



Local conformal autoencoder for standardized data coordinates

Erez Peterfreund^{a,1}, Ofir Lindenbaum^{b,1}, Felix Dietrich^c, Tom Bertalan^d, Matan Gavish^a, Ioannis G. Kevrekidis^d, and Ronald R. Coifman^{b,2}

^aSchool of Computer Science and Engineering, Hebrew University of Jerusalem, Jerusalem 9190401, Israel; ^bProgram in Applied Mathematics, Yale University, New Haven, CT 06520; ^cDepartment of Informatics, Technical University of Munich, 80333 Munich, Germany; and ^dDepartment of Chemical and Biomolecular Engineering, Johns Hopkins University, Baltimore, MD 21218

Contributed by Ronald R. Coifman, September 20, 2020 (sent for review July 14, 2020; reviewed by Richard Baraniuk and Guillermo Sapiro)

We propose a local conformal autoencoder (LOCA) for standardized data coordinates. LOCA is a deep learning-based method for obtaining standardized data coordinates from scientific measurements. Data observations are modeled as samples from an unknown, nonlinear deformation of an underlying Riemannian manifold, which is parametrized by a few normalized, latent variables. We assume a repeated measurement sampling strategy, common in scientific measurements, and present a method for learning an embedding in \mathbb{R}^d that is isometric to the latent variables of the manifold. The coordinates recovered by our method are invariant to diffeomorphisms of the manifold, making it possible to match between different instrumental observations of the same phenomenon. Our embedding is obtained using LOCA, which is an algorithm that learns to rectify deformations by using a local z-scoring procedure, while preserving relevant geometric information. We demonstrate the isometric embedding properties of LOCA in various model settings and observe that it exhibits promising interpolation and extrapolation capabilities, superior to the current state of the art. Finally, we demonstrate LOCA's efficacy in single-site Wi-Fi localization data and for the reconstruction of three-dimensional curved surfaces from two-dimensional projections.

manifold learning | autoencoder | dimensionality reduction | canonical coordinates

1. Introduction

Reliable, standardized tools for analyzing complex measurements are crucial for science in the data era. Experimental data often consist of multivariate observations of a physical object that can be represented as an unknown Riemannian manifold. A key challenge in data analysis involves converting the observations into a meaningful and, hopefully, intrinsic parametrization of this manifold. For example, in astrophysics, one is interested in a representation that is coherent with the material composition of stars, based on measurable, high-dimensional spectroscopic data (1, 2). This type of challenge has typically been studied under the broader umbrella of dimensionality reduction and manifold learning, where numerous algorithmic solutions have been proposed (3–12). These methods rely on statistical or geometrical assumptions and aim to reduce the dimension, while preserving different affinities of the observed high-dimensional data.

In this paper, we focus on data obtained from several observation modalities measuring a complex system. These observations are assumed to lie on a path-connected manifold, which is parameterized by a small number of latent variables. We assume that the measurements are obtained via an unknown nonlinear measurement function observing the inaccessible manifold. The task is then to invert the unknown measurement function, so as to find a representation that provides a standardized parametrization of the manifold. In general, this form of blind inverse problem may not be feasible. Fortunately, in many cases, one can exploit a localized measurement strategy, suggested in ref. 13,

to extract an embedding into internally standardized (z-scored) latent variables.

Toward a useful formulation of our problem, we note that in numerous real-world scenarios, it is possible to capture data using a localized burst sampling strategy (14–21). As motivation for this type of burst sampling, we describe a toy experiment (Fig. 1). Consider the task of recovering the geometry of a curved two-dimensional (2D) homogeneous surface in three dimensions using a laser beam which heats the surface locally at several positions. Here a burst is realized through the brief local isotropic propagation of heat around each laser impact location (each data point), which can be visualized as a local ellipse by a thermal camera. Now, the task is to recover the curved geometry of the surface in three dimensions using the collection of observed local 2D ellipses.

More generally, our strategy is realized by measuring such brief bursts, which are modeled as local isotropic perturbations added to each state in the inaccessible latent manifold. The bursts provide information on the local variability in the neighborhood of each data point. Thus, they can be used to estimate the Jacobian (modulo an orthogonal transformation, as we will discuss) of the unknown measurement function. The authors of ref. 13 use such bursts and suggest a scheme to recover a

Significance

A fundamental issue in empirical science is the ability to calibrate between different types of measurements/observations of the same phenomenon. This naturally suggests the selection of canonical variables, in the spirit of principal components, to enable matching/calibration among different observation modalities/instruments. We develop a method for extracting standardized, nonlinear, intrinsic coordinates from measured data, leading to a generalized isometric embedding of the observations. This is achieved through a local burst data acquisition strategy that allows us to capture the local z-scored structure. We implement this method using a local conformal autoencoder architecture and illustrate it computationally. The proposed embedding is fast, parallelizable, easy to implement using existing open-source neural network implementations and exhibits surprising interpolation and extrapolation capabilities.

Author contributions: O.L., M.G., I.G.K., and R.R.C. designed research; E.P., O.L., F.D., and T.B. performed research; E.P., O.L., and F.D. contributed new reagents/analytic tools; E.P., O.L., F.D., and T.B. analyzed data; and E.P., O.L., M.G., I.G.K., and R.R.C. wrote the paper.

Reviewers: R.B., Rice University; and G.S., Duke University.

The authors declare no competing interest.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹E.P. and O.L. contributed equally to this work.

²To whom correspondence may be addressed. Email: coifman-ronald@yale.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2014627117/-DCSupplemental>.

First published November 23, 2020.

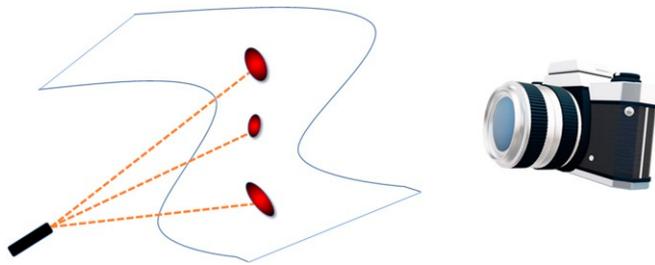


Fig. 1. An example motivating LOCA: learning the stereographic shape of a surface. We consider a laser beam used to locally heat the surface at several positions. A thermal camera measures the brief isotropic propagation of heat around each location. By scanning a thermal image, we can identify the neighborhoods of each position, which we define as our “bursts.” LOCA uses these bursts to invert the projection and recover a scaled version of the curved surface.

representation that is invariant to the unknown transformation. Specifically, they use a local Mahalanobis metric, combined with eigenvectors of an anisotropic Laplacian, in order to extract a desired embedding.

Solutions such as ref. 13 and extensions such as refs. 20 and 22–25 can be used. There remain, however, several challenges: 1) they require a dense sampling of the deformed manifold, 2) they deform the representation due to inherent boundary effects, and 3) they do not extend easily to unseen samples. To overcome these limitations, we introduce the concept of a local conformal autoencoder (LOCA), which is a deep learning-based algorithm, specifically suited to burst measurements. LOCA is realized using an encoder–decoder pair, where the encoder attempts to find a mapping such that each burst is locally whitened (z-scored). By taking into account a reconstruction loss, our decoder ensures that crucial geometric information is preserved in the embedding. We have found LOCA to be scalable, easy to implement, and parallelizable using the existing deep-learning open-source implementations. We provide empirical evidence that the LOCA embedding is approximately isometric to the latent manifold and extrapolates reliably to unseen samples. We discuss a scheme to automatically tune the minimal embedding dimension of LOCA and demonstrate its precision in two real data problems.

The contributions in this paper are as follows: 1) We show that the localized sampling strategy (bursts at a given scale) generates a consistent Riemannian structure; under certain conditions, it also allows inverting the unknown measurement function (modulo a shift and orthogonal transformation). 2) We present a two-step optimization scheme for learning a parametric mapping, which is approximately an isometry of the latent manifold. 3) We observe that this optimization problem can be approximately solved by training a deep autoencoder. 4) We present empirical evidence that the resulting algorithm allows surprisingly accurate interpolation and extrapolation. While most popular nonlinear dimension-reduction algorithms cannot be feasibly applied to large datasets, we note that the algorithm proposed here can be used to embed large datasets using available deep-learning infrastructure.

1.1. Reproducibility Advisory. All figures in this paper, including those in *SI Appendix*, are fully reproducible using the deposited code supplement (<https://purl.stanford.edu/zt044bg9296>). In the code supplement we also offer a Python implementation of the proposed LOCA algorithm.

2. Problem Settings

2.1. The Burst Measurement Strategy. Consider first, for simplicity, the case where the latent domain for our system of interest is a

path-connected domain in a Euclidean space $\mathcal{X} = \mathbb{R}^d$. We call \mathcal{X} the latent space. Observations of the system consist of samples captured by a measurement device given as a nonlinear function $f: \mathcal{X} \rightarrow \mathcal{Y}$, where \mathcal{Y} is the ambient, or measurement space. Even if f is invertible, it is generally not feasible to identify f^{-1} without access to \mathcal{X} . Here we assume that 1) f is smooth and injective and that 2) multiple, slightly perturbed versions of the physical system point in \mathcal{X} give rise to multiple (slightly perturbed) measurements in \mathcal{Y} . In this notation, by exploiting a specific type of local perturbation, we develop a method to recover a standardized version \mathcal{X} from \mathcal{Y} (up to an approximately isometric transformation, which, for Euclidean spaces, would be a rigid transformation).

Consider N data points (burst centers), denoted $x_1, \dots, x_N \in \mathbb{R}^d$ in the latent space. Assume that all these points lie on a path-connected, d -dimensional subdomain of \mathcal{X} ; we will later discuss the restriction to Riemannian manifolds with a smaller dimension than the full space. Importantly, we do not have direct access to the latent space \mathcal{X} . Samples in the latent space \mathcal{X} , which can be thought of as latent states, are pushed forward to the ambient space \mathcal{Y} via the unknown deformation f . Let $y_i = f(x_i)$ (for $i = 1, \dots, N$) so that $y_1, \dots, y_N \in \mathbb{R}^D$. We assume that the observed burst around y_i consists of perturbed versions of the latent state x_i , pushed through the unknown deformation f . Formally, for fixed $1 \leq i \leq N$, let $y_i^{(1)}, \dots, y_i^{(M)}$ be independent and identically distributed samples of the random variable

$$Y_i = f(X_i) \in \mathbb{R}^D, \quad [1]$$

where $X_i \sim \mathcal{N}_d(x_i, \sigma^2 I_d)$ for $i = 1, \dots, N$ are independent random variables. The data available to the scientist consists of N sets of observed states, where each set, indexed by $1 \leq i \leq N$, is $\{y_i^{(j)}\}_{j=1, \dots, M}$. We assume that $\sigma \ll 1$ or, alternatively, that σ is sufficiently small such that the differential of f practically does not change within a ball of radius σ around any point. Such sufficiently small σ allows us to capture the local neighborhoods of the states at this measurement scale on the latent manifold. While we assume Gaussian perturbations, note that other isotropic distributions that satisfy this condition could also be used to model the perturbations.

Let us explore the implications of this localized sampling strategy for learning a representation that is consistent with \mathcal{X} . Specifically, our goal is to construct an embedding $\rho: \mathbb{R}^D \rightarrow \mathbb{R}^s$ that maps the observations y_i , so that the image of $\rho \circ f$ is isometric to \mathcal{X} when σ is known. In our Euclidean setting, such an isometric embedding should satisfy

$$\|\rho(y_i) - \rho(y_j)\|_2 = \|x_i - x_j\|_2, \text{ for any } i, j = 1, \dots, N. \quad [2]$$

We note that if σ is not known, we will relax Eq. 2 by allowing a global scaling of the embedding. In general, Eq. 2 means that we are only looking for a representation that preserves the pairwise Euclidean distances between the latent samples, rather than obtaining their actual values. More specifically, a ρ that satisfies Eq. 2 is not unique and is defined up to an isometric transformation of the data. We refer to representations which satisfy Eq. 2 up to errors smaller than σ as “isometries.”

3. Related Work

The problem of finding an isometric embedding was also studied in ref. 26. The paper proposes an algorithm to embed a manifold with dimension d into a space of dimension $s \geq d$. The method in ref. 26 is based on ref. 12 and uses a discrete version of the Laplace–Beltrami operator, as in ref. 7, to estimate the metric of the desired embedding. To ensure that the embedding is isometric to the observed samples \mathcal{Y} , the authors proposed a loss function that measures the deviation between the push-forward metric on the observed space \mathcal{Y} and the restricted Euclidean

metric induced by the embedding. The embedding then is refined by gradient descent optimization with respect to this proposed loss. The approach successively approximates a Nash embedding with respect to the observed space, \mathcal{Y} , and requires that the manifold is densely sampled at all scales.

In this work, we use bursts to learn an embedding that corrects the deformation f and isometrically represents the inaccessible manifold \mathcal{X} . The idea of using bursts, or data neighborhoods, to learn a standardized reparametrization of data was first suggested in ref. 13. The authors assume the data are obtained via some unknown nonlinear transformation of latent independent variables. Locally, the distortion caused by the transformation is corrected by inverting a Jacobian that is estimated from the covariances of the observed bursts. This allows the authors to define a local Mahalanobis metric (which is affine invariant). Then, this metric is used to construct an anisotropic intrinsic Laplacian operator. Finally, the eigenvectors of this Laplacian provide the independent components and are used as a canonical embedding. This framework was extended in several studies such as refs. 15, 22, 23, and 25.

The work of these authors can be improved in three directions. First, they require inverting a covariance matrix in the ambient space. Second, they suffer from deformation near the boundaries. Third, they typically do not provide an embedding function that can be naturally extended over the entire data domain and beyond; instead, they provide a specific mapping for the existing training samples. This last direction means that to embed test data, methods such as refs. 14 and 27 can be employed. The mapping approximations based on these methods are limited and cannot extend further than a small neighborhood around each training point. Furthermore, even though the provided embedding is unique, it is not isometric to the latent variables. We present a method that alleviates these shortcomings and empirically demonstrate that it extracts a canonical representation that is isometric to the latent variables.

The work that is perhaps most related to this study was recently presented by ref. 19. The authors consider using bursts to develop a method for finding an embedding that is isometric to the latent variables. They build upon Isomap (28) and use two neural networks to refine the Isomap-based embedding. The first neural network is used in order to obtain a continuous model for estimating the covariance $C(Y_i)$. The covariances are used for calculating local Mahalanobis-based distances, which are fed into Isomap to obtain an initial embedding. Next, they train an additional neural network to correct the Isomap embedding so that the Euclidean distances will approximate the local Mahalanobis distances. In this paper, we take a different and, we believe, a more systematic/general approach by presenting a simple encoder–decoder pair (Fig. 2) that is directly applicable to samples in the observed, high-dimensional space. Specifically, our approach provides a parametric mapping that allows us to extend the embedding to new unseen samples naturally. Furthermore, we learn the inverse mapping, which could be used to generate new samples by interpolating in the latent space.

4. Deriving an Alternative Isometry Objective

Without access to samples from \mathcal{X} , the objective described in Eq. 2 does not provide any information for extracting ρ . Here we reformulate this objective by utilizing the special stochastic sampling scheme presented in Section 2 and relate it to the differential equation for the embedding described in Lemma 1. We start by plugging the unknown measurement function into Eq. 2; then, we can approximate its left-hand side using a first-order Taylor expansion

$$\begin{aligned} \|\rho(y_i) - \rho(y_j)\|_2 &= \|\rho \circ f(x_i) - \rho \circ f(x_j)\|_2 \\ &\approx \|\mathbf{J}_{\rho \circ f}(x_i)(x_j - x_i)\|_2. \end{aligned}$$

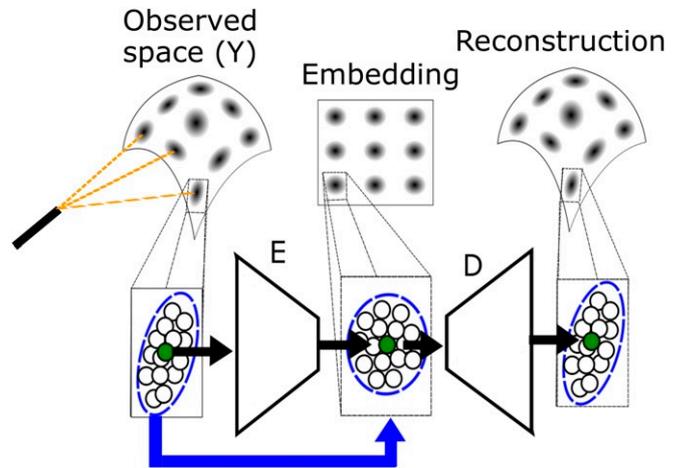


Fig. 2. An illustration of the LOCA. The observation space (\mathcal{Y}) is assumed to model a nonlinear deformation of the inaccessible manifold (\mathcal{X}). We attempt to invert the unknown measurement function. Here we utilize the bursts sampling strategy illustrated in Fig. 1. LOCA is a special type of autoencoder, consisting of an encoder (E) parametrized by ρ and a decoder (D) parametrized by γ (see Section 5). The autoencoder receives a set of points along with corresponding neighborhoods; each neighborhood is depicted as a dark oval point cloud (at the top of the figure). At the bottom, we zoom in onto a single anchor point y_i (green) along with its corresponding neighborhood Y_i (bounded by a blue ellipsoid). The encoder attempts to whiten each neighborhood in the embedding space, while the decoder tries to reconstruct the input.

Hence, by neglecting higher-order terms, we can define the following objective

$$\mathbf{J}_{\rho \circ f}(x_i)^T \mathbf{J}_{\rho \circ f}(x_i) = \mathbf{I}_d, \text{ for } i = 1, \dots, N, \quad [3]$$

which allows us to evaluate the isometric property of ρ .

Now we want to relate the Jacobian in Eq. 3 to the measurable properties of the observations Y_1, \dots, Y_N , pushed forward using the encoder, ρ . Specifically, we can rely on the following lemma to approximate the derivatives of the unknown function f at each point x_1, \dots, x_N . Lemma 1 is proved in SI Appendix, section 1:

Lemma 1. Let $d, s \in \mathbb{N}$, where $s \geq d$. Let $g: \mathcal{X} \rightarrow \mathcal{Z}$ be a function, where $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Z} = \mathbb{R}^s$. Let $x \in \mathcal{X}$ and $\sigma \in \mathbb{R}_+$. Define a random variable $X \sim \mathcal{N}(x, \sigma^2 \mathbf{I}_d)$. If the function satisfies $g \in C^3$ and is injective, there exists a $\sigma \in \mathbb{R}_+$ such that the covariance of the transformed random variable $Z = g(X)$ is related to the Jacobian of g at x via

$$\mathbf{J}_g(x) \mathbf{J}_g(x)^T = \frac{1}{\sigma^2} \mathbf{C}(Z) + O(\sigma^2).$$

Moreover,

$$\frac{1}{\sigma^2} \mathbf{C}(Z) \xrightarrow{\sigma \rightarrow 0} \mathbf{J}_g(x_i) \mathbf{J}_g(x)^T.$$

By setting $g \equiv \rho \circ f$, Lemma 1 provides a relation between the Jacobian of $\rho \circ f$ and the covariance in the embedding space. Specifically, this translates to a system of differential equations for the Jacobian of an isometric (Nash) embedding

$$\mathbf{J}_{\rho \circ f}(x_i) \mathbf{J}_{\rho \circ f}(x_i)^T = \frac{1}{\sigma^2} \mathbf{C}(\rho(Y_i)) + O(\sigma^2). \quad [4]$$

When $s = d$, meaning ρ embeds the data in \mathbb{R}^d , we can link the approximation of objective Eq. 2 to Eq. 4 by

$$\frac{1}{\sigma^2} \mathbf{C}(\rho(Y_i)) = \mathbf{I}, \text{ for any } i = 1, \dots, N. \quad [5]$$

Thus, we can evaluate the embedding function at each point without gaining access to the latent states of the system.

Algorithm 1: LOCA: local conformal autoencoder**Input:** Observed clouds $Y_i, i = 1, \dots, N$.**Output:** θ_e and θ_d - the weights of the encoder ρ and decoder γ neural networks.1: **for** $t = 1, \dots, T$ **do**

2: Compute the whitening loss

$$L_{white} = \frac{1}{N} \sum_{i=1}^N \left\| \frac{1}{\sigma^2} \widehat{C}(\rho(Y_i)) - I_d \right\|_F^2$$

3: Update $\theta_e := \theta_e - \eta \nabla_{\theta_e} L_{white}$

4: Compute the reconstruction loss

$$L_{recon} = \frac{1}{N \cdot M} \sum_{i,m=1}^{N,M} \left\| y_i^{(m)} - \gamma(\rho(y_i^{(m)})) \right\|_2^2$$

5: Update $\theta_e := \theta_e - \eta \nabla_{\theta_e} L_{recon}$ and
 $\theta_d := \theta_d - \eta \nabla_{\theta_d} L_{recon}$ **5. Local Conformal Autoencoder**

We now introduce the LOCA, with training *Algorithm 1*. Our method is based on optimizing two loss terms: the first is defined based on Eq. 5, using what we refer to as a whitening loss,

$$L_{white}(\rho) = \frac{1}{N} \sum_{i=1}^N \left\| \frac{1}{\sigma^2} \widehat{C}(\rho(Y_i)) - I_d \right\|_F^2, \quad [6]$$

where ρ is an embedding function and $\widehat{C}(\rho(Y_i))$ is the empirical covariance over a set of M realizations $\rho(y_i^{(1)}), \dots, \rho(y_i^{(M)})$,

where $y_i^{(1)}, \dots, y_i^{(M)}$ are realizations of the random variable Y_i .

As f is invertible on its domain, an embedding function ρ that approximates f^{-1} should be invertible as well. The invertibility of ρ means that there exists an inverse mapping $\gamma: \mathbb{R}^d \rightarrow \mathcal{Y}$, such that $y_i = \gamma(\rho(y_i))$ for any $i \in [N]$. This additional objective helps remove undesired ambiguities (which may occur for insufficient sampling). By imposing an invertibility property on ρ , we effectively regularize the solution of ρ away from noninvertible functions. To impose invertibility, we define our second loss term, referred to as reconstruction loss:

$$L_{recon}(\rho, \gamma) = \frac{1}{N \cdot M} \sum_{i,m} \left\| y_i^{(m)} - \gamma(\rho(y_i^{(m)})) \right\|_2^2. \quad [7]$$

We suggest finding an isometric embedding based on an autoencoder, where ρ will be defined as the encoder and γ as the decoder. We construct solutions to Eqs. 6 and 7 with a neural network ansatz $\rho = h_e^{(L)}$ and $\gamma = h_d^{(L)}$ consisting of L layers each, such that

$$h_e^{(\ell)}(y) = \sigma_e \left(W_e^{(\ell-1)} h_e^{(\ell-1)}(y) + b^{(\ell-1)} \right), \quad \ell = 1, \dots, L,$$

$$h_d^{(\ell)}(z) = \sigma_d \left(W_d^{(\ell-1)} h_d^{(\ell-1)}(z) + b_d^{(\ell-1)} \right), \quad \ell = 1, \dots, L,$$

where $h_e^{(0)}(y) = y$ and $h_d^{(0)}(z) = z$. Here W_e^ℓ, b_e^ℓ and W_d^ℓ, b_d^ℓ are the weights and biases at layer ℓ of the encoder and decoder, respectively. The functions σ_e, σ_d are nonlinear activations applied individually to each input coordinate. As the activation function can have a limited image, we recommend removing the nonlinear activation for $\ell = L$.

We propose to find ρ and γ by alternating between a stochastic gradient descent on Eqs. 6 and 7. It is important to note that the main objective that we are trying to optimize for is based on Eq. 6; therefore, Eq. 7 can be viewed as a regularization term. A

pseudo-code of this procedure appears in *Algorithm 1*. To prevent overfitting, we propose an early stopping procedure (29, 30) by evaluating the loss terms on a validation set. In *Section 6*, we demonstrate different properties of the proposed LOCA algorithm using various geometric example manifolds.

Note that functions that perfectly satisfy our objectives are not unique; i.e., for any solution ρ we can define an equivalent solution $\bar{\rho}$ that will attain the same loss. Specifically, we can define it by $\bar{\rho}(y) = U y + c$ for any $y \in \mathcal{Y}$, where $U \in O(d)$ and $c \in \mathbb{R}^d$.

In summary, 1) we collect distorted neighborhoods of a fixed size (σ) around data points of the system of interest. 2) We embed/encode the data in a low-dimensional Euclidean space, so that these neighborhoods are standardized or z-scored. 3) The embedding is decoded back to the original measurements, to regularize the encoder. In *Section 6*, we demonstrate that 4) the encoder is invariant to the measurement modality [up to errors of $O(\sigma^2)$ and modulo an orthogonal transformation and shift]. Below, we further demonstrate that 5) the parametric form of the embedding enables reliable interpolation and extrapolation.

6. Properties of LOCA

In this section, we evaluate the properties of the proposed embedding ρ by generating various synthetic datasets. We compare the extracted embedding provided by LOCA (described in *Section 4*) with the embeddings of alternative methods such as diffusion maps (DM) (7) and anisotropic diffusion maps (A-DM) (13) denoted as ψ and ϕ , respectively. Note that the DM algorithm does not use the burst data, while the A-DM algorithm uses it in order to construct a local Mahalanobis metric. We present the details of the exact implementation for each of the methods in *SI Appendix, section 2*.

6.1. LOCA Creates an Isometric Embedding. We first evaluate the isometric quality of the proposed embedding ρ with respect to the true inaccessible structure of \mathcal{X} . Here we follow the setting in ref. 13, where the intrinsic latent coordinates are independent, specifically distributed by $U[0, 1]^2$. Based on this distribution, we sample $N = 2,000$ anchor points x_i along with bursts X_i . Each burst consists of $M = 200$ points sampled independently from $\mathcal{N}_2(x_i, \sigma^2 I_2)$, where $\sigma = 0.01$.

We now define the nonlinear transformation, $f_1: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ (as in ref. 13), to the ambient space by

$$f_1(x) = \begin{pmatrix} x[0] + x[1]^3 \\ -x[0]^3 + x[1] \end{pmatrix}, \quad [8]$$

for any $x \equiv (x[0], x[1])^\top \in \mathbb{R}^2$. In Fig. 3, we present measurements from \mathcal{X} with the corresponding measurements of \mathcal{Y} , where $\mathcal{Y} = f(\mathcal{X})$. To illustrate the local deformation caused by f_1 , we overlay the samples with clouds around five different positions (see green dots). Next, we apply LOCA (described in *Algorithm 1*) to compute an embedding ρ that satisfies Eqs. 6 and 7. We evaluate the isometric quality of LOCA by comparing the pairwise Euclidean distances in the embedding space ρ to the Euclidean distances in the latent space \mathcal{X} . For comparison, we apply DM and A-DM (which also uses the bursts) to \mathcal{Y} and plot the pairwise Euclidean distances in the embedding vs. the corresponding Euclidean distances in the latent space. Here we evaluate isometry up to a scaling, as DM and A-DM use eigenvectors (that are typically normalized). The scaling is optimized to minimize the stress defined by

$$\text{Stress}(\mathbf{g}) = \frac{1}{N(N-1)} \sum_{i,j=1}^N (D_x(x_i, x_j) - D_g(y_i, y_j))^2, \quad [9]$$

where \mathbf{g} is some embedding function from \mathcal{Y} and $D_g(y_i, y_j) = \|\mathbf{g}(y_i) - \mathbf{g}(y_j)\|_2$. Specifically, the stress values for LOCA and the

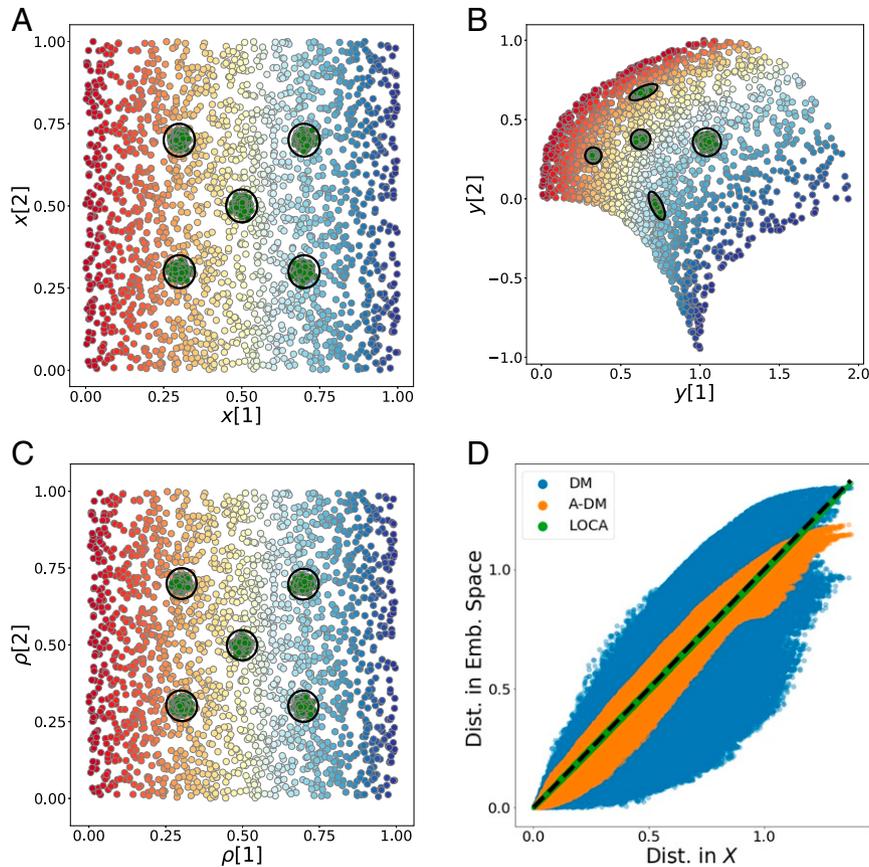


Fig. 3. Evaluating the isometric quality of the embedding (the setting is detailed in A). (A) The points $x_i, i = 1, \dots, N = 2,000$, in the latent space of the system of interest. (B) The pushed-forward points as measurements $y_i = 1, \dots, N$, created by applying the nonlinear transformation f_1 (described in Eq. 8). We observe bursts around each sample (based on the burst model described in Section 2). To illustrate this burst sampling scheme, we overlay the points with additional green samples generated by bursts at five different positions. (C) The calibrated embedding ρ (using an orthogonal transformation and a shift) overlaid with the corresponding green bursts. Here we calibrated the embedding merely for visualization purposes. The colors in these figures correspond to the values of $x[1]$. (D) Euclidean distances between pairs of points in the latent space plotted versus the corresponding Euclidean distance in the embedding space. The corresponding distances for DM and for A-DM are also shown in color, scaled with a factor which minimizes the stress (defined in Eq. 9).

scaled versions of DM and A-DM are $1.5 \cdot 10^{-5}$, 0.03, and 0.002, respectively. As evident from the stress values and from Fig. 3, LOCA provides an embedding that is approximately isometric to \mathcal{X} (up to an orthogonal transformation and shift).

6.2. The Encoder Is Observed to Extend Reliably to Unseen Samples.

In the next experiment, we evaluate the out-of-sample extension capabilities of LOCA. The experiment is based on the same nonlinear transformation described in Section 6.1. We sample $N = 2,000$ points from a partial region of the latent representation \mathcal{X} , specifically described by $[0, 1]^2 \setminus [0.1, 0.9]^2$. In Fig. 4, we present the framed sampling regions along with the corresponding observed framed regions in \mathcal{Y} (see black and green frames in Fig. 4 A and B). To generate the bursts, we follow the settings presented in Section 6.1 and refer to them as our training set. The test set is defined by an additional $2 \cdot 10^4$ samples generated as in Section 6.1 from $[-.025, 1.025]^2$ in \mathcal{X} pushed forward by f_1 .

In Fig. 4C, we quantify the interpolation and extrapolation capabilities of LOCA by presenting the extracted embedding along with the corrected frame. To further evaluate the quality of this embedding, we compare the pairwise distances in Fig. 4D (as described in Section 6.1). This comparison (presented in Fig. 4D) demonstrates the merit of using LOCA for extrapolating the embedding to unseen samples. Interestingly, the actual stress val-

ues of LOCA in the interpolation region, on the frame and in the extrapolation region, are all approximately 10^{-4} .

6.3. The Decoder Is Observed to Extend Reliably to Unseen Samples.

In this experiment we evaluate the out-of-sample capabilities of LOCA's decoder. While in Section 6.2 we trained LOCA and evaluated the quality of the encoder on unseen data, here we focus on the performance of the decoder. Specifically, we apply the decoder to unseen samples from the embedding space. Each unseen sample in the embedding space is created using linear interpolation. We now provide the exact details of this evaluation.

For this interpolation experiment, we use the same LOCA model trained in Section 6.2 on the framed data. We further generate $N = 400$ points in the interior boundary of the frame, represented by the green dots in Fig. 5. Next, we perform a linear interpolation between horizontal and vertical pairs; see, for example, the colored lines in Fig. 5A. The data are then pushed forward using the nonlinear transformation described in Eq. 8, as shown in Fig. 5B. We embed the training samples along with the green frame using LOCA; a calibrated version of the embedding space is shown in Fig. 5C (up to a shift and an orthogonal transformation). Then, we perform an additional interpolation in the embedding space using the same corresponding pairs as were used in the latent space (Fig. 5A). Finally,

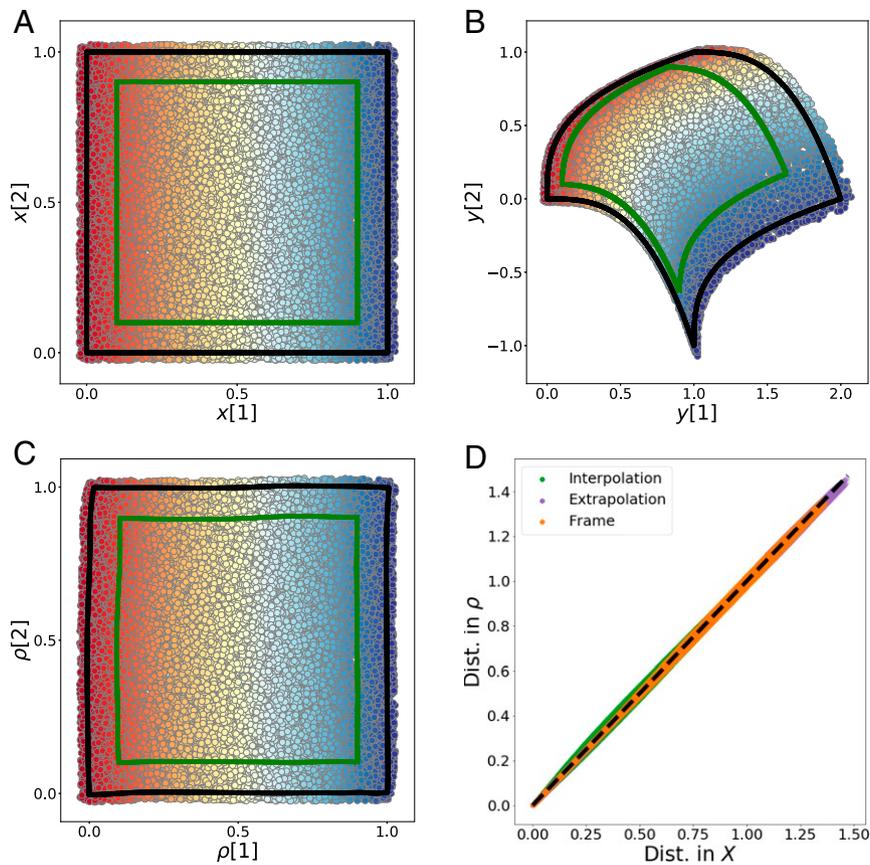


Fig. 4. Evaluating the out-of-sample performance of the encoder, detailed in Section 6.2. (A) The latent space of interest X . Our training region is bounded here by the black and green frames. The interpolation region lies within the green frame, while the extrapolation regime lies outside the black frame. (B) The observed space Y with corresponding regions of interest. (C) The calibrated embedding ρ (using an orthogonal transformation and a shift) with corresponding regions of interest. The colors in these figures correspond to the values of $x[1]$. Here we calibrated the embedding merely for visualization purposes. (D) The Euclidean distances between pairs of points in the latent space versus the corresponding Euclidean distance in the embedding space.

we apply the decoder to the embedding of the training samples and to the newly interpolated samples, which are presented in Fig. 5D. As evident in this figure, the reconstructed points faithfully capture the mushroom-shaped manifold. The mean squared error between the pushed-forward interpolated points and the decoded interpolated points is $2.3 \cdot 10^{-4}$, with an SD of $2.4 \cdot 10^{-4}$. This experiment demonstrates that LOCA may also be used as a generative model, by reconstructing new points generated using interpolation in the embedding space.

6.4. LOCA on a Curved Manifold. Here we examine a more challenging configuration, generalizing our original Euclidean problem setting. The latent space is now taken to be a k -dimensional manifold that resides in \mathbb{R}^d , where $d > k$ and d is the minimal dimension required to embed the manifold in a Euclidean space isometrically. Interestingly, we consider an observation process such that the observation dimension, D , is smaller than d . To clarify, this means that the measurement process can involve projections to a lower dimension.

We consider a manifold that covers three quarters of a 2D unit sphere in \mathbb{R}^3 , where the training points admit the following form:

$$\mathbf{x} = \begin{pmatrix} \sin(\alpha) \cos(\beta) \\ \sin(\alpha) \sin(\beta) \\ \cos(\alpha) \end{pmatrix} \quad \beta \in [0, 2\pi), \alpha \in [\pi/3, \pi]. \quad [10]$$

The manifold is embedded in \mathbb{R}^3 but has an intrinsic dimension of 2. This requires us to revisit our definition of bursts, discussed

in Eq. 1. Specifically, we assume that the bursts are confined to the manifold. Here we approximate this constraint in the form of random variables Z_i , obtained using a local isotropic Gaussian with a 2D covariance $\sigma^2 \mathbf{I}_2$, defined on the tangent plane to the point.

We consider $N = 491$ states of the system \mathbf{x}_i , $i = 1, \dots, N$, which are generated on a uniform grid using the Fibonacci sphere sampling scheme (31) for points with $\alpha \in [\pi/3, 5\pi/6]$. We define each burst X_i using $M = 400$ points sampled from our 2D isotropic Gaussian defined by the tangent plane around \mathbf{x}_i with $\sigma = 0.01$. Now, in order to create the observed samples \mathbf{y} , we apply the stereographic projection to \mathbf{x} by projecting each point from \mathcal{X} onto a 2D space defined by

$$\mathbf{y} = \begin{pmatrix} \frac{x[1]}{1-x[3]} \\ \frac{x[2]}{1-x[3]} \end{pmatrix}. \quad [11]$$

The transformation can be thought of as a projection onto the plane $\mathbb{R}^2 \times \{1\}$; an illustration of the stereographic projection appears in Fig. 6A. The training bursts in the latent space and the observed space appear in Fig. 6B and C, respectively.

We apply DM, A-DM, and LOCA to embed the data in a 3D space. The difference between the pairwise Euclidean distances (see description in Section 6.1) in each embedding space and the original Euclidean distances along with the extracted embeddings are described in SI Appendix. The stress values for LOCA and the scaled DM and A-DM on the training data are 10^{-3} , 0.18, and $6 \cdot 10^{-3}$, respectively. In order to examine

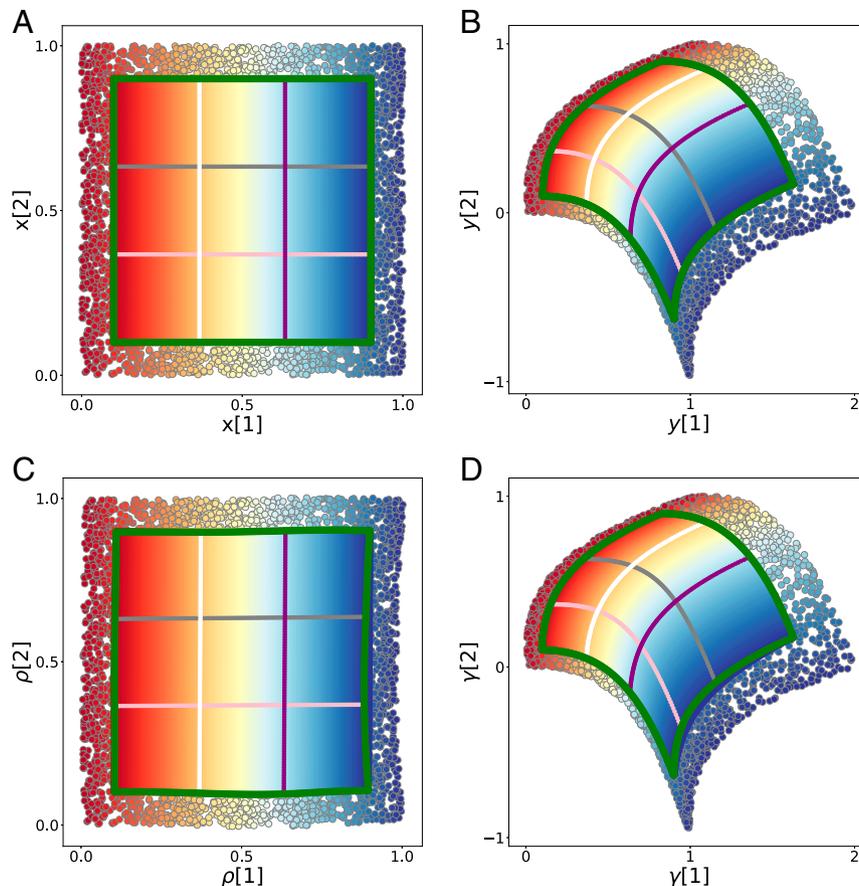


Fig. 5. Evaluating the out-of-sample reconstruction capabilities of LOCA. Here we attempt to generate new points in the ambient space by performing a linear interpolation in the embedding space. A description of the linear interpolation appears in C. (A) The inaccessible latent space. The points surrounding the green frame are the training samples. Interpolation is performed horizontally and vertically between points on the green frame; see, for example, the four colored lines. (B) The pushed-forward data from A based on the nonlinear function f_1 described in Eq. 8. (C) Recovered calibrated embedding ρ . The training samples in the frame and green border are embedded using LOCA. Within the embedded green border we perform an additional linear interpolation, using the same corresponding pairs as were used in the latent space. For example, see the four horizontal and vertical colored lines. (D) The pushed forward data from C by the decoder (γ) learned by LOCA. This experiment demonstrates that LOCA learns a decoding function that is consistent with the unknown transformation, even in a regime that is not covered by training samples.

the interpolation capabilities of LOCA, we generate 55 points using the Fibonacci sphere that satisfies $\alpha \in (5\pi/6, \pi]$. Using the trained model of LOCA, we embed these data and obtain that the stress value is 10^{-4} . Fig. 6 demonstrates that LOCA can well approximate an isometry, even if the dimension of the observations is lower than the minimal embedding dimension needed for the isometry, i.e., $s > D$.

7. Applications

7.1. Flattening a Curved Surface. Our first application is motivated by ref. 32, in which the authors propose a method for estimating the 3D deformation of a 2D object. They focus on the task of autonomous robotic manipulation of deformable objects. Their method uses a stream of images from an RGB-D camera and aligns them to a reference shape to estimate the deformation.

We explore the merits of LOCA for the task of estimating a deformation based on a 2D projection of an object, without using any depth information. A black square-shaped uniform grid has been created with $N = 2,500$ burst centers. Around each center, we generated a burst with $M = 50$ samples drawn from a Gaussian with $\sigma = 0.01$. The collection of 2,500 bursts has been printed on an A-4 white page (Fig. 7A). We manually deformed the printed square and photographed the deformed page from above. The image of the original squared object along with a

2D snapshot of the deformed object appears in Fig. 7B. This experiment complements the motivating example presented in Fig. 1.

To define the anchor points y_i along with corresponding bursts, we first identify the locations of all points by applying a simple threshold filter to the image. Then, we identify the bursts by applying density-based spatial clustering of applications with noise (33). In Fig. 7B we present the identified groups of points (black). Note that some bursts are lost in this process as there is nearly no gap between them in the deformed shape. Here the parameter σ^2 for the whitening loss (Eq. 6) is estimated using the median of the first eigenvalue of the bursts covariances. We apply LOCA and extract the embedding ρ . In Fig. 7C, we present a calibrated version (scaled rigid transformation) of the embedding, $\tilde{\rho}$, overlaid on the latent representation. The transformation is found by minimizing the mean squared error between the underlying representation and the extracted embedding of the four corners of the square. This experiment demonstrates that LOCA corrects the unknown deformation based on the estimated bursts.

7.2. Application to Wi-Fi Localization. Here we evaluate LOCA for the task of geographical localization based on mobile devices. The localization problem involves estimating the geographic

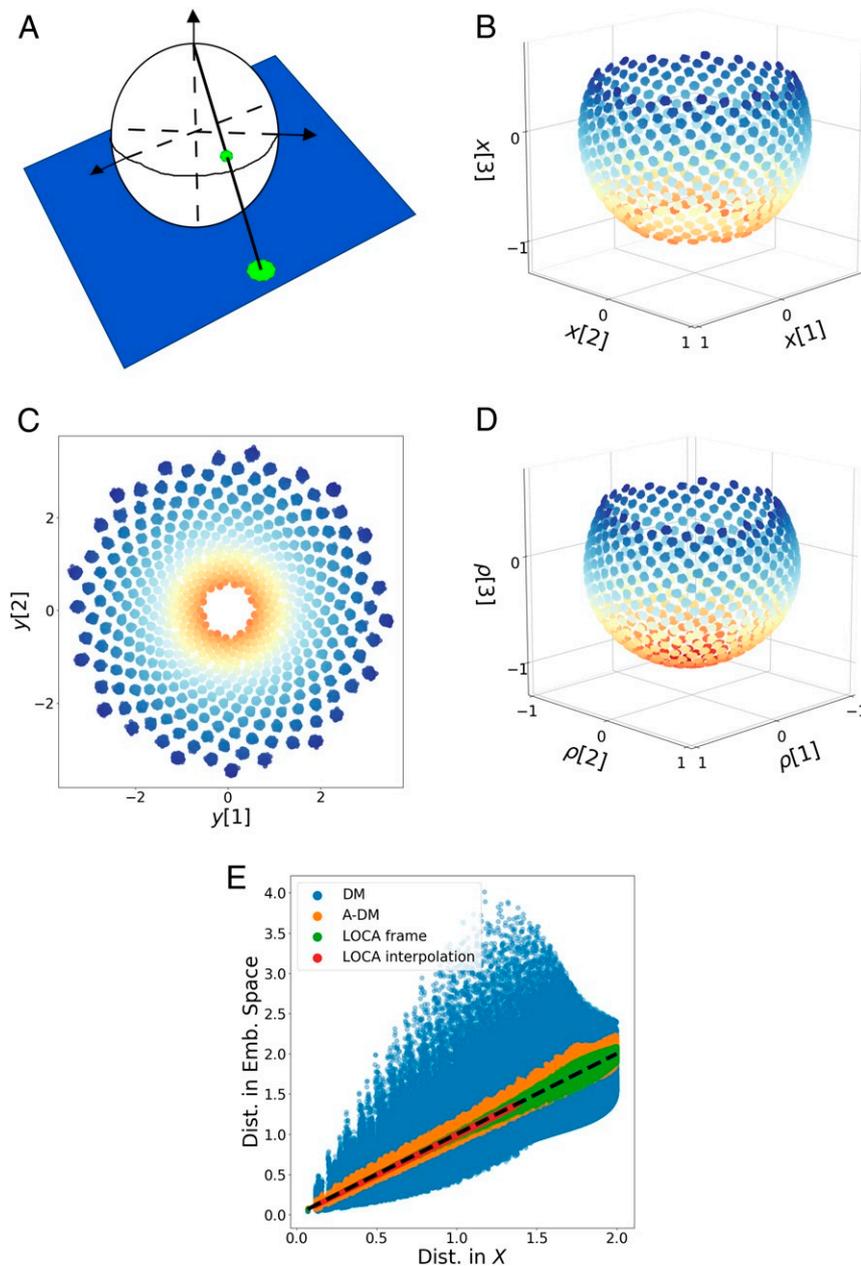


Fig. 6. The stereographic projection experiment (see description in Section 6.4). (A) A schematic illustration of the stereographic projection generating the data. (B) The original latent representation of the bursts employed. (C) The 2D observations of the bursts created using the stereographic projection. The plot contains only the training data, meaning points that satisfy $\alpha \in (\pi/3, 5\pi/6)$, leaving a hole at the south pole. (D) The 3D embedding of these training data, with the missing lower cap ($\alpha \in (\pi/3, \pi]$). The color represents the value of α of each point as defined in Eq. 10. The colors used in B–D correspond to the spherical angle α defined in Eq. 10. (E) The Euclidean distances between pairs of points in the original, 3D latent space versus the corresponding Euclidean distance in the embedding space. Here we compare distances based on the training region (frame) as well as the unseen test region, where $\alpha \in (5\pi/6, \pi]$ (interpolation).

location of a receiver device based on signals sent from multiple Wi-Fi transmitters. This problem has been addressed by modeling the strength of the signal in time and space or based on fingerprints of the signal learned in a supervised setting (34, 35). We address the problem in an unsupervised fashion by applying the proposed LOCA algorithm without employing any physical model.

The experiment is performed by simulating the signal strength of $L = 17$ Wi-Fi transmitters at multiple locations across a model of a room, where each transmitter uses a separate channel. The room is modeled based on a simplified floor plan of the

fourth floor of the unusually shaped Stata Center building at Massachusetts Institute of Technology. We refer to the 2D representation of the room as $\mathcal{X} \subset \mathbb{R}^2$; a schematic of the floor plan with $600 \times 1,000$ pixels appears in Fig. 8 (black line). The $L = 17$ Wi-Fi transmitters are randomly located across the floor plan; we denote each of these locations by $t_\ell \in \mathbb{R}^2$, for any $\ell \in \{1, \dots, 17\}$. Next, we sample $x_i, i = 1, \dots, N$, using $N = 4,000$ anchor points distributed uniformly over \mathcal{X} , and define the amplitude of each measured Wi-Fi signal using a radial basis function (RBF). The RBF decay is monotonic in the distance between the transmitter and the measurement location, so that the amplitude at point x_i

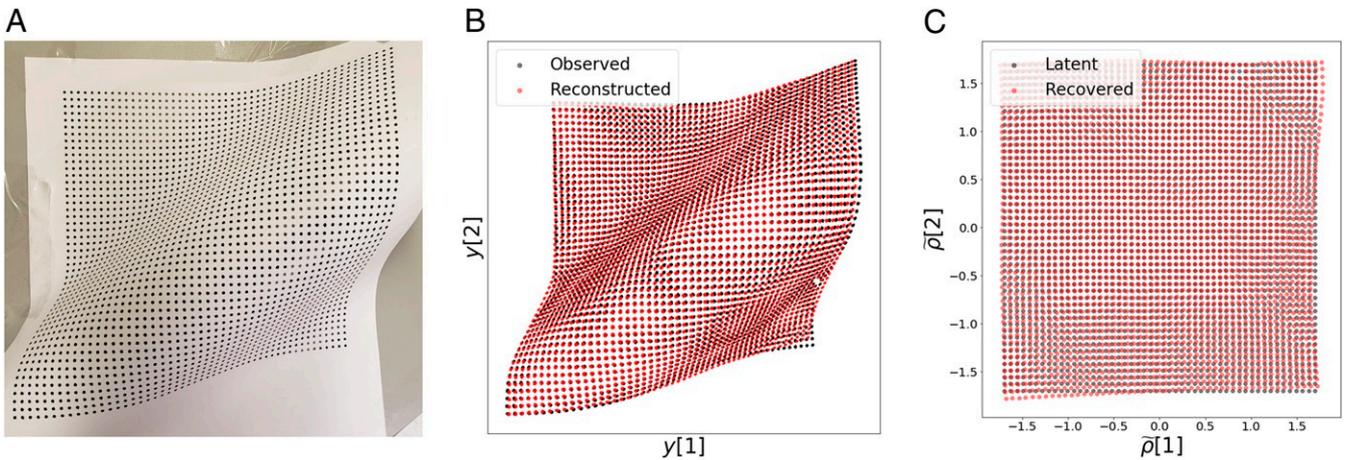


Fig. 7. A LOCA embedding can flatten a deformed object using estimated bursts. (A) An image of the deformed printed square on an A-4 paper. (B) The input training samples used by LOCA (black) that were extracted from A, along with the reconstructed points (red). (C) The calibrated embedding (using an orthogonal transformation and a shift) of the deformed object using LOCA (red) and the underlying representation of points (black). As in the synthetic examples, we use calibration only for visualization purposes. Here LOCA manages to correct the deformation of the local bursts and thus learns a function that approximately uncovers the latent structure of the object.

of the signal of transmitter ℓ is $y_{i,\ell} = \exp(-\|\mathbf{x}_i - \mathbf{t}_\ell\|_2^2 / \epsilon^2)$, where $\epsilon = 600$ pixels. Here the bursts will be defined by a circle of $M = 6$ receivers equally spaced at a radius of $r = 0.5$ pixels around each anchor point \mathbf{x}_i : these six receivers model a circular sensor array as the measurement device.

Next, we apply LOCA and embed the observed vectors of multichannel amplitudes into a 2D space. To demonstrate the performance of LOCA, we calibrate the LOCA embedding to the ground truth floor plan using a shift and scaled orthogonal transformation, as done in *Section 7.1* but using all of the training data. In Fig. 8 we present the scaled, calibrated, 2D embedding $\hat{\rho}$ with the locations of the transmitters and anchor points. The isometric properties of LOCA can be evaluated based on the stress value between distances in the calibrated LOCA embedding and distances in the original floor plan. Here the stress value is 0.33.

8. Discussion

In this paper, we proposed a method and an algorithm for extracting canonical data coordinates from scientific measurements. Our algorithm produces a nonlinear embedding that is approximately isometric to the unknown latent manifold structure of the data. Our method assumes a specific, broadly applicable stochastic sampling strategy and successfully corrects for unknown measurement device deformations. The proposed method constructs a representation that whitens (namely, changes to multivariate z scores) batches of neighboring measurements, which we call bursts. We impose additional constraints to patch together the locally whitened neighborhoods, ensuring a smooth global structure. Finally, the method is implemented using a neural network architecture, namely, an encoder–decoder pair, which we name LOCA. As shown in *Lemma 1*, the covariances of the bursts can be used to estimate the Jacobian of the unknown measurement function. Alternatively, we can replace this estimation with any other type of measurement strategy informative enough to estimate the local Jacobian of the measurement function.

The method can be summarized as follows: 1) we collect distorted neighborhoods of a fixed size of data samples; 2) we embed/encode the data in the lowest dimensional Euclidean space so that these neighborhoods are standardized or z-scored; and 3) the data are decoded from the embedding space to original measurements, enabling interpolation and extrapolation.

We observed that LOCA is invariant to the measurement modality (approximately to second order and modulo a rigid transformation). Under condition of scaling consistency for samples drawn from a Riemannian manifold, the encoder approximates an isometric embedding of the manifold. From an implementation perspective, our method is simpler than existing manifold learning methods, which typically require an eigendecomposition of an N -by- N matrix (N being the number of samples). Indeed, existing implementations of deep neural networks enable a single developer to produce fast, reliable, GPU-based implementations of LOCA.

We provided solid empirical evidence that if the deformation is invertible, then LOCA extracts an embedding that is isometric to the latent variables. Moreover, LOCA exhibits intriguing—indeed promising—interpolation and extrapolation capabilities, as demonstrated in two promising applications. Finally, in *SI Appendix*, we demonstrate that LOCA can be used to register

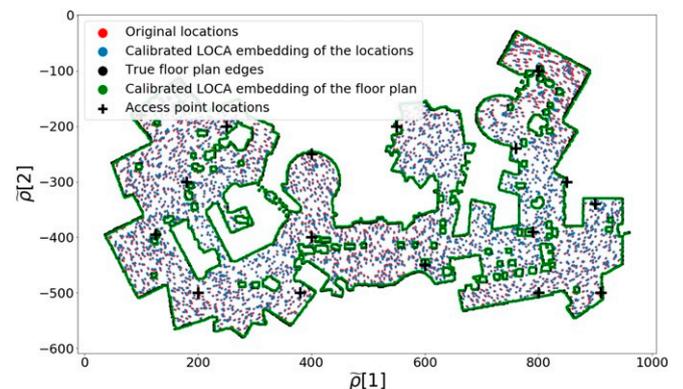


Fig. 8. Application of LOCA to Wi-Fi localization. We use a floor plan model based on the fourth floor of Massachusetts Institute of Technology’s Ray and Maria Stata Center. The edges of the ground truth model appear in black. We simulate $L = 17$ Wi-Fi access points (transmitters), which are presented as black crosses. We use $N = 4,000$ locations depicted as red dots with corresponding $M = 6$ burst samples around them (modeling a circular antenna array). To demonstrate that LOCA’s embedding is coherent with the latent representation, we calibrate the embedding to the true floor plan; see blue dots and green line.

observations from multiple high-dimensional modalities (*SI Appendix*, section 4 and Fig. S4).

Data Availability. The code and data supplement are available online at the Stanford Digital Repository (<https://purl.stanford.edu/zt044bg9296>).

1. S. Alam *et al.*, The eleventh and twelfth data releases of the Sloan Digital Sky Survey: Final data from SDSS-III. *Astrophys. J. Suppl.* **219**, 12 (2015).
2. V. F. Calderon, A. A. Berlind, Prediction of galaxy halo masses in SDSS DR7 via a machine learning approach. *Mon. Not. R. Astron. Soc.* **490**, 2367–2379 (2019).
3. H.-P. Deutsch, “Principal component analysis” in *Derivatives and Internal Models* (Springer, 2002), pp. 539–547.
4. J. B. Kruskal, Nonmetric multidimensional scaling: A numerical method. *Psychometrika* **29**, 115–129 (1964).
5. J. B. Tenenbaum, V. de Silva, J. C. Langford, A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319–2323 (2000).
6. M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.* **15**, 1373–1396 (2003).
7. R. R. Coifman, S. Lafon, Diffusion maps. *Appl. Comput. Harmon. Anal.* **21**, 5–30 (2006).
8. S. T. Roweis, L. K. Saul, Nonlinear dimensionality reduction by local linear embedding. *Science*, **290**, 2323–2326.
9. P. Baldi, “Autoencoders, unsupervised learning, and deep architectures” in *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, I. Guyon, G. Dror, V. Lemaire, G. Taylor, D. Silver, Eds. (JMLR, 2012), pp. 37–49.
10. L. van der Maaten, G. Hinton, Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
11. L. McInnes, J. Healy, N. Saul, L. Großberger, Umap: Uniform manifold approximation and projection. *J. Open Source Softw.* **3**, 861 (2018).
12. D. Perrault-Joncas, M. Meila, Non-linear dimensionality reduction: Riemannian metric estimation and the problem of geometric discovery. arXiv:1305.7255 (30 May 2013).
13. A. Singer, R. R. Coifman, Non-linear independent component analysis with diffusion maps. *Appl. Comput. Harmon. Anal.* **25**, 226–239 (2008).
14. N. Rabin, R. R. Coifman, “Heterogeneous datasets representation and learning using diffusion maps and Laplacian pyramids” in *Proceedings of the 2012 SIAM International Conference on Data Mining*, J. Ghosh, H. Liu, I. Davidson, C. Domeniconi, C. Kamath, Eds. (SIAM, 2012), pp. 189–199.
15. R. Talmon, I. Cohen, S. Gannot, “Supervised source localization using diffusion kernels” in *2011 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)* (IEEE, 2011), pp. 245–248.
16. D. Kushnir, A. Haddad, R. R. Coifman, Anisotropic diffusion on sub-manifolds with application to earth structure classification. *Appl. Comput. Harmon. Anal.* **32**, 280–294 (2012).
17. H.-T. Wu, R. Talmon, Y.-L. Lo, Assess sleep stage by modern signal processing techniques. *IEEE Trans. Biomed. Eng.* **62**, 1159–1168 (2014).
18. X. Wang, H. Mao, H. Hu, Z. Zhang, Crack localization in hydraulic turbine blades based on kernel independent component analysis and wavelet neural network. *Int. J. Comput. Intell. Syst.* **6**, 1116–1124 (2013).
19. A. Schwartz, R. Talmon, Intrinsic isometric manifold learning with application to localization. *SIAM J. Imag. Sci.* **12**, 1347–1391 (2019).
20. A. Singer, R. Erban, I. G. Kevrekidis, R. R. Coifman, Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 16090–16095 (2009).
21. S. A. Shevchik, B. Meylan, G. Violakis, K. Wasmer, 3D reconstruction of cracks propagation in mechanical workpieces analyzing non-stationary acoustic mixtures. *Mech. Syst. Signal Process.* **119**, 55–64 (2019).
22. T. Ronen, R. R. Coifman, Empirical intrinsic geometry for nonlinear modeling and time series filtering. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 12535–12540 (2013).
23. C. J. Silva, T. Ronen, N. Rabin, R. R. Coifman, I. G. Kevrekidis, Nonlinear intrinsic variables and state reconstruction in multiscale simulations. *J. Chem. Phys.* **139**, 184109 (2013).
24. F. P. Kemeth *et al.*, An emergent space for distributed data with hidden internal order through manifold learning. *IEEE Access* **6**, 77402–77413 (2018).
25. M. Salhov *et al.*, Multi-view kernel consensus for data analysis. *Appl. Comput. Harmon. Anal.* **49**, 208–228 (2020).
26. J. McQueen, M. Meila, D. Joncas, “Nearly isometric embedding by relaxation” in *Advances in Neural Information Processing Systems* (2016), pp. 2631–2639.
27. R. R. Coifman, S. Lafon, Geometric harmonics: A novel tool for multiscale out-of-sample extension of empirical functions. *Appl. Comput. Harmon. Anal.* **21**, 31–52 (2006).
28. J. B. Tenenbaum, V. D. Silva, J. C. Langford, A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319–2323 (2000).
29. M. Li, M. Soltanolkotabi, S. Oymak, “Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks” in *International Conference on Artificial Intelligence and Statistics* (JMLR, 2020), pp. 4313–4324.
30. B. Ronen, D. Jacobs, Y. Kasten, S. Kritchman, “The convergence rate of neural networks for learned functions of different frequencies” in *Advances in Neural Information Processing Systems* (2019), pp. 4761–4771.
31. R. Swinbank, R. J. Purser, Fibonacci grids: A novel approach to global modelling. *Q. J. R. Meteorol. Soc.* **132**, 1769–1793 (2006).
32. T. Han, X. Zhao, P. Sun, J. Pan, Robust shape estimation for 3D deformable object manipulation. *Commun. Inf. Syst.* **18**, 107–124 (2018).
33. E. Schubert, J. Sander, M. Ester, H. P. Kriegel, X. Xu, Dbscan revisited, revisited: Why and how you should (still) use dbscan. *ACM Trans. Database Syst.* **42**, 19 (2017).
34. A. Jaffe, M. Wax, Single-site localization via maximum discrimination multipath fingerprinting. *IEEE Trans. Signal Process.* **62**, 1718–1728 (2014).
35. N. Alikhani, S. Amiranloo, V. Moghtadaiee, S. A. Ghorashi, “Fast fingerprinting based indoor localization by wi-fi signals” in *2017 7th International Conference on Computer and Knowledge Engineering (ICCKE)* (IEEE, 2017), pp. 241–246.