



Visualizing probabilistic models and data with Intensive Principal Component Analysis

Katherine N. Quinn^{a,1}, Colin B. Clement^a, Francesco De Bernardis^a, Michael D. Niemack^a, and James P. Sethna^a

^aDepartment of Physics, Cornell University, Ithaca, NY 14853-2501

Edited by William Bialek, Princeton University, Princeton, NJ, and approved May 31, 2019 (received for review October 5, 2018)

Unsupervised learning makes manifest the underlying structure of data without curated training and specific problem definitions. However, the inference of relationships between data points is frustrated by the “curse of dimensionality” in high dimensions. Inspired by replica theory from statistical mechanics, we consider replicas of the system to tune the dimensionality and take the limit as the number of replicas goes to zero. The result is intensive embedding, which not only is isometric (preserving local distances) but also allows global structure to be more transparently visualized. We develop the Intensive Principal Component Analysis (InPCA) and demonstrate clear improvements in visualizations of the Ising model of magnetic spins, a neural network, and the dark energy cold dark matter (Λ CDM) model as applied to the cosmic microwave background.

manifold learning | information theory | probabilistic models | probabilistic data | visualization

Visualizing high-dimensional data is a cornerstone of machine learning, modeling, big data, and data mining. These fields require learning faithful and interpretable low-dimensional representations of high-dimensional data and, almost as critically, producing visualizations which allow interpretation and evaluation of what was learned (1–4). Unsupervised learning, which infers features from data without manually curated data or specific problem definitions (5), is especially important for high-dimensional, big data applications in which specific models are unknown or impractical. For high dimensions, the relative distances between features become small and most points are orthogonal to one another (6). A trade-off between preserving local and global structure must often be made when inferring a low-dimensional representation. Classic manifold learning techniques include linear methods such as principal component analysis (PCA) (7) and multidimensional scaling (MDS) (8), which preserve global structure but at the cost of obscuring local features. Existing nonlinear manifold learning techniques, such as t-distributed stochastic network embedding (t-SNE) (9) and diffusion maps (10), preserve the local structure while maintaining only some qualitative global patterns such as large clusters. The uniform manifold approximation (UMAP) (11) better preserves topological structures in data, a global property.

In this article, we develop a nonlinear manifold learning technique which achieves a compromise between preserving local and global structure. We accomplish this by developing an isometric embedding for general probabilistic models based on the replica trick (12). Taking the number of replicas to zero, we reveal an intensive property—an information density characterizing the distinguishability of distributions—ameliorating the canonical orthogonality problem and “curse of dimensionality.” We then describe a simple, deterministic algorithm that can be used for any such model, which we call Intensive Principal Component Analysis (InPCA). Our method quantitatively captures global structure while preserving local distances. We first apply InPCA to the canonical Ising model of magnetism, which inspired the zero-replica limit. Next, we show how InPCA can capture and summarize the learning

trajectory of a neural network. Finally, we visualize the dark energy cold dark matter (Λ CDM) model as applied to the cosmic microwave background (CMB), using InPCA, t-SNE, and diffusion maps.

Model Manifolds of Probability Distributions

Any measurement obtained from an experiment with uncertainty can generally be understood as a probability distribution. For example, when some data x are observed with normally distributed noise ξ of variance σ^2 , under experimental conditions θ_j , a model is expressed as

$$x = f(\theta_j) + \xi \quad \text{where} \quad \mathcal{L}(\xi) \sim \mathcal{N}(0, \sigma^2), \quad [1]$$

and $f(\theta_i)$ is a prediction given the experimental conditions. This relationship is equivalent to saying that the probability of measuring data x given some conditions θ is

$$\mathcal{L}(x | \theta) \sim \mathcal{N}(f(\theta), \sigma^2). \quad [2]$$

More complicated noise profiles with asymmetry or correlations can be accommodated with this picture. Measurements without an underlying model can also be seen as distributions, where a measurement x_i with uncertainty σ can induce a probability $\mathcal{L}(x | x_i, \sigma)$ of observing new data x .

Significance

We introduce Intensive Principal Component Analysis (InPCA), a widely applicable manifold-learning method to visualize general probabilistic models and data. Using replicas to tune dimensionality in high-dimensional data, we use the zero-replica limit to discover a distance metric, which preserves distinguishability in high dimensions, and an embedding with superior visualization performance. We apply InPCA to the model of cosmology which predicts the angular power spectrum of the cosmic microwave background, allowing visualization of the space of model predictions (i.e., different universes).

Author contributions: K.N.Q., C.B.C., F.D.B., M.D.N., and J.P.S. designed research; K.N.Q. performed research; K.N.Q. contributed new reagents/analytic tools; K.N.Q. analyzed data; and K.N.Q., C.B.C., F.D.B., M.D.N., and J.P.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

Data deposition: All code used to generate the figures is available through public repositories on Github. Fig. 1 and Fig. 2 can be generated from code on https://github.com/katnquinn/Ising_ModelManifold, Fig. 3 can be generated from code on <https://github.com/katnquinn/IntensiveEmbedding>, Fig. 4 from code found on <https://github.com/katnquinn/1Spin>, and Fig. 5 from code on https://github.com/katnquinn/CMB_ModelManifold and from Code for Anisotropies in the Microwave Background software found at <https://camb.info/>.

¹To whom correspondence may be addressed. Email: knq2@cornell.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1817218116/-DCSupplemental.

Published online June 24, 2019.

We define a probabilistic model $\mathcal{L}(\mathbf{x}|\boldsymbol{\theta})$, the likelihood of observing data \mathbf{x} given parameters $\boldsymbol{\theta}$. The model manifold is defined as the set of all possible predictions, $\{\mathcal{L}(\mathbf{x}|\boldsymbol{\theta}_i)\}$, which is a surface parameterized by the model parameters $\{\boldsymbol{\theta}_i\}$. The parameter directions related to the longest distances along the model manifold have been shown to predict emergent behavior (how microscopic parameters lead to macroscopic behavior) (13). We will see that InPCA orders its principal components by the length of the model manifold along their direction, highlighting global structure. The boundaries of the model manifold represent simplified models which retain predictive power (14), and the constraint of data lying near the model manifold has been used to optimize experimental design (15). In this article, we study the Ising model, which defines probabilities of spin configurations given interaction strengths; a neural network, which predicts the probability of an image representing a single handwritten digit given weights and biases; and Λ CDM, which predicts the distribution of CMB radiation given fundamental constants of nature.

Hypersphere Embedding

We promised an embedding which both is isometric and preserves global structures. We satisfy the first promise by considering the hypersphere embedding

$$\{z_{\mathbf{x}}(\boldsymbol{\theta}_i)\} = \left\{ 2\sqrt{\mathcal{L}(\mathbf{x}|\boldsymbol{\theta}_i)} \right\}, \quad [3]$$

where the normalization constraint of $\mathcal{L}(\mathbf{x}|\boldsymbol{\theta})$ forces $z_{\mathbf{x}}$ to lie on the positive orthant of a sphere. A natural measure of distance on the hypersphere is the Euclidean distance, in this case also known as the Hellinger divergence (16)

$$\begin{aligned} d^2(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) &= \|z(\boldsymbol{\theta}_1) - z(\boldsymbol{\theta}_2)\|^2 \\ &= 8 \left(1 - \sqrt{\mathcal{L}(\mathbf{x}|\boldsymbol{\theta}_1)} \cdot \sqrt{\mathcal{L}(\mathbf{x}|\boldsymbol{\theta}_2)} \right)^2, \end{aligned} \quad [4]$$

where \cdot represents the inner product over \mathbf{x} . Now we can see that the hypersphere embedding is isometric: The Euclidean metric of this embedding is equal to the Fisher information metric \mathcal{I} of the model manifold (17),

$$d^2(z_i, z_i + dz_i) = \sum_i dz_i dz_i = \sum_{kl} \mathcal{I}_{kl} d\theta_k d\theta_l. \quad [5]$$

The Fisher information metric (FIM) is the natural metric of the model manifold (18), so the hypersphere embedding preserves the local structure of the manifold defined by $\mathcal{L}(\mathbf{x}|\boldsymbol{\theta})$.

As the dimension of the data increases, almost all features become orthogonal to each other, and most measures of distance lose their ability to discriminate between the smallest and largest distances (19). For the hypersphere embedding, we see that as the dimension of \mathbf{x} increases, the inner product in the Hellinger distance of Eq. 4 becomes smaller as the probability is distributed over more dimensions. In the limit of large dimension, all nonidentical pairs of points become orthogonal and equidistant around the hypersphere (a constant distance $\sqrt{8}$ apart), frustrating effective dimensional reductions and visualization.

To illustrate this problem with the hypersphere embedding, consider the Ising model, which predicts the likelihood of observing a particular configuration of binary random variables (spins) on a lattice. The probability of a spin configuration is determined by the Boltzmann distribution and is a function of a local pairwise coupling and a global applied field. The dimension is determined by the number of spin configurations, 2^N ,

where N is the number of spins. Holding temperature fixed at one, we vary h and J : external magnetic field ($h \in (-1.3, 1.3)$) and nearest-neighbor coupling ($J \in (-0.4, 0.6)$), using a Monte Carlo method weighted by Jeffrey's prior to sample 12,000 distinct points. From the resulting set of parameters, we compute $X_{ij} = \{z_i(\boldsymbol{\theta}_j)\}$ using the Boltzmann distribution and visualize the model manifold in the N -sphere embedding of Eq. 3 by projecting the predictions onto the first three principal components of X . Fig. 1A shows this projection of the model manifold of a 2×2 Ising model which is embedded in 2^4 dimensions. Fig. 1B shows a larger, 4×4 Ising model, of dimension 2^{16} . As the dimension is increased from 2^4 to 2^{16} , we see the points starting to wrap around the hypersphere, becoming increasingly equidistant and less distinguishable.

A natural way to increase the dimensionality of a probabilistic model is by drawing multiple samples from the distribution. If D is the dimension of \mathbf{x} , then N identical draws from the distribution will have dimension D^N . The more samples drawn, the easier it is to distinguish between distributions, mimicking the curse of dimensionality for large systems. We see this demonstrated for our Ising model in Fig. 1C, where we drew four replica samples from the same model. Note that compared with the original 2×2 model, the model manifold of the four-replica 2×2 model “wraps” more around the hypersphere, just like the larger, 4×4 Ising model. High-dimensional systems have “too much information,” in the same way that large numbers of samples have too much information. In the next section, we consider the contraposition of the insight that a large number of replicas lead to the curse of dimensionality and discover an embedding which not only is isometric but also ameliorates the high-dimensional wrapping around the n -sphere.

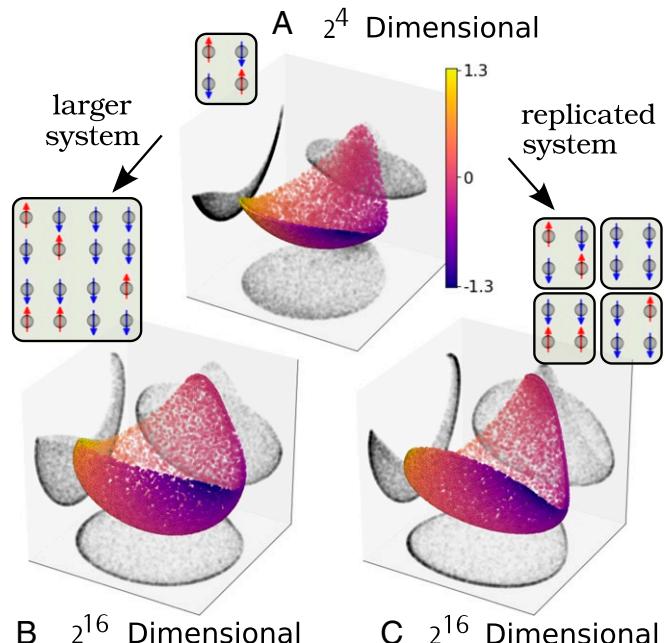


Fig. 1. (A–C) Hypersphere embedding, illustrating an embedding of the 2D Ising model. Points were generated through a Monte Carlo sampling and visualized by projecting the probability distributions onto the first three principal components (28). The points are colored by magnetic field strength. As the system size increases from 2×2 to 4×4 , the orthogonality problem is demonstrated by an increase in “wrapping” around the hypersphere. This effect can also be produced by instead considering four replicas of the original system, motivating the replica trick which takes the embedding dimension or number of replicas to zero.

Replica Theory and the Intensive Embedding

We saw in Fig. 1 that increasing the dimension of the data led to a saturation of the distance function Eq. 4. This problem is referred to as the loss of relative contrast or the concentration of distances (19), and to overcome it requires a non-Euclidean distance function, discussed below. In the previous section we saw the same saturation of distance could be achieved by adding replicas, increasing the embedding dimension. Fig. 2A shows this process taken to an extreme: the model manifold of the 2×2 Ising model with the number of replicas taken to infinity. All of the points cluster together, obscuring the fact that the underlying manifold is 2D. To cure the abundance of information which makes all points on the hypersphere equidistant, we seek an intensive distance, such as the distance per number of replicas observed. Next, because the limit of many replicas artificially leads to the same symptoms of the curse of dimensionality, we consider the limit of zero replicas, a procedure which is often used in the study of spin glasses and disordered systems (20). Fig. 2B shows the result of this analysis, the intensive embedding, where the distance concentration has been cured, and the inherent 2D structure of the Ising model has been recovered.

To find the intensive embedding, we must first find the distance between replicated models. The likelihood for N replicas of a system is given by their product

$$\mathcal{L}(\{x_1, \dots, x_N\} | \theta)^{(N)} = \mathcal{L}(x_1 | \theta) \cdots \mathcal{L}(x_N | \theta), \quad [6]$$

where the set $\{x_1, \dots, x_N\}$ represents the observed data in the replicated systems. Writing the inner product or cosine angle between two distributions as

$$\langle \theta_1; \theta_2 \rangle = \sqrt{\mathcal{L}(x | \theta_1)} \cdot \sqrt{\mathcal{L}(x | \theta_2)}, \quad [7]$$

and using Eq. 4, the distance per replica d_N^2 between two points on the model manifold is

$$d_N^2(\theta_1, \theta_2) = \frac{d^2(\theta_1, \theta_2)}{N} = -8 \frac{\langle \theta_1; \theta_2 \rangle^N - 1}{N}. \quad [8]$$

We are now poised to define the intensive distance by taking the number of replicas to zero:

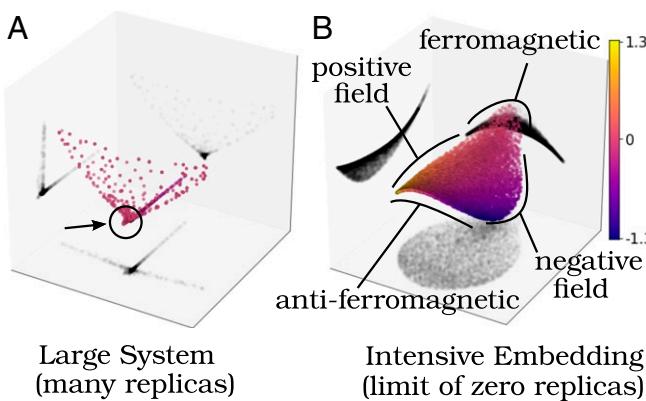


Fig. 2. Replicated Ising model illustrating the derivation of our intensive embedding. All points are colored by magnetic field strength. (A) Large dimensions are characterized by large system sizes; here we mimic a 128×128 Ising model which is of dimension 2^{128^2} . The orthogonality problem becomes manifest as all points are effectively orthogonal, producing a useless visualization with all points clustered in the cusp. (B) Using replica theory, we tune the dimensionality of the system and consider the limit as the number of replicas goes to zero. In this way, we derive our intensive embedding. Note that the z axis reflects a negative-squared distance, a property which allows violations of the triangle inequality and is discussed in the text.

$$d_I^2(\theta_1, \theta_2) = \lim_{N \rightarrow 0} d_N^2(\theta_1, \theta_2) = -8 \log \langle \theta_1; \theta_2 \rangle. \quad [9]$$

The last equality is achieved using the standard trick in replica theory, $(x^N - 1)/N \rightarrow \log x$ as $N \rightarrow \infty$, which is a basis trick used to solve challenging problems in statistical physics (20). The trick is most evident using the identity $x^N = \exp(\log Nx) \approx 1 + N \log x$. One can check that the intensive distance is isometric,

$$d_I^2(\theta, \theta + \delta\theta) = \delta\theta^\alpha \delta\theta^\beta g_{\alpha\beta} = \delta\theta^\alpha \delta\theta^\beta \mathcal{I}_{\alpha\beta}, \quad [10]$$

where again \mathcal{I} is the Fisher information metric in Eq. 5, so that we can be confident the intensive embedding distance preserves local structures.

Importantly, the intensive distance does not satisfy the triangle inequality (and is thus non-Euclidean): The distance between points on the hypersphere can go to infinity, rather than lie constrained to the finite radius of the hypersphere embedding. Because of this, the intensive embedding can overcome the loss of relative contrast (19) discussed at the beginning of this section. Distances in the intensive embedding maintain distinguishability in high dimensions, as illustrated in Fig. 2B, wherein the 2D nature of the Ising model has been recovered. We hypothesize that this process, which cures the curse of dimensionality for models with too many samples, will also cure it for models with intrinsically high dimensionality. The intensive distance obtained here is proportional to the Bhattacharyya distance (21). Considering the zero-replica limit of the Hellinger divergence, we discovered a way to derive the Bhattacharyya distance. The importance of this is discussed further in the following section.

Connection to Least Squares. Consider the concrete and canonical paradigm of models $f_i(\theta)$ with data points x_i and additive white Gaussian noise, usually called a nonlinear least-squares model. The likelihood $\mathcal{L}(x | \theta)$ is defined by

$$-\log \mathcal{L}(x | \theta) = \sum_i \frac{(f_i(\theta) - x_i)^2}{2\sigma_i^2} + \log \mathcal{Z}(\theta), \quad [11]$$

where \mathcal{Z} sets the normalization. A straightforward evaluation of the intensive distance given by Eq. 9 finds for the case of nonlinear least squares that

$$d_I^2(\theta_1, \theta_2) = \sum_i \frac{(f_i(\theta_1) - f_i(\theta_2))^2}{\sigma_i^2}, \quad [12]$$

so that the intensive distance is simply the variance-scaled Euclidean distance between model predictions.

Intensive Principal Component Analysis

Classical PCA takes a set of data examples and infers features which are linearly uncorrelated. (7). The features to be analyzed with PCA are compared via their Euclidean distance. Can we generalize this comparison to use our intensive embedding distance? Given a matrix of data examples $X \in \mathbb{R}^{m \times p}$ (with features along the rows), PCA first requires the mean-shifted matrix $M_{ij} = X_{ij} - \bar{X}_i = PX$, where $P_{ij} = \delta_{ij} - 1/p$ is the mean-shift projection matrix and p is the number of sampled points. The covariance and its eigenvalue decomposition are then

$$\text{cov}(X, X) = \frac{1}{p} M^T M = X^T P P X = V \Sigma V^T, \quad [13]$$

where the orthogonal columns of the matrix V are the natural basis onto which the rows of M are projected,

$$MV = (UDV^T)V = UD = U\sqrt{\Sigma}, \quad [14]$$

where the columns of $U\sqrt{\Sigma}$ are called the principal components of the data X .

The principal components can also be obtained from the cross-covariance matrix, MM^T , since

$$MM^T = PXX^TP = (UDV^T)(UDV^T)^T = U\Sigma U^T. \quad [15]$$

The eigenbasis U of the cross-covariance is the natural basis for the components of the data, and the eigenbasis V of the covariance is the natural basis of the data points. For us this flexibility is invaluable, as the cross-covariance is more natural for expressing the distances between distributions of different parameters.

Writing our data matrix as $X_{ij} = z_i(\theta_j)$ using Eq. 3 for replicated systems, the cross-covariance is

$$\begin{aligned} (MM^T)_{ij}^{(N)} &= (PXX^TP)_{ij} \\ &= (z(\theta_i) - \bar{z}) \cdot (z(\theta_j) - \bar{z}) \\ &= 4 \langle \theta_i; \theta_j \rangle^N + \frac{4}{p^2} \sum_{k, k'=1}^p \langle \theta_k; \theta_{k'} \rangle^N \\ &\quad - \frac{4}{p} \sum_{k=1}^p \left(\langle \theta_i; \theta_k \rangle^N + \langle \theta_j; \theta_k \rangle^N \right), \end{aligned} \quad [16]$$

where \bar{z} is the average over all sampled parameters, and we used the definition of z in Eq. 6. As with the intensive embedding, we can take the limit as the number of replicas goes to zero to find

$$W_{ij} = \lim_{N \rightarrow 0} \frac{1}{N} (MM^T)_{ij}^{(N)}. \quad [17]$$

Explicitly, the intensive cross-covariance matrix is

$$\begin{aligned} W_{ij} &= 4 \log \langle \theta_i; \theta_j \rangle + \frac{4}{p^2} \sum_{k, k'=1}^p \log \langle \theta_{k'}; \theta_k \rangle \\ &\quad - \frac{4}{p} \sum_{k=1}^p (\log \langle \theta_i; \theta_k \rangle + \log \langle \theta_j; \theta_k \rangle) \end{aligned} \quad [18]$$

$$= (PLP)_{ij}, \quad [19]$$

where $L_{i,j} = 4 \log \langle \theta_i; \theta_j \rangle$ and P is the same projection matrix as defined above. In taking the limit of zero replicas, the structure of the cross-covariance has transformed

$$PXX^TP \xrightarrow{N \rightarrow 0} PLP, \quad [20]$$

and thus the symmetric Wishart structure is lost. It is therefore possible to obtain negative eigenvalues in this decomposition, which give rise to imaginary components in the projections. Note the similarity between the form of this cross-covariance and the double-centered distance matrix used in PCA and multidimensional scaling (MDS). This arises because both InPCA and PCA/MDS rely on mean shifting the input data before finding an eigenbasis. Thus, we view InPCA as a natural generalization of PCA to probability distributions and MDS to non-Euclidean embeddings.

In summary, InPCA is achieved by the following procedure: (i) Compute the cross-covariance matrix from a set of probability samples: Compute W_{ij} as derived in Eq. 18. (ii) Compute the eigenvalue decomposition $W = U\Sigma U^T$. (iii) Compute the coordinate projections, $T = U\sqrt{\Sigma}$. (iv) Plot the projections using the columns of T .

Neural Network MNIST Digit Classifier. To demonstrate the utility of InPCA, we use it to visualize the training of a two-layer convolution neural network (CNN), constructed using TensorFlow (22), trained on the MNIST dataset of hand-written digits (23). A set of 55,000 images was used to train the network, which was then used to predict the likelihood that an additional set of 10,000 images is classified each as a specific digit between 0 and 9. We use softmax (24) to calculate the probabilities from the category estimates supplied by the network. The CNN defines the likelihood $\mathcal{L}(x|\theta)$ that some input image θ contains the image of a particular handwritten digit x . The InPCA projections of the CNN output in Fig. 3 visualize the clustering learned by the CNN as a function of the number of learning epochs. The initialized network's model manifold shows no knowledge of the digits (colored dots), but as training commences, the network separates digits into separate regions of its manifold (Movie S1). InPCA can be used as a fast, interpretable, and deterministic method for qualitatively evaluating what a neural network has learned.

Properties of the Intensive Embedding and InPCA

The space characterized by our intensive embedding has two weird properties: First, it is formally 1D, yet there are multiple orthogonal directions upon which it can be projected; and second, it is Minkowski-like, in that it has negative squared distances, violating the triangle inequality. We posit that, fundamentally, this second property is what allows InPCA to cure the orthogonality problem.

We begin with a discussion of the the 1D nature of the embedding space. The embedding dimension is given by D^N , where D is the original dimension of data x and N is the number of replicas. In the case of noninteger replicas the space becomes “fractional” in dimension and in the limit of zero replicas ultimately goes to one. However, it is still possible to obtain projections themselves along the dominant components of this space, by leveraging the cross-covariance instead of the covariance, summarized in step *ii* of our algorithm. Visualizations produced by InPCA are cross-sections of a space of the dimension equal to the

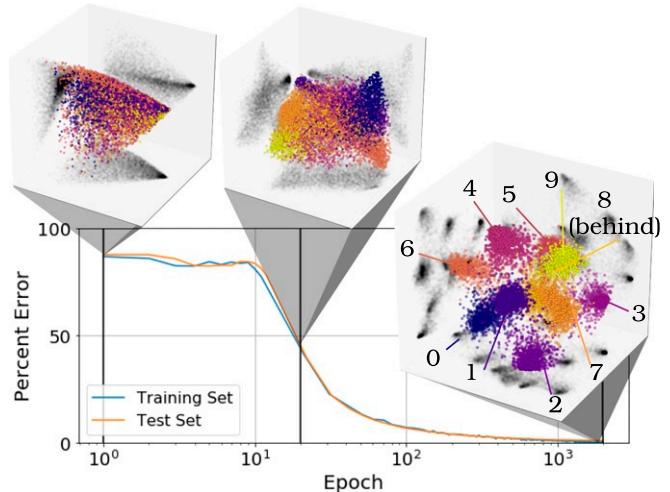


Fig. 3. Stages of training a CNN. Each point in the 3D projections represents one of 10,000 test images supplied to the CNN (29). At the first epoch, the neural network is untrained and so is unable to reliably classify images, with about a 90% error rate—an effect reflected in the cloud of points. As training progresses and error rate decreases, the cloud begins to cluster as shown by InPCA at the 20th epoch. Finally, when completely trained, the clustered regions are manifest at the 2000th epoch with 10 clusters representing the 10 digits.

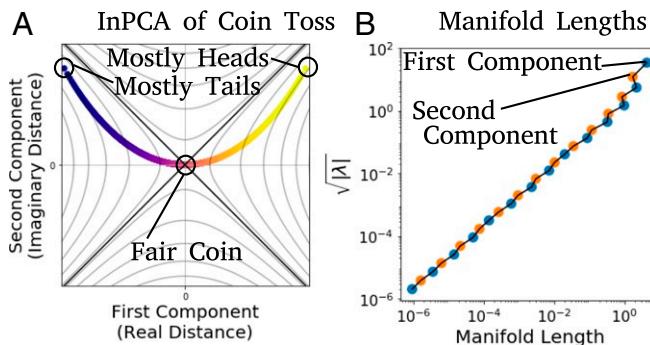


Fig. 4. InPCA visualization of biased coins (30). (A) The first two InPCA components correspond to the coin bias and variance, yet the first one is real and the second one is imaginary (the aspect ratio between axes is one). The contour lines represent constant distances from a fair coin and are hyperbolas: Points can be a finite distance from a fair coin yet an infinite distance from each other. (B) The ordered eigenvalues correspond to the manifold lengths, illustrating the hierarchical nature of the components extracted from InPCA.

number of sampled points of the model manifold p , instead of the dimension D or D^N .

In the limit of zero replicas in Eq. 18, the positive-definite, Wishart structure of the cross-covariance matrix is lost. It is therefore possible to have negative squared distances. The non-Euclidean nature of the embedding (flat, but Minkowski-like) does not suffer from the concentration of distances which plagues Euclidean measures in high dimensions, thus allowing the model manifold to be “unwound” from the N -sphere and for InPCA to produce useful, low-dimensional representations.

Finally, the eigenvalues of InPCA correspond to the cross-sectional widths of the model manifold. We see this quite explicitly with the following example of a biased coin (specifically, in Fig. 4B) where the eigenvalues extracted from InPCA map directly to the manifold widths measured along the direction of the corresponding InPCA eigenvector. Therefore, we see that InPCA produces a hierarchy of directions, ordered by the global widths of the model manifold. Note that, as with classical PCA, this correspondence depends on how faithfully the model manifold was originally sampled; that is, InPCA can tell you about the structure of the manifold only from observed points.

Biased Coins. To illustrate the properties of InPCA, we use it to visualize a simple probabilistic model, that of a simple biased coin. A biased coin has one parameter, the odds ratio of heads to tails, and so forms a 1D manifold. Fig. 4A shows the first two InPCA components for the manifold of biased coins, for 2,000 sampled points with probabilities uniformly spread between 0 and 1 (excluding the endpoints, since they are orthogonal and thus are infinitely far apart). The two extracted InPCA components correspond to the bias and variance of the coin, respectively. The hierarchy of components extracted from InPCA therefore corresponds to known features of the model (i.e., they are meaningful).

The importance of the negative-squared distances is illustrated in Fig. 4. The contour lines representing constant distances from a fair coin and are hyperbolas: Points can be a finite distance from a fair coin yet an infinite distance from each other. As two oppositely biased coins become increasingly biased, their distance from each other can go to infinity (because an outcome of a coin which always lands on heads will never be the same as an outcome of a coin which always land on tails) yet all points remain a finite distance from a fair coin. Note that all points are in the left and right portions of Fig. 4A, representing net positive distances (the intensive pairwise distances are all positive).

Comparing with t-SNE and Diffusion Maps. We compare our manifold learning technique to two standard methods, t-SNE and the diffusion maps, by applying each one to the six-parameter Λ CDM cosmological model predictions of the CMB. The Λ CDM predicts $\mathcal{L}(x|\theta)$, where x represents fluctuations in the CMB, and θ are the different cosmological parameters (i.e., it predicts the angular power spectrum of temperature and polarization anisotropies in sky maps of the CMB). Observations of the CMB from telescopes on satellites, balloons, and the ground provide thousands of independent measurements from large angular scales to a few arcminutes, which are used to fit model parameters. Here we consider only CMB observations from the 2015 Planck data release (25). The Λ CDM model we consider has six parameters, the Hubble constant (H_0) which we sampled in a range of $20-100 \text{ km}\cdot\text{s}^{-1}\cdot\text{Mpc}^{-1}$, the physical baryon density ($\Omega_b h^2$) and the physical cold dark matter density ($\Omega_c h^2$) both sampled from 0.0009 to 0.8, the primordial fluctuation amplitude (A_s) sampled from 10^{-11} to 10^{-8} , the scalar spectral index (η) sampled from 0 to 0.98, and the optical depth at reionization (τ) sampled from 0.001 to 0.9.

To determine the likelihood functions, we use the Code for Anisotropies in the Microwave Background software package to generate power spectra (26). We perform a Monte Carlo sampling of 50,000 points around the best-fit parameters provided by the 2015 Planck data release (25), with sample weights based on the intensive distance to the best fit.

In Fig. 5 we show the first two components of the manifold embedding for InPCA, t-SNE, and diffusion maps. To apply t-SNE and the diffusion map to probabilistic data we must provide a distance. We therefore use our intensive distance, from Eq. 9, for consistency and ease of comparison. In all three cases, the first component from each method is directly related to the primordial fluctuation amplitude A_s , which reflects the amplitude of density fluctuations in the early universe and is the dominant feature in real data (25). The second InPCA component

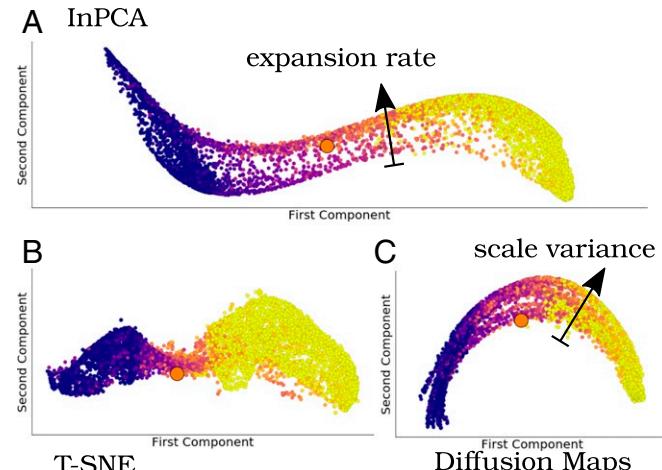


Fig. 5. Model manifold of the six-parameter Λ CDM cosmological model predictions of temperature and polarization power spectra in the CMB using InPCA, t-SNE, and the diffusion map. Axes reflect the true aspect ratio from extracted components in all cases. Here the model manifold is colored by the primordial fluctuation amplitude, the most prominent feature in CMB data. (A) InPCA extracts, as the first and second component, this amplitude term as well as the Hubble constant. These parameters control the two most dominant features in the Planck data and so reflect a physically meaningful hierarchy of importance. In contrast, (B) t-SNE extracts only the amplitude term and (C) the diffusion map extracts the amplitude term and a different parameter, the scalar spectral index η , which reflects the scale variance of the density fluctuations in the early universe. In all plots, the orange point represents our universe, as represented by Planck 2015 data.

predicts the Hubble constant, whereas the diffusion map predicts the scalar spectral index (a reflection of the size variance of primordial density fluctuations). In all cases, the projected components were plotted against the corresponding parameters to determine correlations, such as how one can see that A_s corresponds with the first component in all three cases. Detailed plots and correlation coefficients for all three methods are provided in *SI Appendix*.

Such stark differences between manifold learning methods are surprising, as all techniques aim to extract important features in the data distribution, i.e., important geometric features in the manifolds. Given the ranges of sampled parameters, one would expect the variation in the Hubble constant to relate in some way to one of the dominant components, as it does for InPCA. Figures illustrating the effect of different parameters are provided in *SI Appendix*, following results from ref. 27.

There are two important differences between InPCA and other methods. First, InPCA has no tunable parameters and yields a geometric object defined entirely by the model distribution. For example, t-SNE embeddings rely on parameters such as the perplexity, a learning rate, and a random seed (yielding non-deterministic results), and the diffusion maps rely on a diffusion parameter and choice of diffusion operator, all of which must be

manually optimized to obtain good results. Second, t-SNE and diffusion maps embed manifolds in Euclidean spaces in a way which aims to preserve local features. However, InPCA seeks to preserve both global and local features, by embedding manifolds in a non-Euclidean space.

Summary

In this article, we introduce an unsupervised manifold learning technique, InPCA, which captures low-dimensional features of general, probabilistic models with wide-ranging applicability. We consider replicas of a probabilistic system to tune its dimensionality and consider the limit of zero replicas, deriving an intensive embedding that ameliorates the canonical orthogonality problem. Our intensive embedding provides a natural, meaningful way to characterize a symmetric distance between probabilistic data and produces a simple, deterministic algorithm to visualize the resulting manifold.

ACKNOWLEDGMENTS. We thank Mark Transtrum for guidance on algorithms and for useful conversations. We thank Pankaj Mehta for pointing out the connection to MDS. K.N.Q. was supported by a fellowship from the Natural Sciences and Engineering Research Council of Canada, and J.P.S. and K.N.Q. were supported by the National Science Foundation (NSF) through NSF Grants DMR-1312160 and DMR-1719490. M.D.N. was supported by NSF Grant AST-1454881.

1. M. F. De Oliveira, H. Levkowitz, From visual data exploration to visual data mining: A survey. *IEEE Trans. Visualization Comput. Graphics* **9**, 378–394 (2003).
2. S. Liu, D. Maljovec, B. Wang, P. T. Bremer, V. Pascucci, Visualizing high-dimensional data: Advances in the past decade. *IEEE Trans. Visualization Comput. Graphics* **23**, 1249–1268 (2017).
3. J. A. Lee, M. Verleysen, *Nonlinear Dimensionality Reduction* (Springer, New York, NY, 2007).
4. A. Zimek, E. Schubert, H. P. Kriegel, A survey on unsupervised outlier detection in high-dimensional numerical data. *Stat. Anal. Data Mining ASA Data Sci. J.* **5**, 363–387 (2012).
5. K. P. Murphy, *Machine Learning: A Probabilistic Perspective* (The MIT Press, 2012).
6. H. P. Kriegel, P. Kröger, A. Zimek, Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering. *ACM Trans. Knowl. Discov. Data* **3**, 1–58 (2009).
7. H. Hotelling, Analysis of a complex of statistical variables into principal components. *J. Educ. Psychol.* **24**, 417–441 (1933).
8. W. S. Torgerson, Multidimensional scaling: I. Theory and method. *Psychometrika* **17**, 401–419 (1952).
9. L. van der Maaten, G. Hinton, Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
10. R. R. Coifman et al., Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 7426–7431 (2005).
11. L. McInnes, J. Healy, J. Melville, Umap: Uniform manifold approximation and projection for dimension reduction. arXiv:1802.03426 (6 December 2018).
12. M. Mézard, G. Parisi, M. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, 1986).
13. B. B. Machta, R. Chachra, M. K. Transtrum, J. P. Sethna, Parameter space compression underlies emergent theories and predictive models. *Science* **342**, 604–607 (2013).
14. M. K. Transtrum, P. Qiu, Model reduction by manifold boundaries. *Phys. Rev. Lett.* **113**, 098701 (2014).
15. M. K. Transtrum et al., Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *J. Chem. Phys.* **143**, 010901 (2015).
16. E. Hellyer, Neue Begründung der Theorie quadratischer Formen von unendlichvielen Veränderlichen. *J. Reine Angew. Math.* **136**, 210–271 (1909).
17. M. Gromov, In a search for a structure, part 1: On entropy. *Entropy* **17**, 1273–1277 (2013).
18. S. Amari, H. Nagaoka, *Translations of Mathematical Monographs: Methods of Information Geometry* (Oxford University Press, 2000), vol. 191.
19. K. Beyer, J. Goldstein, R. Ramakrishnan, U. Shaft, “When is “nearest neighbor” meaningful?” in *Database Theory—ICDT’99*, C. Beeri, P. Buneman, Eds. (Springer Berlin Heidelberg, Berlin, Heidelberg, Germany, 1999), pp. 217–235.
20. G. Parisi, Infinite number of order parameters for spin-glasses. *Phys. Rev. Lett.* **43**, 1754–1756 (1979).
21. A. Bhattacharyya, On a measure of divergence between two multinomial populations. *Sankhyā Indian J. Stat. (1933–1960)* **7**, 401–406 (1946).
22. M. Abadi et al., TensorFlow: Large-scale machine learning on heterogeneous systems. <https://www.tensorflow.org/>. Accessed 1 December 2017.
23. Y. LeCun, C. Cortes, C. J. Burges, “MNIST database”. <http://yann.lecun.com/exdb/mnist/>. Accessed 1 December 2017.
24. C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, New York, NY, 2006).
25. Planck Collaboration, Planck 2015 results - i. Overview of products and scientific results. *A&A* **594**, A1 (2016).
26. A. Lewis, A. Challinor, A. Lasenby, Efficient computation of cosmic microwave background anisotropies in closed Friedmann-Robertson-Walker models. *Astrophys. J.* **538**, 473–476 (2000).
27. W. Hu, CMB tutorials. <http://background.uchicago.edu/>. Accessed 1 August 2018.
28. K. Quinn, Ising Model Manifold. GitHub. <https://github.com/katnquinn/Ising-ModelManifold>. Deposited 23 July 2018.
29. K. Quinn, Intensive Embedding. GitHub. <https://github.com/katnquinn/IntensiveEmbedding>. Deposited 11 March 2019.
30. K. Quinn, 1 Spin. GitHub. <https://github.com/katnquinn/1Spin>. Deposited 13 March 2019.