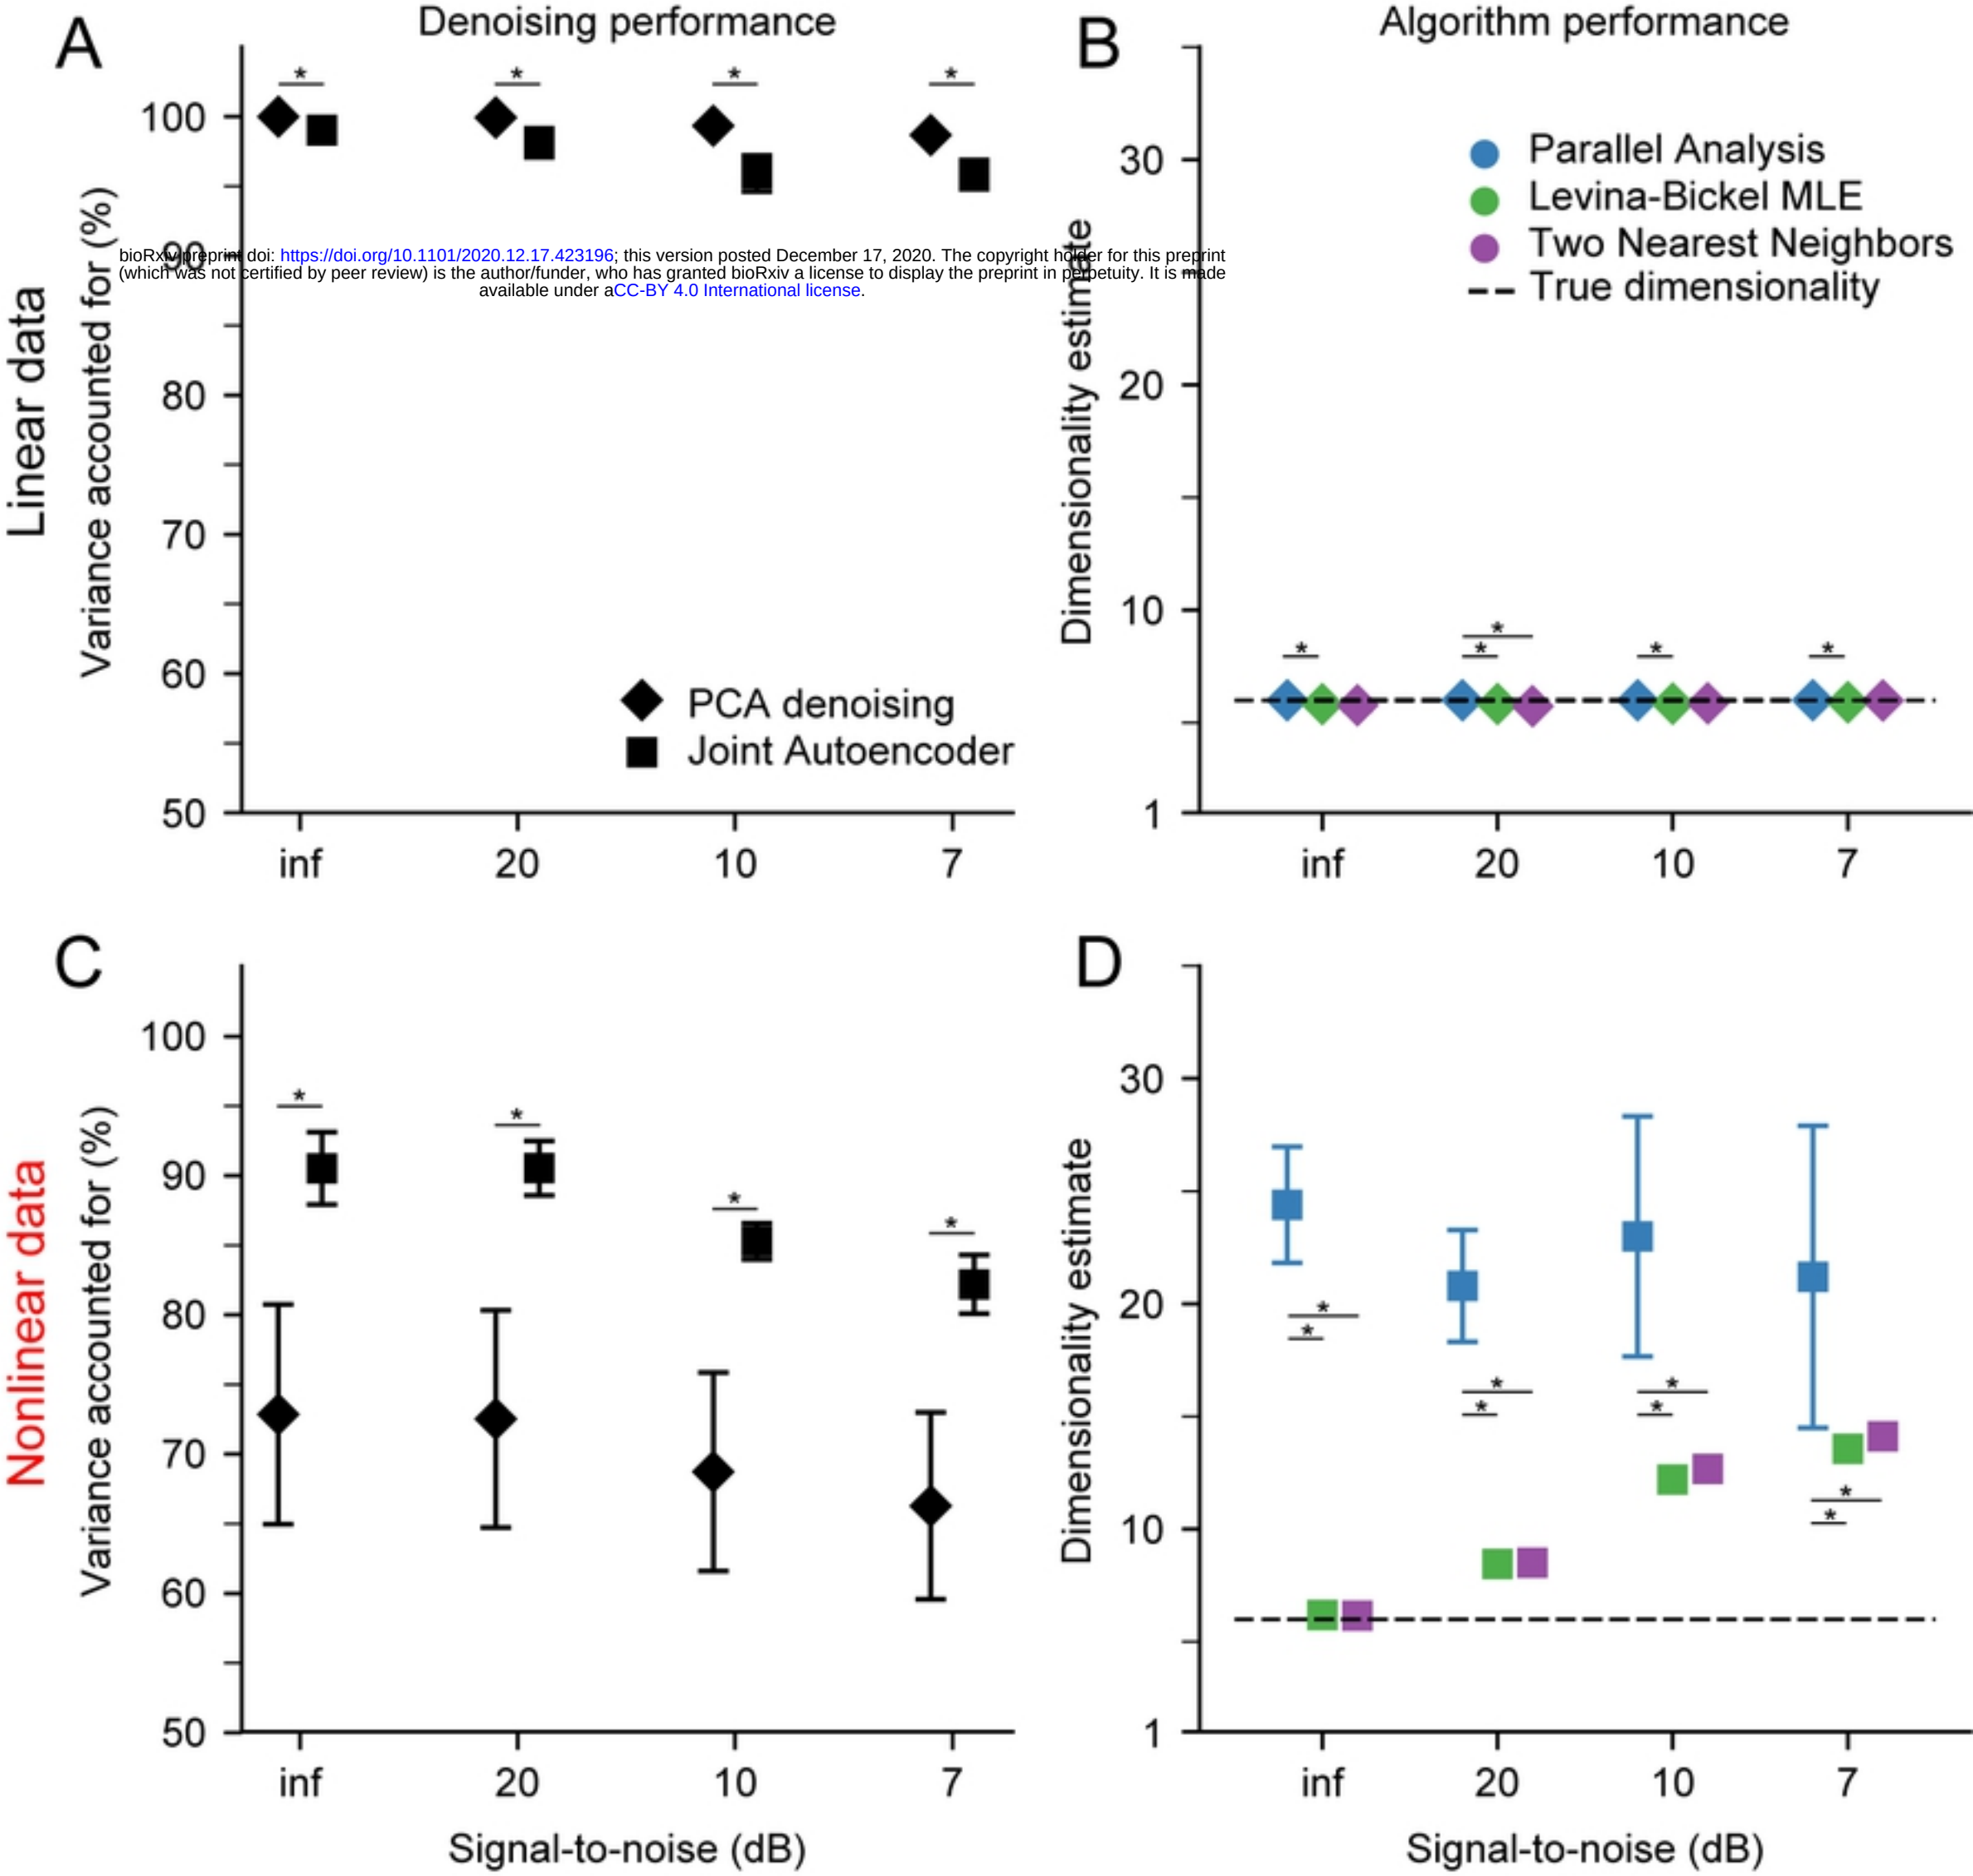

Fig8