
Product Manifold Learning

Sharon Zhang
Princeton University

Amit Moscovich
Princeton University

Amit Singer
Princeton University

Abstract

We consider problems of dimensionality reduction and learning data representations for continuous spaces with two or more independent degrees of freedom. Such problems occur, for example, when observing shapes with several components that move independently. Mathematically, if the parameter space of each continuous independent motion is a manifold, then their combination is known as a product manifold. In this paper, we present a new paradigm for nonlinear independent component analysis called manifold factorization. Our factorization algorithm is based on spectral graph methods for manifold learning and the separability of the Laplacian operator on product spaces. Recovering the factors of a manifold yields meaningful lower-dimensional representations and provides a new way to focus on particular aspects of the data space while ignoring others. We demonstrate the potential use of our method for an important and challenging problem in structural biology: mapping the motions of proteins and other large molecules using cryo-electron microscopy datasets.

1 Introduction

Consider a data-generating process F which maps vectors of latent (unobserved) variables to observations,

$$(\theta_1, \dots, \theta_m) \xrightarrow{F} \mathbf{x} \in \mathbb{R}^D. \quad (1)$$

Our focus is on the case where the latent variables are low-dimensional and the observations are high-dimensional vectors. As an illustrative example, consider an industrial articulated robot where the latent

variables $\theta_1, \dots, \theta_m$ correspond to the angles of its m rotary joints, and the observed vector \mathbf{x} is a set of measurements recorded by an external monitoring system. When F is deterministic (i.e. without noise), full-rank, smooth and injective, then the space of observations is a submanifold of \mathbb{R}^D (Lee, 2012, Theorem 4.14). This observation underlies the field of *manifold learning* (Tenenbaum et al., 2000; Belkin and Niyogi, 2003; Coifman and Lafon, 2006; Talmon et al., 2013). A key problem in this field is dimensionality reduction, where a set of high-dimensional observations $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^D$ are mapped to a low-dimensional space, preferably one whose dimension is not much larger than the dimension of the latent space.

In this work, we explore the idea of manifold learning when the latent space is a *product manifold*. If each latent variable θ_i lies on a manifold \mathcal{M}_i of dimension d_i then the latent space is the cartesian product

$$\mathcal{M} = \mathcal{M}_1 \times \dots \times \mathcal{M}_m \quad (2)$$

which is a manifold of dimension $d = d_1 + \dots + d_m$. We refer to the decomposition (2) as a *manifold factorization* of \mathcal{M} and to each of the manifolds \mathcal{M}_i as *factors*. Learning in the latent space is achieved by a data-dependent manifold factorization algorithm whose input is the output of spectral graph-based methods for manifold learning (Belkin and Niyogi, 2003; Coifman and Lafon, 2006). These algorithms output a set of N Laplacian eigenvectors which approximate the Laplacean eigenfunctions of the manifold. On a product manifold the eigenfunctions are nothing but products of eigenfunctions on the factor manifolds. Hence, we can try to find the eigenvectors of the factor manifolds by seeking two small sets of vectors whose cartesian elementwise products produce all of the N eigenvectors.

Instead of having each observation map to a point on a high-dimensional manifold, our approach enables each observation to be mapped to several points on separate low-dimensional spaces. Since each factor may encode a different aspect of the data, this opens the door to exciting new methods for data visualization and representation. For example, one can use it to visualize observations which come from a high-dimensional latent

spaces as a set of points on m low-dimensional manifolds that correspond to the factors of $\mathcal{M}_1, \dots, \mathcal{M}_m$, then determine which aspects of the data space correspond to what factor, and ultimately focus only on the factors of interest.

1.1 Related Work

In this paper, we build on spectral embedding methods for dimensionality reduction and data representation (Belkin and Niyogi, 2003; Coifman and Lafon, 2006). These methods construct a data-dependent graph Laplacian which approximates the Laplacian operator on the data manifold. Furthermore, the eigenvectors of the graph Laplacian converge to the eigenfunctions of the manifold Laplacian. By mapping the observations to their eigenvector coordinates one can obtain a low-dimensional embedding of the data space (Bates, 2014).

One work that is closely related to ours is the spectral method of (Singer, 2006a) for linear independent component analysis (ICA), as it is also based on the separability of the Laplacian on a product space. This was later extended to a method for non-linear ICA that assumes that the latent space has a product structure and that the Jacobian of the mapping $(\theta_1, \dots, \theta_m) \mapsto \mathbf{x}$ is either known or can be estimated in some way (Singer and Coifman, 2008). Recently, (Rodolà et al., 2019) applications of spectral representations on product manifolds to shape analysis problems such as finding correspondences.

In the neural-network literature, learning *disentangled representations* has drawn interest in recent years (Reed et al., 2014; Locatello et al., 2019; Tran et al., 2017; Siddharth et al., 2017; Kim and Mnih, 2018). Unlike our work, which makes strict assumptions on the data manifold, the disentanglement literature seems to treat broader classes of data-generating processes and does not have a formal notion of what constitutes a disentangled representation.

2 Theoretical background

We begin with a brief review of some useful properties of the Laplace operator, and then examine how the Helmholtz equation separates over product domains in \mathbb{R}^d . Lastly, we go over results that relate the graph Laplacian to the continuous Laplacian, and introduce the connection to diffusion maps.

2.1 Properties of the Laplace operator

Let f be a real-valued, twice-differentiable function and \mathcal{M} a compact d -dimensional manifold through

out this section. Let $\nabla = \nabla_{\mathcal{M}}$ be the gradient and $\Delta = \Delta_{\mathcal{M}} = \nabla_{\mathcal{M}} \cdot \nabla_{\mathcal{M}}$ the manifold Laplacian (or Laplace-Beltrami operator). We consider the solutions of the Helmholtz equation with Neumann boundary conditions

$$-\Delta f(\mathbf{x}) = \lambda f(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{M} \quad (3)$$

$$\nabla f(\mathbf{x}) \cdot \nu(\mathbf{x}) = 0, \quad \forall \mathbf{x} \in \partial \mathcal{M} \quad (4)$$

where ν is a normal unit vector to the boundary. The functions f are known as the *Neumann eigenfunctions*, and together with the corresponding eigenvalues λ they comprise the spectrum of the Laplacian operator on \mathcal{M} . In this paper, we will only consider the particular Neumann boundary condition specified in (4), so we refer to the Neumann eigenfunctions simply as the eigenfunctions. A well-known property of the eigenfunctions is the following.

Theorem 1. *Let $\{f_k\}_{k=1}^{\infty}$ be the Neumann eigenfunctions, and $\{\lambda_k\}_{k=1}^{\infty}$ be the corresponding eigenvalues of a compact domain $\Omega \subset \mathbb{R}^d$. Then*

- (i) $0 = \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$ and $\lambda_k \rightarrow \infty$ as $k \rightarrow \infty$.
- (ii) The eigenfunctions $\{f_m\}_{m=1}^{\infty}$ form a complete orthonormal basis for $L^2(\Omega)$.

For a complete proof and discussion on generalizations, see (Folland, 1976). In particular, the second result of Theorem 1 tells us that we can express any square-integrable function over the data manifold in terms of the eigenfunctions. Analogs of these results also hold for more general Riemannian manifolds.

2.2 The Laplacian over product spaces

As we shift to the case when our data comes from a product manifold \mathcal{M} , we can take advantage of some key separability properties of the Laplacian. For simplicity, we treat the case when $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$.

Proposition 1. *Let $f_1 : \mathcal{M}_1 \rightarrow \mathbb{R}$ and $f_2 : \mathcal{M}_2 \rightarrow \mathbb{R}$ be twice-differentiable functions such that $\Delta_{\mathcal{M}_1} f_1 = \lambda_1 f_1$ and $\Delta_{\mathcal{M}_2} f_2 = \lambda_2 f_2$. Take $\pi_i : \mathcal{M} \rightarrow \mathcal{M}_i$ to be the projection of \mathcal{M} onto \mathcal{M}_i for $i = 1, 2$. We can then define the natural extension of f_i to \mathcal{M} via $g_i = f_i \circ \pi_i$. It follows that $\Delta_{\mathcal{M}}(g_1 g_2) = (\lambda_1 + \lambda_2) g_1 g_2$.*

For a proof, see for example (Canzani, 2013, Section 4.6). Proposition 1 tells us that the product of the eigenfunctions of \mathcal{M}_1 and \mathcal{M}_2 are eigenfunctions of \mathcal{M} and that the corresponding eigenvalue is equal to the sum of the eigenvalues in \mathcal{M}_1 and \mathcal{M}_2 . Note that this result easily extends to the general product of $m > 2$ Riemannian manifolds.

Thus far, we can conclude that products of eigenfunctions of manifold factors are eigenfunctions of

the product manifold. The converse is actually also true, and this can be shown by an application of the Stone-Weierstrass theorem (Canzani, 2013). Hence, the eigenfunctions of $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$ are precisely

$$\{f_i g_j \mid \Delta_{\mathcal{M}_1} f_i = \lambda_i f_i, \Delta_{\mathcal{M}_2} g_j = \mu_j g_j\} \quad (5)$$

with corresponding eigenvalues $\{\lambda_i + \mu_j\}_{i,j}$. For the $m > 2$ case, the eigenfunctions of \mathcal{M} are of the form

$$f^{(k_1, \dots, k_m)} = \prod_{i=1}^m f_{k_i}^{(i)}, \quad (6)$$

with corresponding eigenvalues

$$\lambda^{(k_1, \dots, k_m)} = \sum_{i=1}^m \lambda_{k_i}^{(i)}. \quad (7)$$

Note that since $\lambda_0^{(i)} = 0$ and $f_0^{(i)} = (\sqrt{\text{vol}(\mathcal{M}_i)})^{-1}$ for all $i = 1, \dots, m$, we must have $\lambda^{(0)} = 0$ and $f^{(0)} = \prod_i (\sqrt{\text{vol}(\mathcal{M}_i)})^{-1}$ as well. This is consistent with (1).

Example 1 (2D rectangle). *Consider the domain*

$$\mathcal{M} = [0, a] \times [0, b] \subset \mathbb{R}^2. \quad (8)$$

The Laplacian eigenfunctions with Neumann boundary conditions are

$$f_{m,n}(x, y) = \frac{1}{2} \cos\left(\frac{m\pi}{a}x\right) \cos\left(\frac{n\pi}{b}y\right), \quad (9)$$

with eigenvalues

$$\lambda_{m,n} = \lambda_m + \lambda_n = \pi^2 \left(\frac{m^2}{a^2} + \frac{n^2}{b^2} \right). \quad (10)$$

The eigenfunctions are precisely the product of the eigenfunctions on the closed intervals $[0, a]$ and $[0, b]$, which are $\cos(\frac{m\pi}{a}x)$ and $\cos(\frac{n\pi}{b}y)$. The eigenvalues of the product space are also the sums of the corresponding eigenvalues of the intervals.

For the Laplacian eigenfunctions and eigenvalues on other domains, see the review by Grebenkov and Nguyen (2013).

2.3 Graph Laplacians and diffusion maps

The results of the previous sections hold for the continuous Laplacian, but in practice we only have limited observable samples $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^D$. Thus, we must turn to the discrete analog of the Laplace operator, the graph Laplacian. To start, define the fully-connected graph (G, E, V) whose edge weight matrix $W \in \mathbb{R}^{n \times n}$ is described by

$$W_{ij} = \exp(-\|x_i - x_j\|^2/\epsilon). \quad (11)$$

We then compute the random walk matrix $A = D^{-1}W$, where D is the diagonal degree matrix,

$$D_{ii} = \sum_{j=1}^n W_{ij}. \quad (12)$$

The matrix A can be viewed as a transition probability matrix, and the i th entry of the j th column of A^t describes the probability cloud of a random walk on G ending up in position j at time t , given a starting position at i . The diffusion map (Coifman and Lafon, 2006; Belkin and Niyogi, 2003) is a mapping $\phi : V \rightarrow \mathbb{R}^{n-1}$ given by

$$\phi_t(v_i) = [\lambda_2^t \phi_2(i), \lambda_3^t \phi_3(i), \dots, \lambda_n^t \phi_n(i)]$$

Here, ϕ_j are the right eigenvectors of A and $\phi_j(i)$ denotes the i th coordinate of ϕ_j . The coordinates of the embedding are known as diffusion coordinates. It can be shown that the diffusion distance, which is defined as the Euclidean distance in the diffusion coordinates, corresponds to the difference between different probability clouds (Coifman and Lafon, 2006). Alternatively, if we view the observed data as samples from a smooth signal on the manifold, then the eigenvectors can also be interpreted as a Fourier basis for these signals (Coifman and Lafon, 2006; Lee and Izbecki, 2016).

The standard graph Laplacian is defined as $L = D - W$. We can also define a normalized version called the random walk Laplacian, which is $L_{rw} = D^{-1}L = I - A$. The random walk Laplacian has several properties that are discussed in the next section which make it the most suitable approximation. Computing the eigenvectors and eigenvalues of L_{rw} directly is computationally inefficient, but we can compute them indirectly via the spectrum of the symmetric Laplacian, $L_{sym} = D^{-1/2}LD^{-1/2}$. The spectrum of L_{rw} can then be derived directly from the spectrum of L_{sym} (Von Luxburg, 2007). A condensed version of this procedure can also be found in the thesis of Lafon (2004).

2.4 Limit of the graph Laplacian

The graph Laplacian constructed from an i.i.d. sample of points on a manifold converges to a linear differential operator. For the standard kernel-based construction of the graph and when the samples are drawn uniformly the limiting operator is the Laplace-Beltrami operator $\Delta_{\mathcal{M}}$ (Hein et al., 2005; Singer, 2006b; Belkin and Niyogi, 2008). If the data instead has a non-uniform density $p(\mathbf{x})$, the graph Laplacian converges to the Fokker-Planck operator on the manifold, which has an additional drift term given by the log-density $U(x) = -2 \log p(x)$ (Nadler et al., 2005; Coifman and Lafon, 2006),

$$\mathcal{L}f = \Delta_{\mathcal{M}}f - \nabla U \cdot \nabla f. \quad (13)$$

Alternative graph constructions, such as the k -nearest-neighbor graph also converge to a Fokker-Planck operator (Ting et al., 2010).

In addition to the pointwise convergence, several works have proved the spectral convergence of the graph Laplacian. i.e. the convergence of the eigenvalues and eigenvectors of graph Laplacians (von Luxburg et al., 2008; Rosasco et al., 2010; García Trillos and Slepčev, 2018; García Trillos et al., 2020).

3 Method

As discussed in Section 2.2, each eigenfunction φ_k on the product space $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$ is the product of an eigenfunction φ_i on \mathcal{M}_1 and an eigenfunction φ_j on \mathcal{M}_2 . A corollary of the convergence of the graph Laplacian (see Section 2.4) is the following: given a large sample of points on \mathcal{M} (or any isometric embedding of it), if we construct the graph Laplacian from the sample then each Laplacian eigenvector φ_k should be approximately equal (up to normalization) to an elementwise product $\varphi_i\varphi_j$ where φ_i and φ_j approximate eigenfunctions on \mathcal{M}_1 and \mathcal{M}_2 respectively. If either φ_i or φ_j correspond to an eigenfunction with eigenvalue $\lambda = 0$ we shall say φ_k is a *factor* eigenvector. Otherwise, we call $\varphi_k \approx \varphi_i\varphi_j$ a *product* eigenvector. Note that for data sampled uniformly from a connected manifold, the only eigenfunctions with $\lambda = 0$ are the constant functions.

The main idea behind our method is to first differentiate between the factor eigenvectors and the product eigenvectors and then to divide the factor eigenvectors into two sets that correspond to the eigenspace of \mathcal{M}_1 and \mathcal{M}_2 . This is based on the observation that if $\varphi_k \approx \varphi_i\varphi_j$ then this is a hint that φ_i and φ_j belong to different factor manifolds. Mathematically, our method is based on two assumptions:

Assumption 1. *The mapping $(\theta_1, \dots, \theta_m) \mapsto \mathbf{x}$ is an isometric embedding into \mathbb{R}^D .*

Remark 1. Embeddings of Riemannian manifolds into Euclidean space exists for any Riemannian manifold (Nash, 1954; Kuiper, 1955). As isometric embeddings preserve the Riemannian metric, the spectral properties are not affected by the specific embedding.

Assumption 2. *The latent variables $\theta_1, \dots, \theta_m$ are drawn independently of each other.*

Remark 2. This assumption may be relaxed by using the diffusion maps normalization for the graph Laplacian (Coifman and Lafon, 2006). In contrast to the random-walk graph Laplacian, which converges to a density-dependent Fokker-Planck operator, with the diffusion maps normalization the graph Laplacian converges to the Laplace-Beltrami operator which does

not depend on the density.

We now present an algorithm for factoring the set of eigenvectors of the graph Laplacian.

3.1 Factorizing product eigenvectors

First, we find for each eigenvector φ_k the pair φ_i, φ_j whose elementwise product is closest. This “closeness” is measured by the absolute cosine similarity,

$$S(\varphi_k, \varphi_i\varphi_j) := \frac{|\langle \varphi_k, \varphi_i\varphi_j \rangle|}{\|\varphi_k\| \|\varphi_i\varphi_j\|} \in [0, 1]. \quad (14)$$

We take the absolute value of the dot product because of the sign ambiguity of the computed eigenvectors. For each k , we find the combination (i, j) with the highest similarity.

To avoid computing the element-wise product for all $k(k-1)/2$ combinations, we skip triplets (i, j, k) for which $|\lambda_i + \lambda_j - \lambda_k| > \delta$ for some certain eigenvalue threshold $\delta > 0$, which is specified as a parameter. This “eigenvalue criterion” is based on the results of Proposition 1. In terms of the quality of the match, the eigenvalue criterion may only act as a coarse filter. Thus, we additionally filter out triplets whose highest similarity is less than some threshold γ , which we call the “similarity criterion.” In total, we store at most N combinations, one for every $1 \leq k \leq N$. This process is summarized in Algorithm 1.

Algorithm 1: Identification of individual factors

Data: Eigenvectors $\{\varphi_1, \dots, \varphi_N\}$ of L_{rw} .

Result: List of triplets (i, j, k) where $\varphi_k \approx \varphi_i\varphi_j$ and their corresponding similarity scores.

```

for  $k \leftarrow 1 \dots N$  do
     $\maxS \leftarrow 0$ ;
    for  $i, j < k$  do
        if  $|\lambda_i + \lambda_j - \lambda_k| < \delta$  and
            $S(\varphi_k, \varphi_i\varphi_j) > \maxS$  then
             $\maxS \leftarrow S(\varphi_k, \varphi_i\varphi_j)$ ;
             $i_{\max} \leftarrow i$ ;
             $j_{\max} \leftarrow j$ ;
        end
    end
    if  $\maxS > \gamma$  then
        add  $(i_{\max}, j_{\max}, k)$  to triplets;
    end
end
    
```

The result of the algorithm is a list of index triplets (i, j, k) , indicating that the best factorization we could find for φ_k is $\varphi_i\varphi_j$. Next, we use the triplets to divide the factor eigenvectors into two sets that correspond to the manifolds \mathcal{M}_1 and \mathcal{M}_2 .

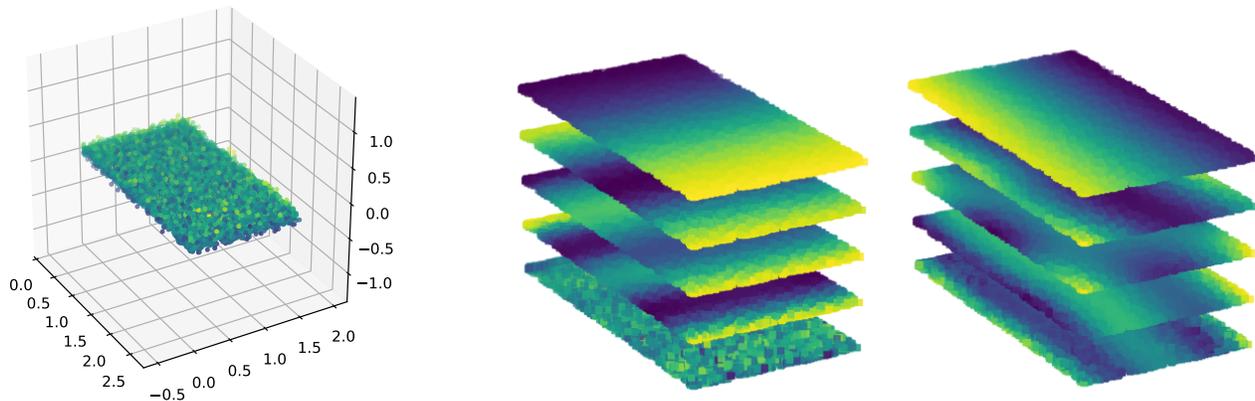


Figure 1: **Rectangle with noise.** The algorithm was run on 10,000 samples with $\delta = 0.5$, $\gamma = 0.85$, $K = 0$. For clarity, the eigenvectors shown are plotted on the non-noisy ground truth x and y coordinate plane. **Left.** The original data, with depth of noise indicated by color. **Middle.** The first five factor eigenvectors associated with the line $[0, 1 + \sqrt{\pi}]$ as determined by the algorithm. **Right.** The first five factor eigenvectors associated with the line $[0, 1.5]$ as determined by the algorithm.

3.2 Assigning the factor eigenvectors to factor manifolds

We can use the list of triplets to split the candidate factor eigenvectors into two separate subsets using a MAX-CUT algorithm.

To do so, we store the computed similarity scores in a $T \times T$ separability matrix C , where T is the number of eigenvectors which appear as a factor eigenvector in any triplet. Each entry of C is defined by

$$C_{i,j} = S(\varphi_k, \varphi_i \varphi_j), \quad (15)$$

and note that C is upper-triangular with zero diagonal. Here, the similarity score is repurposed as a representation of the factorization assignments between each pair of factor eigenvectors in a triplet. The separation of two eigenvectors into eigenspaces of distinct manifold factors should be prioritized depending on the similarity of the triplet in which they appear. This naturally lends to a graph-cut approach to determine the ideal separation.

We use the symmetrized separability matrix, $C + C^T$, as the input to our MAX-CUT algorithm to obtain two groups of factor eigenvectors. Since the original MAX-CUT problem is NP-hard, we use the semi-definite program (SDP) relaxation described by (Goemans and Williamson, 1995).

3.3 Implementation details and runtime

The entire algorithm was implemented in Python, and the MAX-CUT SDP was implemented using CVXPY (Diamond and Boyd, 2016). The first part of the algorithm involves computing the diffusion map of the

data. The runtime for this step is dependent on the Scikit-learn randomized SVD algorithm (Pedregosa et al., 2011; Martinsson et al., 2011). In the subsequent step, we must find the best triplet for each of $k = 1, \dots, N$ out of all $k(k-1)/2$ triplets. Taking into account the element-wise vector multiplication at each step, this takes $O(nN^3)$ in the worst case. In practice, however, the eigenvalue criterion eliminates the multiplication computation for the vast majority of triplets. The last part of the algorithm consists of iterating through the list of triplets to form the dissimilarity matrix C and then using a MAX-CUT SDP solver. Table 1 gives a breakdown of runtimes for each of these parts on 10,000 samples from a rectangle with three-dimensional noise.

Table 1: Running times (in seconds) for each part of the algorithm on a 2017 2.9 GHz Quad-Core Intel Core i7 processor. The experiment was run for various N on 10,000 samples from a rectangle $\mathcal{M} = [0, 1 + \sqrt{\pi}] \times [0, 1.5]$ with Gaussian noise $\mathcal{N}(0, 0.1)$ in the z -direction, where N is the total number of computed eigenvectors. The algorithm parameters were set to $\delta = 2.0$, $\gamma = 0.75$. For each N , we take the average runtime over five trials. The third column refers to Algorithm 1 and the fourth column refers to the construction of the separability matrix and the MAX-CUT SDP described in Section 3.2.

N	Diffusion map	Alg. 1	Alg. 2
50	5.55	1.24	0.03
100	5.98	7.54	0.026
200	8.30	43.19	0.03
400	12.49	251.68	0.044

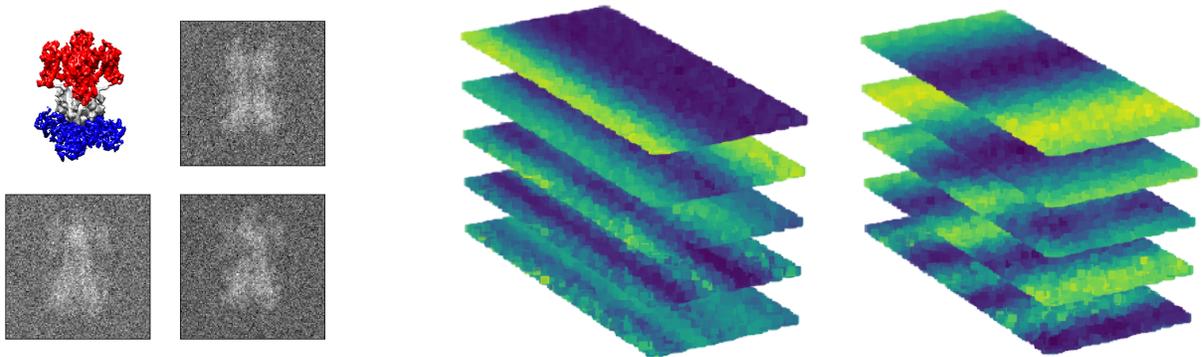


Figure 2: **Two-factor noisy cryo-EM data.** The algorithm was run on 10,000 images with the parameters $\delta = 1.0$, $\gamma = 0.80$. For clarity, the eigenvectors shown are plotted on the ground-truth values of the rotation angles $[0, 360]$ and the x -stretch $[-20, 20]$. **Left.** A synthetic cryo-EM image of the potassium ion channel with two independently deformable subunits is shown in the top-left. The red portion rotates randomly around the z -axis, and the blue portion stretches independently along the x -axis. Simulated cryo-EM images are shown in the other squares. **Middle.** The factor eigenvectors associated with x -stretch. **Right.** The factor eigenvectors associated with the rotation.

4 Simulations

We tested the algorithm on two different types of data: uniform samples from a noisy rectangle and synthetic cryo-EM data.

4.1 2D noisy rectangle

As a baseline, we ran the algorithm on data sampled from a rectangle that lies in \mathbb{R}^3 . The data was generated by taking random samples from a product of lines $\ell_1 = [0, \sqrt{\pi} + 1]$ and $\ell_2 = [0, 1.5]$, and then adding Gaussian noise in the z -direction. We ran two different trials on the rectangle data: first, using 10,000 samples with Gaussian noise applied in the z -direction; second, repeating the first trial except with additional Gaussian noise added in the x and y directions.

For each trial, we use $\gamma = 0.85$ and $\delta = 0.5$. The results of the separations are shown in Figure 1. The two groups of eigenvectors clearly show cosine waves running alongside the x and y directions, which precisely matches the actual eigenfunctions for this domain.

4.2 Analysis of gateway criteria

Our algorithm has two main input parameters which act as thresholds: δ for the eigenvalue criterion and γ for the similarity criterion. A question one might have is, for a given product eigenvector φ_k , how much of an outlier will the similarity of the true combination of factor eigenvectors $\varphi_i \varphi_j$ be compared to the similarities of all other pairs? To investigate this, we create scatterplots of the similarities against the eigenvalue

criteria error over all pairs for various product eigenvectors, shown in Figure 3. We see that in most cases the similarity score for the chosen triplet is indeed an outlier from the rest of the distribution. Unsurprisingly, for product eigenvectors with lower eigenvalues the highest similarity stands out more significantly, as this part of the spectrum is more stable. As the eigenvalue of the product eigenvector increases, the best pair, or the one that has the highest similarity score and satisfies the eigenvalue criterion, becomes slightly more ambiguous. However, this ambiguity tends to be between relatively few candidates. This confirms that the eigenvalue and similarity criteria work well in conjunction.

4.3 Single-particle cryo-EM data

A particularly compelling application of manifold factorization is for cryo-electron microscopy (cryo-EM). Cryo-EM is a technique used for imaging proteins and other macromolecules that has fostered major advances in structural biology (Kühlbrandt, 2014), thus garnering it recognition by the 2017 Nobel Prize in Chemistry (Cressey and Callaway, 2017). More recently, cryo-EM has played a significant role in unraveling the structural properties of the novel coronavirus SARS-CoV-2 (Zimmer, 2020).

In cryo-EM, molecular samples are rapidly frozen in a thin layer of ice, thus capturing them in their native states (Dubochet et al., 1988). The preservation of this natural state is key for better understanding of their biological function. Once frozen, a transmission electron microscope is used to produce images of

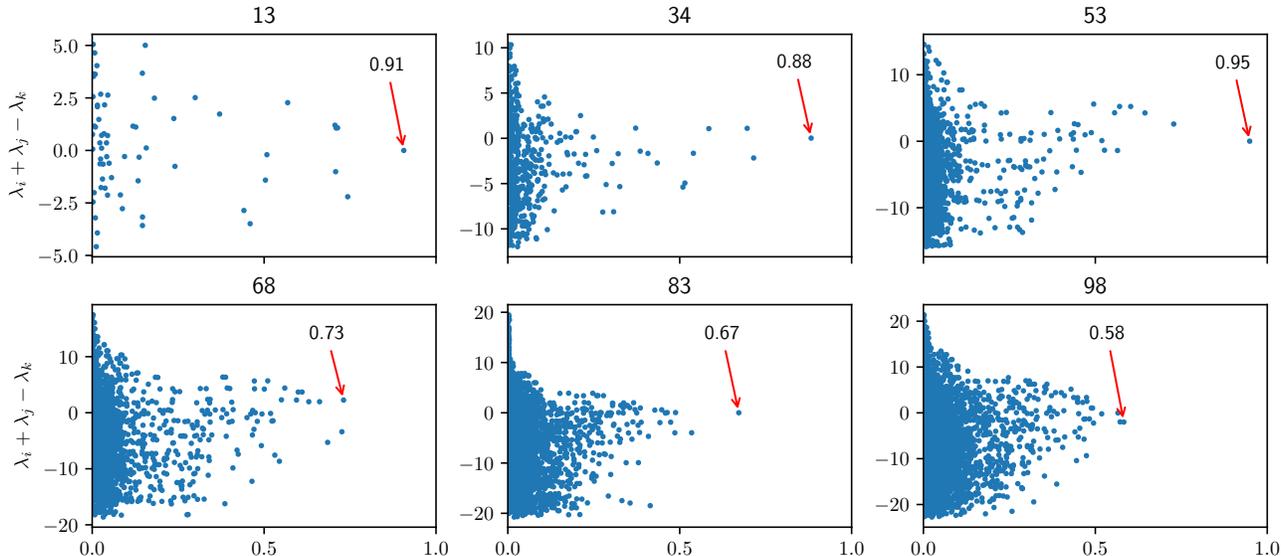


Figure 3: **Analysis of eigenvalue criterion and similarity score for the rectangle in 3D.** We choose six eigenvectors uniformly from the set of product eigenvectors and plot the eigenvalue criteria error and similarity scores for all possible triplets. Indices at the top identify the product eigenvector to which a column corresponds. The highest similarity is marked by a red arrow. The similarity is measured along the x -axis from 0 to 1.

the molecules, which are 2D tomographic projections of their electrostatic potential (Vulović et al., 2013). These 2D images are then used to reconstruct, through a complicated series of computational steps, a 3D reconstruction of the molecule.

The classical cryo-EM reconstruction problem is posed as follows: given n tomographic projection images I_1, \dots, I_n of a particular molecule at random (unknown) orientations, how can we recover the three-dimensional structure of the molecule? A core obstacle to this problem is handling the low signal-to-noise ratio present in the micrographs (Singer and Sigworth, 2020). This task is made even more difficult if we consider the *heterogeneity* problem, which has been the subject of recent literature (Frank, 2018; Nakane et al., 2018; Andén and Singer, 2018; Sorzano et al., 2019; Lederman et al., 2020; Zhong et al., 2020). Here, a separate challenge arises from structural variations arising from motion within the molecule. Namely, for different images I_i we will likely capture different molecular conformations. In particular, several works have focused on the application of diffusion maps for the analysis of cryo-EM datasets with continuous heterogeneity (Dashti et al., 2014; Schwander et al., 2014; Moscovich et al., 2020; Zelesko et al., 2020; Dashti et al., 2020).

For molecules that exhibit two or more independent continuous motions, the molecular shape manifold of 3D electrostatic densities is a product space, so one can try using our algorithm to factor it. For our

experiments we used the model of a potassium ion channel with two independently deformable parts: a spinning subunit with 360° of rotation, and a subunit that stretches in the xy -plane (see Figure 2). To create the images, we sampled 10,000 conformations of the molecule with uniformly random angles and x, y stretches ranging from $[-20, 20]$ in either direction. All the images were projected from a single view. This models the computational pipeline of (Dashti et al., 2020), which computes a diffusion maps embedding for each viewing direction separately. Finally, we use PCA with four components to preprocess the raw images and transform them so that they are zero-centered with unit variance.

We perform three different trials: first, a dataset of molecules with rotations and stretches in the x direction only and low noise; then, a dataset of molecules with rotations and stretches in the x -direction only and high noise; and finally, a dataset of molecules with rotations and stretches in both the x and y directions and low noise. For each trial, use $\gamma = 0.80$ and $\delta = 1.0$. The results for the second trial are shown in Figure 2. The eigenvectors corresponding to the x -stretch manifold exhibit the ground truth cosine waves, and the eigenvectors corresponding to the rotation manifold (a circle) exhibit the ground truth sine waves. Note that even though the latent space is a product space, Assumption 1 does not strictly hold due to the additive Gaussian noise. Nonetheless, our method works nicely on this dataset.

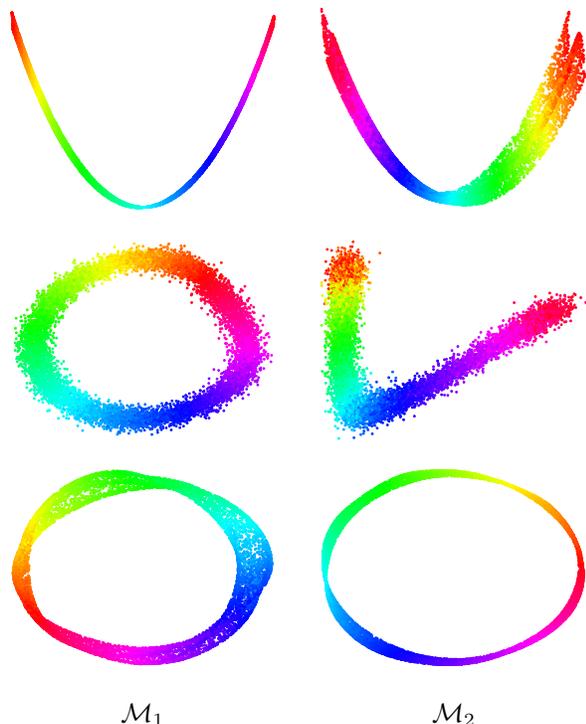


Figure 4: **Laplacian eigenmap embeddings.** Each scatterplot shows 10,000 samples of the first two non-trivial eigenvectors found by the algorithm. The colors represent the ground truth data in the parameter space. **Top row.** Samples from a rectangle with z -direction noise. The algorithm was run with $\delta = 0.5$, $\gamma = 0.85$. **Middle row.** Samples from noisy cryo-EM data. The algorithm was run with $\delta = 1.0$, $\gamma = 0.85$. **Bottom row.** Samples from a torus. The algorithm was run with $\delta = 1.0$ and $\gamma = 0.6$.

4.4 Laplacian eigenmap embeddings

We can visualize the manifold factors by plotting the Laplacian eigenmap embeddings. Figure 4 shows the eigenmap embeddings for three different experiments: the rectangle with noise in the z -direction; noisy cryo-EM data with a rotational component and a stretch component in the x -direction; and a torus. In each case, the embeddings are in agreement with the actual structure of the underlying domain.

These embeddings show that the results of the algorithm can be used to accomplish two objectives. First, we can use representative factor eigenvectors from each manifold factor to get a more condensed representation of the data, thus achieving dimensionality reduction. This is similar to using diffusion coordinates for dimensionality reduction, but diffusion coordinates generally appear in no particular order. Additionally, by separating the factor eigenvectors according to the sepa-

rate manifold factors, we also “de-mix” these diffusion coordinates into a more interpretable order.

5 Conclusion

In this paper, we presented a spectral method for non-linear dimensionality reduction and data representation. Our method is based on manifold factorization and applies to data sets with two or more independent degrees of freedom. We tested our algorithm on electron-microscope snapshots of a protein with two independently moving subunits and demonstrated that we can recover the motion manifolds for each subunit. Hence, our approach shows promise as a tool for the analysis of macromolecules with continuous degrees of freedom.

While our discussion and experiments are limited to the two-factor case, our algorithm can be extended to the general case of K independent latent variables by replacing MAX-CUT with MAX-K-CUT (Newman, 2018) or similar procedures. Another interesting direction for future work is to study the potential of our approach for analyzing broader classes of data sets, other than product spaces, where the Laplacian operator is still separable.

Reproducibility

Code for reproducing the figures in this paper can be found at: <https://github.com/sxzhang25/product-manifold-learning>

References

- Joakim Andén and Amit Singer. Structural Variability from Noisy Tomographic Projections. *SIAM Journal on Imaging Sciences*, 11(2):1441–1492, 2018. doi:10.1137/17M1153509.
- Jonathan Bates. The embedding dimension of Laplacian eigenfunction maps. *Applied and Computational Harmonic Analysis*, 37(3):516–530, 2014. doi:10.1016/j.acha.2014.03.002.
- Mikhail Belkin and Partha Niyogi. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Computation*, 15(6):1373–1396, 2003. doi:10.1162/089976603321780317.
- Mikhail Belkin and Partha Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, 2008. doi:10.1016/j.jcss.2007.08.006.
- Yaiza Canzani. Analysis on Manifolds via the Laplacian (lecture notes), 2013.
- Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computa-*

- tional Harmonic Analysis*, 21(1):5–30, 2006. doi:[10.1016/j.acha.2006.04.006](https://doi.org/10.1016/j.acha.2006.04.006).
- Daniel Cressey and Ewen Callaway. Cryo-electron microscopy wins chemistry Nobel. *Nature*, 550(7675):167–167, 2017. doi:[10.1038/nature.2017.22738](https://doi.org/10.1038/nature.2017.22738).
- Ali Dashti et al. Trajectories of the ribosome as a Brownian nanomachine. *Proceedings of the National Academy of Sciences*, 111(49):17492–17497, 2014. doi:[10.1073/pnas.1419276111](https://doi.org/10.1073/pnas.1419276111).
- Ali Dashti et al. Retrieving functional pathways of biomolecules from single-particle snapshots. *Nature Communications*, 11(1):4734, 2020. doi:[10.1038/s41467-020-18403-x](https://doi.org/10.1038/s41467-020-18403-x).
- Steven Diamond and Stephen Boyd. CVXPY: A Python-Embedded Modeling Language for Convex Optimization. Technical report, 2016. <http://www.cvxpy.org/>.
- Jacques Dubochet et al. Cryo-electron microscopy of vitrified specimens. *Quarterly Reviews of Biophysics*, 21(2):129–228, 1988. doi:[10.1017/S0033583500004297](https://doi.org/10.1017/S0033583500004297).
- Gerald B. Folland. *Introduction to Partial Differential Equations*. Princeton University Press, 1976. ISBN 9780691213033. doi:[10.2307/j.ctvzsmfng](https://doi.org/10.2307/j.ctvzsmfng).
- Joachim Frank. New Opportunities Created by Single-Particle Cryo-EM: The Mapping of Conformational Space. *Biochemistry*, 57(6):888–888, 2018. doi:[10.1021/acs.biochem.8b00064](https://doi.org/10.1021/acs.biochem.8b00064).
- Nicolás García Trillos and Dejan Slepčev. A variational approach to the consistency of spectral clustering. *Applied and Computational Harmonic Analysis*, 45(2):239–281, 2018. doi:[10.1016/j.acha.2016.09.003](https://doi.org/10.1016/j.acha.2016.09.003).
- Nicolás García Trillos, Moritz Gerlach, Matthias Hein and Dejan Slepčev. Error Estimates for Spectral Convergence of the Graph Laplacian on Random Geometric Graphs Toward the Laplace–Beltrami Operator. *Foundations of Computational Mathematics*, 20(4):827–887, 2020. doi:[10.1007/s10208-019-09436-w](https://doi.org/10.1007/s10208-019-09436-w).
- Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. doi:[10.1145/227683.227684](https://doi.org/10.1145/227683.227684).
- Denis S. Grebenkov and Binh-Thanh Nguyen. Geometrical Structure of Laplacian Eigenfunctions. *SIAM Review*, 55(4):601–667, 2013. doi:[10.1137/120880173](https://doi.org/10.1137/120880173).
- Matthias Hein, Jean-Yves Audibert and Ulrike von Luxburg. From Graphs to Manifolds – Weak and Strong Pointwise Consistency of Graph Laplacians. In *International Conference on Computational Learning Theory (COLT)*, pages 470–485, 2005. doi:[10.1007/11503415_32](https://doi.org/10.1007/11503415_32).
- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *International Conference on Machine Learning (ICML)*, volume 6, pages 4153–4171, 2018.
- W. Kühlbrandt. The Resolution Revolution. *Science*, 343(6178):1443–1444, 2014. doi:[10.1126/science.1251652](https://doi.org/10.1126/science.1251652).
- Nicolaas H. Kuiper. On C1-isometric imbeddings. I. *Indagationes Mathematicae (Proceedings)*, 58:545–556, 1955. doi:[10.1016/S1385-7258\(55\)50075-8](https://doi.org/10.1016/S1385-7258(55)50075-8).
- Stephane S. Lafon. *Diffusion Maps and Geometric Harmonics*. PhD thesis, Yale University, 2004.
- Roy R. Lederman, Joakim Andén and Amit Singer. Hyper-molecules: on the representation and recovery of dynamical structures for applications in flexible macro-molecules in cryo-EM. *Inverse Problems*, 36(4):044005, 2020. doi:[10.1088/1361-6420/ab5ede](https://doi.org/10.1088/1361-6420/ab5ede).
- Ann B. Lee and Rafael Izbicki. A spectral series approach to high-dimensional nonparametric regression. *Electronic Journal of Statistics*, 10(1):423–463, 2016. doi:[10.1214/16-EJS1112](https://doi.org/10.1214/16-EJS1112).
- John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer New York, 2012. ISBN 978-1-4419-9981-8. doi:[10.1007/978-1-4419-9982-5](https://doi.org/10.1007/978-1-4419-9982-5).
- Francesco Locatello et al. Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations. *International Conference on Machine Learning (ICML)*, pages 4114–4124, 2019.
- Per Gunnar Martinsson, Vladimir Rokhlin and Mark Tygert. A randomized algorithm for the decomposition of matrices. *Applied and Computational Harmonic Analysis*, 30(1):47–68, 2011. doi:[10.1016/j.acha.2010.02.003](https://doi.org/10.1016/j.acha.2010.02.003).
- Amit Moscovich, Amit Halevi, Joakim Andén and Amit Singer. Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes. *Inverse Problems*, 36(2):024003, 2020. doi:[10.1088/1361-6420/ab4f55](https://doi.org/10.1088/1361-6420/ab4f55).
- Boaz Nadler, Stephane Lafon, Ronald R. Coifman and Ioannis G. Kevrekidis. Diffusion Maps, Spectral Clustering and Eigenfunctions of Fokker-Planck operators. In *Advances in Neural Information Processing Systems (NIPS)*, pages 955–962, 2005.
- Takanori Nakane, Dari Kimanius, Erik Lindahl and Sjors HW Scheres. Characterisation of molecular

- motions in cryo-EM single-particle data by multi-body refinement in RELION. *eLife*, 7:1–18, 2018. doi:[10.7554/eLife.36861](https://doi.org/10.7554/eLife.36861).
- John Nash. C^1 Isometric Imbeddings. *The Annals of Mathematics*, 60(3):383–396, 1954. doi:[10.2307/1969840](https://doi.org/10.2307/1969840).
- Alantha Newman. Complex semidefinite programming and max-k-cut. In *Symposium on Simplicity in Algorithms (SOSA)*, volume 61, pages 13:1–13:11, 2018. doi:[10.4230/OASIS.SOSA.2018.13](https://doi.org/10.4230/OASIS.SOSA.2018.13).
- Fabian Pedregosa et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Scott Reed, Kihyuk Sohn, Yuting Zhang and Honglak Lee. Learning to disentangle factors of variation with manifold interaction. *International Conference on Machine Learning (ICML)*, pages 1431–1439, 2014.
- Emanuele Rodolà, Zorah Löhner, Alex M. Bronstein, Michael M. Bronstein and Justin Solomon. Functional Maps Representation On Product Manifolds. *Computer Graphics Forum*, 38(1):678–689, 2019. doi:[10.1111/cgf.13598](https://doi.org/10.1111/cgf.13598).
- Lorenzo Rosasco, Mikhail Belkin and Ernesto De Vito. On Learning with Integral Operators. *Journal of Machine Learning Research*, 11:905–934, 2010.
- Peter Schwander, Russell Fung and Abbas Ourmazd. Conformations of macromolecules and their complexes from heterogeneous datasets. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1647):1–8, 2014. doi:[10.1098/rstb.2013.0567](https://doi.org/10.1098/rstb.2013.0567).
- N Siddharth et al. Learning Disentangled Representations with Semi-Supervised Deep Generative Models. In *Conference on Neural Information Processing Systems (NIPS)*, 2017.
- A. Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006a. doi:[10.1016/j.acha.2006.03.004](https://doi.org/10.1016/j.acha.2006.03.004).
- A. Singer. Spectral independent component analysis. *Applied and Computational Harmonic Analysis*, 21(1):135–144, 2006b. doi:[10.1016/j.acha.2006.03.003](https://doi.org/10.1016/j.acha.2006.03.003).
- Amit Singer and Ronald R. Coifman. Non-linear independent component analysis with diffusion maps. *Applied and Computational Harmonic Analysis*, 25(2):226–239, 2008. doi:[10.1016/j.acha.2007.11.001](https://doi.org/10.1016/j.acha.2007.11.001).
- Amit Singer and Fred J. Sigworth. Computational Methods for Single-Particle Electron Cryomicroscopy. *Annual Review of Biomedical Data Science*, 3(1):163–190, 2020. doi:[10.1146/annurev-biodatasci-021020-093826](https://doi.org/10.1146/annurev-biodatasci-021020-093826).
- Carlos Oscar S. Sorzano et al. Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy. *Acta Crystallographica Section F Structural Biology Communications*, 75(1):19–32, 2019. doi:[10.1107/S2053230X18015108](https://doi.org/10.1107/S2053230X18015108).
- Ronen Talmon, Israel Cohen, Sharon Gannot and Ronald R. Coifman. Diffusion Maps for Signal Processing: A Deeper Look at Manifold-Learning Techniques Based on Kernels and Graphs. *IEEE Signal Processing Magazine*, 30(4):75–86, 2013. doi:[10.1109/MSP.2013.2250353](https://doi.org/10.1109/MSP.2013.2250353).
- Joshua B. Tenenbaum, Vin de Silva and John C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323, 2000. doi:[10.1126/science.290.5500.2319](https://doi.org/10.1126/science.290.5500.2319).
- Daniel Ting, Ling Huang and Michael Jordan. An Analysis of the Convergence of Graph Laplacians. In *International Conference on Machine Learning (ICML)*, 2010.
- Luan Tran, Xi Yin and Xiaoming Liu. Disentangled Representation Learning GAN for Pose-Invariant Face Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1283–1292, 2017. doi:[10.1109/CVPR.2017.141](https://doi.org/10.1109/CVPR.2017.141).
- Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. doi:[10.1007/s11222-007-9033-z](https://doi.org/10.1007/s11222-007-9033-z).
- Ulrike von Luxburg, Mikhail Belkin and Olivier Bousquet. Consistency of spectral clustering. *The Annals of Statistics*, 36(2):555–586, 2008. doi:[10.1214/009053607000000640](https://doi.org/10.1214/009053607000000640).
- Miloš Vulović et al. Image formation modeling in cryo-electron microscopy. *Journal of Structural Biology*, 183(1):19–32, 2013. doi:[10.1016/j.jsb.2013.05.008](https://doi.org/10.1016/j.jsb.2013.05.008).
- Nathan Zelesko, Amit Moscovich, Joe Kileel and Amit Singer. Earthmover-Based Manifold Learning for Analyzing Molecular Conformation Spaces. In *International Symposium on Biomedical Imaging (ISBI)*, pages 1715–1719, 2020. doi:[10.1109/ISBI45749.2020.9098723](https://doi.org/10.1109/ISBI45749.2020.9098723).
- Ellen D. Zhong, Tristan Bepler, Joseph H. Davis and Bonnie Berger. Reconstructing continuous distributions of 3D protein structure from cryo-EM images. In *International Conference on Learning Representations (ICLR)*, pages 1–20, 2020.
- Carl Zimmer. The Coronavirus Unveiled, oct 2020, The New York Times. <http://nyti.ms/2G06PDV>.