Novelty Assessment Report

Paper: Estimating Dimensionality of Neural Representations from Finite Samples
PDF URL: https://openreview.net/pdf?id=iM4o9a83F7
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-27

Abstract

The global dimensionality of a neural representation manifold provides rich insight into the computational process underlying both artificial and biological neural networks. However, all existing measures of global dimensionality are sensitive to the number of samples, i.e., the number of rows and columns of the sample matrix. We show that, in particular, the participation ratio of eigenvalues, a popular measure of global dimensionality, is highly biased with small sample sizes, and propose a bias-corrected estimator that is more accurate with finite samples and with noise. On synthetic data examples, we demonstrate that our estimator can recover the true known dimensionality. We apply our estimator to neural brain recordings, including calcium imaging, electrophysiological recordings, and fMRI data, and to the neural activations in a large language model and show our estimator is invariant to the sample size. Finally, our estimators can additionally be used to measure the local dimensionalities of curved neural manifolds by weighting the finite samples appropriately.

Core Task Landscape

This paper addresses: **Estimating Global Dimensionality of Neural Representation Manifolds from Finite Samples**

A total of **16 papers** were analyzed and organized into a taxonomy with **13 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Intrinsic Dimensionality Estimation Methods**
- **Neural Representation Dimensionality in Biological Systems**
- **Neural Representation Dimensionality in Artificial Systems**
- **Theoretical Frameworks for Neural Dimensionality**
- **Manifold Learning and Reconstruction Methods**
- **Domain-Specific Manifold Dimensionality Applications**

Complete Taxonomy Tree

- Estimating Global Dimensionality of Neural Representation Manifolds from Finite Samples Survey Taxonomy
- Intrinsic Dimensionality Estimation Methods
  - Finite-Sample Bias Correction Techniques ★ (2 papers)
  - [0] Estimating Dimensionality of Neural Representations from Finite Samples (Anon et al., 2026) View paper
  - [4] Intrinsic dimension estimation for locally undersampled data (Erba, 2019) View paper
  - General Estimation Approaches (3 papers)
  - [8] Manifold-adaptive dimension estimation revisited (Benko Zsigmond, 2020) View paper
  - [14] Intrinsic Dimension Estimation. (Adam Block, 2021) View paper
  - [15] IAN: Iterated Adaptive Neighborhoods for Manifold Learning and Dimensionality Estimation. (Luciano Dyballa, 2023) View paper
  - Bayesian Nonparametric Dimensionality Estimation (1 papers)
  - [16] Bayesian Manifold Regression (Yun Yang, 2013) View paper
- Neural Representation Dimensionality in Biological Systems
  - Multi-Electrode and Population Recording Analysis (2 papers)
  - [2] Estimating the dimensionality of the manifold underlying multi-electrode neural recordings (Ege Altan, 2021) View paper
  - [9] Statistical Analysis of Large-Scale Neural Representations (Chun, 2025) View paper
  - Cortical Object Representation Dimensionality (1 papers)
  - [6] Dimensionality of object representations in monkey inferotemporal cortex (Sidney R. Lehky, 2014) View paper
  - Hippocampal Cognitive Map Variability (1 papers)
  - [10] Probing variability in a cognitive map using manifold inference from neural dynamics (Ryan Low, 2018) View paper
- Neural Representation Dimensionality in Artificial Systems
  - Deep Network Image Representation Dimensionality (1 papers)
  - [1] On the intrinsic dimensionality of image representations (Sixue Gong, 2019) View paper
  - Intrinsic Dimension-Guided Network Compression (1 papers)
  - [13] Deep Convolutional Neural Network Compression based on the Intrinsic Dimension of the Training Data (Abir Mohammad Hadi, 2024) View paper
- Theoretical Frameworks for Neural Dimensionality (1 papers)
  - [5] A theory of multineuronal dimensionality, dynamics and measurement (Peiran Gao, 2017) View paper
- Manifold Learning and Reconstruction Methods
  - Neural Network-Based Manifold Flattening (1 papers)

- [7] Representation learning via manifold flattening and reconstruction (Psenka, 2024) View paper
    - Geometric Clustering on Nonlinear Manifolds (1 papers)
    - [3] Geometric analysis of nonlinear manifold clustering (Tianjiao Ding, 2024) View paper
    - Manifold-Based Uncertainty Quantification (1 papers)
    - [12] Neural active manifolds: nonlinear dimensionality reduction for uncertainty quantification (Zanoni, 2024) View paper
- Domain-Specific Manifold Dimensionality Applications (1 papers)
    - [11] CH2O coordinate manifolds and their analysis for the paper "A global view of reactive coordinate manifolds from nonlinear dimensionality reduction" (Rappoport, 2023) View paper

Narrative

Core task: estimating global dimensionality of neural representation manifolds from finite samples. This field addresses a fundamental challenge in neuroscience and machine learning—determining the intrinsic dimensionality of high-dimensional neural activity or learned representations when only limited observations are available. The taxonomy reveals several complementary perspectives: one branch focuses on intrinsic dimensionality estimation methods themselves, developing algorithms that can handle finite-sample biases and adapt to local manifold structure (e.g., Intrinsic Dimension Undersampled Data[4], Manifold Adaptive Dimension Estimation[8]). Other branches examine neural representation dimensionality in biological systems (Dimensionality Manifold Neural Recordings[2], Dimensionality Inferotemporal Cortex[6]) and artificial systems (Intrinsic Dimensionality Image Representations[1], CNN Compression Intrinsic Dimension[13]), while theoretical frameworks (Multineuronal Dimensionality Theory[5], Statistical Neural Representations[9]) provide formal grounding. Additional branches cover manifold learning techniques for reconstruction (Manifold Flattening Reconstruction[7]) and domain-specific applications spanning cognitive maps to chemical reaction coordinates.

A central tension across these branches concerns how to reliably estimate dimensionality when sample sizes are modest relative to ambient dimensionality—a ubiquitous constraint in neural recordings and representation analysis. Many studies grapple with finite-sample biases that can systematically overestimate or underestimate true dimensionality, particularly when data lie on curved or variable-density manifolds. Estimating Dimensionality Finite Samples[0] sits squarely within the methodological branch addressing finite-sample bias correction techniques, closely aligned with work like Intrinsic Dimension Undersampled Data[4] that explicitly tackles undersampling challenges. Compared to broader manifold learning approaches (Geometric Nonlinear Manifold Clustering[3]) or domain-specific applications, this work emphasizes rigorous statistical correction to yield accurate global dimensionality estimates despite sampling limitations—a critical step for interpreting neural coding capacity and representational geometry across both biological and artificial systems.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Intrinsic dimension estimation for locally undersampled data

Authors: Erba, Vittorio, Gherardi Marco, Rotondo, Pietro | **Year/Venue**: 2019 • Scientific Reports | **URL**: View paper

Abstract

Identifying the minimal number of parameters needed to describe a dataset is a challenging problem known in the literature as intrinsic dimension estimation. All the existing intrinsic dimension estimators are not reliable whenever the dataset is locally undersampled, and this is at the core of the so called curse of dimensionality. Here we introduce a new intrinsic dimension estimator that leverages on simple properties of the tangent space of a manifold and extends the usual correlation integr...

Relationship Analysis

Both papers belong to the Finite-Sample Bias Correction Techniques category, addressing dimensionality estimation when sample sizes are limited. The candidate paper (Erba et al.) focuses on correcting bias in the correlation integral estimator for locally undersampled data by leveraging tangent space properties and introducing the Full Correlation Integral (FCI) method, whereas the original paper corrects finite-sample bias in the participation ratio estimator by deriving unbiased estimators through averaging over unequal indices. The key difference is that the original paper targets the participation ratio metric with explicit bias correction for both row and column sampling, while the candidate paper develops a geometric method based on correlation integrals with multiscale extensions for curved manifolds.

Contributions Analysis

Overall novelty summary. The paper proposes a bias-corrected estimator for the participation ratio of eigenvalues to measure global dimensionality of neural representation manifolds from finite samples. It resides in the 'Finite-Sample Bias Correction Techniques' leaf, which contains only two papers total (including this one). This places the work in a relatively sparse research direction within the broader taxonomy of 16 papers across 13 leaf nodes. The sibling paper in this leaf also addresses finite-sample correction, suggesting this specific methodological niche—correcting bias in dimensionality measures under limited sampling—is not yet crowded but represents a recognized gap in the field.

The taxonomy tree reveals that neighboring leaves focus on general intrinsic dimension estimation approaches (three papers using nearest-neighbor and correlation-based techniques) and Bayesian nonparametric methods (one paper). These adjacent directions do not explicitly emphasize finite-sample bias correction, instead offering broader algorithmic frameworks. The paper's position bridges methodological development (Intrinsic Dimensionality Estimation Methods branch) with applications to both biological neural recordings and artificial neural networks, connecting to separate branches that examine dimensionality in biological systems (three papers across cortical, hippocampal, and multi-electrode studies) and artificial systems (two papers on deep network representations). This cross-branch applicability distinguishes the work from purely algorithmic or purely empirical studies.

Among 23 candidates examined across three contributions, none were found to clearly refute any contribution. The bias-corrected participation ratio estimator examined 3 candidates with 0 refutable; the noise correction method examined 10 candidates with 0 refutable; and the weighted framework for local dimensionality examined 10 candidates with 0 refutable. This suggests that within the limited search scope—top-K semantic matches plus citation expansion—no prior work was identified that directly anticipates the specific combination of bias correction, noise handling, and weighted local dimensionality estimation proposed here. The noise correction and weighted framework contributions, each examined against 10 candidates, appear particularly distinct from existing approaches in the sampled literature.

Based on the limited search of 23 candidates, the work appears to occupy a methodologically focused niche with modest prior coverage. The taxonomy structure confirms that finite-sample bias correction is an emerging rather than saturated direction, and the contribution-level statistics indicate no substantial overlap with examined prior work. However, this assessment reflects the scope of semantic search and citation expansion, not an exhaustive survey of all dimensionality estimation literature. The cross-applicability to both biological and artificial neural systems, demonstrated empirically, may represent a practical contribution beyond the core methodological novelty.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Bias-corrected estimator for participation ratio of eigenvalues

Description: The authors derive an unbiased estimator of the participation ratio (PR) by correcting finite-sample bias in both the numerator and denominator. This estimator addresses the systematic bias that arises when computing global dimensionality from finite data matrices by averaging only over unequal indices, making it resistant to sample size variations.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Sample size determination for GEE analyses of stepped wedge cluster randomized trials

URL: View paper

Brief Assessment

Sample Size GEE[27] focuses on sample size determination for stepped wedge cluster randomized trials using generalized estimating equations with bias-corrected correlation parameters, not on participation ratio of eigenvalues or dimensionality estimation in neural representations.

2. Critical generalized inverse participation ratio distributions

URL: View paper

Brief Assessment

Inverse Participation Ratio Distributions[28] focuses on statistical properties and finite-size scaling of IPR distributions at critical points in quantum systems, not on deriving unbiased estimators for finite-sample bias correction in neural representation dimensionality.

3. A scale-dependent measure of system dimensionality

URL: View paper

Brief Assessment

Scale Dependent System Dimensionality[26] focuses on scale-dependent dimensionality analysis across different observation scales, not on correcting finite-sample bias in participation ratio estimation. The candidate does not address bias correction for finite samples.

Contribution 2: Noise correction method for dimensionality estimation

Description: The authors present a method to correct bias from additive or multiplicative noise in dimensionality estimation by using two independent trials of the same stimuli and neurons. This approach requires only two trials and achieves bias reduction of $O(1/P + 1/Q)$, more efficient than naive averaging methods.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Accuracy maximization analysis for sensory-perceptual tasks: Computational improvements, filter robustness, and coding advantages for scaled additive noise

URL: View paper

Brief Assessment

Accuracy Maximization Sensory Tasks[38] addresses a different problem domain (sensory-perceptual task optimization with filter learning) rather than dimensionality estimation from neural data. The noise correction in [38] is for filter response models in perceptual tasks, not for correcting bias in dimensionality metrics from neural recordings.

2. Estimating the functional dimensionality of neural representations

URL: View paper

Brief Assessment

Functional Dimensionality Neural Representations[36] mentions noise inflation in fMRI dimensionality estimates but does not present a specific two-trial correction method achieving $O(1/P + 1/Q)$ bias reduction as described in the original paper.

3. High-dimensional geometry of population responses in visual cortex

URL: View paper

Brief Assessment

Population Responses Visual Cortex[29] focuses on visual cortex neural recordings and mentions noise considerations, but does not present a systematic method for correcting bias from additive or multiplicative noise in dimensionality estimation using two independent trials as described in the original paper.

4. Python for information theoretic analysis of neural data

URL: View paper

Brief Assessment

Python Information Theoretic Analysis[35] focuses on information theoretic analysis of neural data using bias correction methods for entropy and mutual information estimation, not on dimensionality estimation or noise correction for participation ratio calculations.

5. Manifold Reconstruction of Differences: A Model-Based Iterative Statistical Estimation Algorithm With a Data-Driven Prior.

URL: View paper

Brief Assessment

Manifold Reconstruction Differences[37] focuses on CT image reconstruction using manifold learning for noise reduction in medical imaging, not on correcting bias in dimensionality estimation from neural data with additive/multiplicative noise.

6. Assessing Neural Network Representations During Training Using Noise-Resilient Diffusion Spectral Entropy

URL: View paper

Brief Assessment

Diffusion Spectral Entropy[32] focuses on entropy-based measures using diffusion geometry for neural network representations, not on bias correction for participation ratio dimensionality estimation using multiple trials of neural data.

7. Disentangling Identifiable Features from Noisy Data with Structured Nonlinear ICA

URL: View paper

Brief Assessment

Structured Nonlinear ICA[33] focuses on identifiable feature disentanglement in generative models with structured dependencies (temporal/spatial), not on correcting bias in dimensionality estimation from neural data using multiple trials.

8. Variable noise and dimensionality reduction for sparse Gaussian processes

URL: View paper

Brief Assessment

Variable Noise Sparse GPs[34] focuses on sparse Gaussian process approximations with variable noise modeling and dimensionality reduction for computational efficiency in GP regression, not on correcting bias from additive/multiplicative noise in dimensionality estimation of neural data using independent trials.

9. Gaussian partial information decomposition: Bias correction and application to high-dimensional data

URL: View paper

Brief Assessment

Gaussian Partial Information Decomposition[30] addresses bias correction in partial information decomposition for Gaussian distributions, not dimensionality estimation from neural representations. The candidate focuses on decomposing information into unique, redundant, and synergistic components between brain regions, while the original paper corrects finite-sample bias in participation ratio estimates for measuring neural manifold dimensionality.

10. A security model for smart grid SCADA systems using stochastic neural network

URL: View paper

Brief Assessment

Smart Grid Security Model[31] focuses on cybersecurity threat detection in SCADA systems using optimization and neural networks for attack classification, not on noise correction methods for dimensionality estimation in neural representations.

Contribution 3: Weighted dimensionality framework for local dimensionality estimation

Description: The authors extend their framework to measure local (intrinsic) dimensionality by introducing sample weighting schemes. This weighted approach enables estimation of dimensionality in local neighborhoods of a manifold and is resistant to noise, unlike existing popular local dimensionality estimators such as TwoNN.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Geometric analysis of nonlinear manifold clustering

URL: View paper

Brief Assessment

Geometric Nonlinear Manifold Clustering[3] focuses on manifold clustering via weighted self-expressive representations for clustering tasks, not on local intrinsic dimensionality estimation of neural representations. The weighted framework in [3] is designed for manifold-preserving clustering, not for measuring local dimensionality in curved manifolds as the original paper proposes.

2. Intrinsic Dimension Estimating Autoencoder (IDEA) Using CancelOut Layer and a Projected Loss

URL: View paper

Brief Assessment

IDEA CancelOut Layer[20] focuses on autoencoder-based intrinsic dimension estimation for manifold reconstruction, not on weighted sample methods for local dimensionality estimation on curved manifolds.

3. A continuous Formulation of intrinsic Dimension.

URL: View paper

Brief Assessment

Continuous Intrinsic Dimension[23] focuses on a continuous formulation of intrinsic dimension with spectral energy variance weighting, not the sample weighting schemes for local neighborhoods described in the original paper's framework.

4. Local smoothing for manifold learning

URL: View paper

Brief Assessment

Local Smoothing Manifold Learning[22] focuses on noise reduction and outlier handling for manifold learning using weighted local linear smoothing, not on local intrinsic dimensionality estimation with sample weighting schemes as described in the original paper.

5. A global view of reactive coordinate manifolds from nonlinear dimensionality reduction

URL: View paper

Brief Assessment

Reactive Coordinate Manifolds[21] focuses on weighted sampling of molecular coordinate manifolds for chemical reaction visualization, not neural representation dimensionality estimation. The weighting schemes serve different purposes in distinct domains.

6. Estimating low dimensional dynamical models for molecules

URL: View paper

Brief Assessment

Low Dimensional Molecular Models[25] focuses on molecular dynamics dimensionality reduction using weighted samples for local neighborhoods in molecular manifolds, not neural representation manifolds or finite-sample bias correction in participation ratio estimation.

7. Fermat Distances: Metric Approximation, Spectral Convergence, and Clustering Algorithms

URL: View paper

Brief Assessment

Fermat Distances[17] focuses on density-weighted path metrics and spectral convergence on manifolds, not on local intrinsic dimensionality estimation with sample weighting schemes for neural representations.

8. Mitigating Class Imbalance in Long-Tailed Visual Recognition Through the Use of Intrinsic Dimensionality

URL: View paper

Brief Assessment

Intrinsic Dimensionality Class Imbalance[24] focuses on using intrinsic dimensionality for class imbalance mitigation in long-tailed visual recognition, not on developing weighted sample methods for local dimensionality estimation on curved manifolds.

9. UCET8: Universal Curvature Equivalence on Noncompact, Bounded, and Lorentzian Geometries

URL: View paper

Brief Assessment

UCET8[19] focuses on curvature equivalence in geometric spaces (noncompact, bounded, Lorentzian geometries), not on dimensionality estimation methods for neural representations or manifolds with sample weighting schemes.

10. Weighted Linear Local Tangent Space Alignment via Geometrically Inspired Weighted PCA for Fault Detection

URL: View paper

Brief Assessment

Weighted Tangent Space Alignment[18] focuses on weighted PCA for local tangent space estimation in fault detection contexts, not on local intrinsic dimensionality estimation with sample weighting schemes resistant to noise as in the original paper's framework.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Estimating Dimensionality of Neural Representations from Finite Samples View paper
- [1] On the intrinsic dimensionality of image representations View paper
- [2] Estimating the dimensionality of the manifold underlying multi-electrode neural recordings View paper
- [3] Geometric analysis of nonlinear manifold clustering View paper
- [4] Intrinsic dimension estimation for locally undersampled data View paper
- [5] A theory of multineuronal dimensionality, dynamics and measurement View paper
- [6] Dimensionality of object representations in monkey inferotemporal cortex View paper
- [7] Representation learning via manifold flattening and reconstruction View paper
- [8] Manifold-adaptive dimension estimation revisited View paper
- [9] Statistical Analysis of Large-Scale Neural Representations View paper
- [10] Probing variability in a cognitive map using manifold inference from neural dynamics View paper
- [11] CH2O coordinate manifolds and their analysis for the paper "A global view of reactive coordinate manifolds from nonlinear dimensionality reduction" View paper
- [12] Neural active manifolds: nonlinear dimensionality reduction for uncertainty quantification View paper
- [13] Deep Convolutional Neural Network Compression based on the Intrinsic Dimension of the Training Data View paper
- [14] Intrinsic Dimension Estimation. View paper
- [15] IAN: Iterated Adaptive Neighborhoods for Manifold Learning and Dimensionality Estimation. View paper
- [16] Bayesian Manifold Regression View paper
- [17] Fermat Distances: Metric Approximation, Spectral Convergence, and Clustering Algorithms View paper
- [18] Weighted Linear Local Tangent Space Alignment via Geometrically Inspired Weighted PCA for Fault Detection View paper
- [19] UCET8: Universal Curvature Equivalence on Noncompact, Bounded, and Lorentzian Geometries View paper
- [20] Intrinsic Dimension Estimating Autoencoder (IDEA) Using CancelOut Layer and a Projected Loss View paper
- [21] A global view of reactive coordinate manifolds from nonlinear dimensionality reduction View paper
- [22] Local smoothing for manifold learning View paper
- [23] A continuous Formulation of intrinsic Dimension. View paper
- [24] Mitigating Class Imbalance in Long-Tailed Visual Recognition Through the Use of Intrinsic Dimensionality View paper
- [25] Estimating low dimensional dynamical models for molecules View paper
- [26] A scale-dependent measure of system dimensionality View paper
- [27] Sample size determination for GEE analyses of stepped wedge cluster randomized trials View paper
- [28] Critical generalized inverse participation ratio distributions View paper
- [29] High-dimensional geometry of population responses in visual cortex View paper
- [30] Gaussian partial information decomposition: Bias correction and application to high-dimensional data View paper
- [31] A security model for smart grid SCADA systems using stochastic neural network View paper
- [32] Assessing Neural Network Representations During Training Using Noise-Resilient Diffusion Spectral Entropy View paper
- [33] Disentangling Identifiable Features from Noisy Data with Structured Nonlinear ICA View paper
- [34] Variable noise and dimensionality reduction for sparse Gaussian processes View paper
- [35] Python for information theoretic analysis of neural data View paper
- [36] Estimating the functional dimensionality of neural representations View paper
- [37] Manifold Reconstruction of Differences: A Model-Based Iterative Statistical Estimation Algorithm With a Data-Driven Prior. View paper
- [38] Accuracy maximization analysis for sensory-perceptual tasks: Computational improvements, filter robustness, and coding advantages for scaled additive noise View paper