
Multi-Frequency Vector Diffusion Maps

Yifeng Fan¹ Zhizhen Zhao¹

Abstract

We introduce multi-frequency vector diffusion maps (MFVDM), a new framework for organizing and analyzing high dimensional datasets. The new method is a mathematical and algorithmic generalization of vector diffusion maps (VDM) and other non-linear dimensionality reduction methods. MFVDM combines different nonlinear embeddings of the data points defined with multiple unitary irreducible representations of the alignment group that connect two nodes in the graph. We illustrate the efficacy of MFVDM on synthetic data generated according to a random graph model and cryo-electron microscopy image dataset. The new method achieves better nearest neighbor search and alignment estimation than the state-of-the-arts VDM and diffusion maps (DM) on extremely noisy data.

1. Introduction

Nonlinear dimensionality reduction methods, such as locally linear embedding (LLE) (Roweis & Saul, 2000), ISOMAP (Tenenbaum et al., 2000), Hessian LLE (Donoho & Grimes, 2003), Laplacian eigenmaps (Belkin & Niyogi, 2002; 2003), and diffusion maps (DM) (Coifman & Lafon, 2006) are invaluable tools for embedding complex data in a low dimensional space and for regression problems on graphs and manifolds. To this end, those methods assume that the high-dimensional data lies on a low dimensional manifold and local affinities in a weighted neighborhood graph are used to learn the global structure of the data. Spectral clustering (Nadler et al., 2006; Von Luxburg, 2007), semi-supervised learning (Zhu, 2006; Goldberg et al., 2009; Yang et al., 2016), out-of-sample extension (Belkin et al., 2006), image denoising (Gong et al., 2010; Singer et al., 2009) share similar geometrical considerations. Those techniques are either directly or indirectly related to the heat kernel for functions on the data. Vector diffusion maps (VDM) (Singer & Wu, 2012) generalizes DM to define heat

kernel for vector fields on the data manifold. The corresponding adjacency matrix is based on edge weights and orthogonal transformations between connected nodes. Using the spectral decomposition of the matrix, VDM defines a metric for the data to indicate the closeness of the data points on the manifold. For some applications, the vector diffusion metric is beneficial, since it takes into account linear transformations, and as a result, it provides a better organization of the data. However, for extremely noisy data, VDM nearest neighbor search may fail at identifying the true nearby points on the manifold. This results in shortcut edges that connect points with large geodesic distances on the manifold.

To address this issue, we introduce a new algorithm called multi-frequency vector diffusion maps (MFVDM) to represent and organize complex high-dimensional data, exhibiting a non-trivial group invariance. To this end, we augment VDM with multiple irreducible representations of the compact group to improve the rotationally invariant nearest neighbor search and the alignment estimation between nearest neighbor pairs, when the initial estimation contains a large number of outliers due to noise. Specifically, we define a set of kernels, denoted by W_k , using multiple irreducible representations of the compact alignment group indexed by integer k and introduce the corresponding frequency- k -VDMs. The MFVDM is constructed by concatenating all the frequency- k -VDMs up to a cut-off k_{\max} . We use the new embeddings to identify nearest neighbors. The eigenvectors of the normalized W_k are used to estimate the pairwise alignments between nearest neighbors. This framework also extends the mathematical theory of cryo-electron microscopy (EM) image analysis (Singer et al., 2011; Hadani & Singer, 2011; Giannakis et al., 2012; Schwander et al., 2012; Dashti et al., 2014). We show that MFVDM outperforms VDM and DM for data sampled from low-dimensional manifolds, when a large proportion of the edge connections are corrupted. MFVDM is also able to improve the nearest neighbor search and rotational alignment for 2-D class averaging in cryo-EM.

2. Preliminaries and Problem Setup

Given a dataset $x_i \in \mathbb{R}^l$ for $i = 1, \dots, n$, we assume that the data lie on or close to a low dimensional smooth manifold \mathcal{X} of intrinsic dimension $d \ll l$. Suppose that \mathcal{G} is a compact Lie group, which has unitary irreducible representations according to Peter-Weyl theorem. The data space \mathcal{X} is closed under \mathcal{G} if for all $g \in \mathcal{G}$ and all $x \in \mathcal{X}$, $g \cdot x \in \mathcal{X}$,

¹Department of Electrical and Computer Engineering, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Illinois, USA. Correspondence to: Yifeng Fan <yifengf2@illinois.edu>.

where ‘ \cdot ’ denotes the group action. The \mathcal{G} -invariant distance between two data points is defined as,

$$d_{ij} = \min_{g \in \mathcal{G}} \|x_i - g \cdot x_j\|, \quad (1)$$

and the associated optimal alignment is,

$$g_{ij} = \arg \min_{g \in \mathcal{G}} \|x_i - g \cdot x_j\|. \quad (2)$$

We assume that the optimal alignment is unique and construct an undirected graph $G = (V, E)$ based on the distances in (1) using the ϵ -neighborhood criterion, i.e. $(i, j) \in E$ iff $d_{ij} < \epsilon$, or κ -nearest neighbor criterion, i.e. $(i, j) \in E$ iff j is one of the κ nearest neighbors of i . The edge weights w_{ij} are defined using a kernel function on the \mathcal{G} -invariant distance $w_{ij} = K_\sigma(d_{ij})$. For example, the Gaussian kernel leads to weights of the form

$$w_{ij} = K_\sigma(d_{ij}) = \exp\left(-\frac{\min_{g \in \mathcal{G}} \|x_i - g \cdot x_j\|^2}{\sigma}\right). \quad (3)$$

The resulting graph is defined on the quotient space $\mathcal{M} := \mathcal{X}/\mathcal{G}$ and is invariant to the group transformation of the individual data points. Under certain conditions, the quotient space \mathcal{M} is also a smooth manifold. We can identify each data point x_i with $v_i \in \mathcal{M}$ and the dimension of \mathcal{M} is lower than the dimension of \mathcal{X} . The unitary irreducible representation of the group g is represented by $\rho_k(g)$. If v_i and v_j are close on the manifold, then the representation $\rho_1(g_{ij})$ of the optimal alignment g_{ij} is an approximation of the local parallel transport operator $P_{x_i, x_j} : T_{v_j}\mathcal{M} \mapsto T_{v_i}\mathcal{M}$ (Singer et al., 2011; Singer & Wu, 2012).

Take cryo-EM imaging as an example, each image is a tomographic projection of a 3D object at an unknown orientation $x \in \text{SO}(3)$ represented by a 3×3 orthogonal matrix $R = [R^1, R^2, R^3]$ satisfying $R^\top R = RR^\top = I$ and $\det R = 1$ (Singer et al., 2011; Hadani & Singer, 2011; Zhao & Singer, 2014). The viewing direction of each image can be represented as a point on the unit sphere, denoted by v ($v = R^3$). The first two columns of the orthogonal matrix R^1 and R^2 correspond to the lifted vertical and horizontal axes of the image in the tangent plane $T_v\mathbb{S}^2$. Therefore, each image can be represented by a unit tangent vector on the sphere and the base manifold is $\mathcal{M} = \text{SO}(3)/\text{SO}(2) = \mathbb{S}^2$. Images with similar v ’s are identified as the nearest neighbors and they can be accurately estimated using (1) from clean images. Registering the centered images corresponds to in-plane rotationally aligning the nearest neighbor images according to (2).

In many applications, noise in the observational data affects the estimations of \mathcal{G} -invariant distances d_{ij} and optimal alignments g_{ij} . This results in shortcut edges in the ϵ -neighborhood graph or κ -nearest neighbor graph, and connects points on \mathcal{M} where the underlying geodesic distances are large.

3. Algorithm

To address this issue of shortcut edges induced by noise, we extend VDM using multiple irreducible representations of the compact alignment group.

3.1. Affinity and mapping

We assume the initial graph G is given along with the optimal alignments on the connected edges. For simplicity and because of our interest in cryo-EM image classification, we focus on $\mathcal{G} = \text{SO}(2)$ and we denote the optimal alignment angle by α_{ij} . The corresponding frequency- k unitary irreducible representations is $e^{ik\alpha_{ij}}$, where $i = \sqrt{-1}$. For points that are nearby on \mathcal{M} , the alignments should have cycle consistency under the clean case, for example, $k(\alpha_{ij} + \alpha_{jl} + \alpha_{li}) \approx 0 \pmod{2\pi}$ for integers $k \in \mathbb{Z}$, if nodes i, j and l are true nearest neighbors. To systematically incorporate the alignment information and impose the consistency of alignments, for a given graph $G = (V, E)$, we construct a set of $n \times n$ affinity matrices W_k ,

$$W_k(i, j) = \begin{cases} w_{ij}e^{2k\alpha_{ij}} & (i, j) \in E, \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

where the edge weights according to (3) are real, $w_{ij} = w_{ji}$ and $\alpha_{ij} = -\alpha_{ji}$ for all $(i, j) \in E$. At frequency k , the weighted degree of node i is:

$$\text{deg}(i) := \sum_{j:(i,j) \in E} |W_k(i, j)| = \sum_{j:(i,j) \in E} w_{ij}, \quad (5)$$

and the degree is identical through all frequencies. We define a diagonal degree matrix D of size $n \times n$, where the i^{th} diagonal entry $D(i, i) = \text{deg}(i)$.

We construct the normalized matrix $A_k = D^{-1}W_k$ which is applied to complex vectors z of length n and each entry $z(i) \in \mathbb{C}$ can be viewed as a vector in $T\mathcal{M}$. The matrix A_k is an averaging operator for vector fields, i.e. $(A_k z)(i) = \frac{1}{\text{deg}(i)} \sum_{j:(i,j) \in E} w_{ij}e^{2k\alpha_{ij}} z(j)$. In our framework, we define affinity between i and j by considering the consistency of the transformations over all paths of length $2t$ that connect i and j . In addition, we also consider the consistencies in the transported vectors at k frequency (see Fig. 1). Intuitively, this means $A_k^{2t}(i, j)$ sums the transformations of all length- $2t$ paths from i to j , and a large value of $|A_k^{2t}(i, j)|$ indicates not only the strength of connection between i and j , but also the level of consistency in the alignment along all connected paths.

We obtain the affinity of i and j by observing the following decomposition:

$$A_k = D^{-1}W_k = D^{-1/2} \underbrace{D^{-1/2}W_k D^{-1/2}}_{S_k} D^{1/2}. \quad (6)$$

Since S_k is Hermitian, it has a complete set of real eigenvalues $\lambda_1^{(k)}, \lambda_2^{(k)}, \dots, \lambda_n^{(k)}$ and eigenvectors $u_1^{(k)}, u_2^{(k)}, \dots, u_n^{(k)}$, where $\lambda_1^{(k)} > \lambda_2^{(k)} > \dots > \lambda_n^{(k)}$. We can express $S_k^{2t}(i, j)$ in terms of the eigenvalues and eigen-




Figure 1. Illustration of multi-frequency edge connection. The \oplus -operation denotes concatenation.

vectors of S_k :

$$S_k^{2t}(i, j) = \sum_{l=1}^n \left(\lambda_l^{(k)} \right)^{2t} u_l^{(k)}(i) \overline{u_l^{(k)}(j)}. \quad (7)$$

Therefore the affinity of i and j at the k^{th} frequency is given by

$$\begin{aligned} |S_k^{2t}(i, j)|^2 &= \sum_{l,r=1}^n \left(\lambda_l^{(k)} \lambda_r^{(k)} \right)^{2t} u_l^{(k)}(i) \overline{u_r^{(k)}(i) u_l^{(k)}(j) u_r^{(k)}(j)} \\ &= \left\langle V_t^{(k)}(i), V_t^{(k)}(j) \right\rangle, \end{aligned} \quad (8)$$

which is expressed by an inner product between two vectors $V_t^{(k)}(i), V_t^{(k)}(j) \in \mathbb{C}^{n^2}$ via the mapping $V_t^{(k)}$:

$$V_t^{(k)} : i \mapsto \left(\left(\lambda_l^{(k)} \lambda_r^{(k)} \right)^t \langle u_l^{(k)}(i), u_r^{(k)}(i) \rangle \right)_{l,r=1}^n. \quad (9)$$

We call this *frequency- k -VDM*.

Truncated mapping: Notice the matrices $I + S_k$ and $I - S_k$ are both positive semi-definite (PSD) due to the following property: $\forall z \in \mathbb{C}^n$ we have

$$\begin{aligned} z^*(I \pm S_k)z &= \\ \sum_{(i,j) \in E} w_{ij} \left| \frac{z(i)}{\sqrt{\deg(i)}} \pm \frac{e^{i\alpha_{ij}} z(j)}{\sqrt{\deg(j)}} \right|^2 &\geq 0. \end{aligned} \quad (10)$$

Therefore all eigenvalues $\{\lambda_i^{(k)}\}_{i=1}^n$ of S_k lie within the interval $[-1, 1]$. Consequently, for large t , most $(\lambda_l^{(k)} \lambda_r^{(k)})^{2t}$ terms in (8) are close to 0, and $|S_k^{2t}(i, j)|^2$ can be well approximated by using only a few of the largest eigenvalues and their corresponding eigenvectors. Hence, we truncate the *frequency- k -VDM* mapping $V_t^{(k)}$ using a cutoff m_k for each frequency k :

$$\hat{V}_t^{(k)} : i \mapsto \left(\left(\lambda_l^{(k)} \lambda_r^{(k)} \right)^t \langle u_l^{(k)}(i), u_r^{(k)}(i) \rangle \right)_{l,r=1}^{m_k}. \quad (11)$$

The affinity of i and j at the frequency k after truncation is given by

$$|\hat{S}_k^{2t}(i, j)|^2 = \left\langle \hat{V}_t^{(k)}(i), \hat{V}_t^{(k)}(j) \right\rangle \approx |S_k^{2t}(i, j)|^2. \quad (12)$$

Remark 1: The truncated mapping not only has the advantage of computational efficiency, but also enhances the robustness to noise since the eigenvectors with smaller eigenvalues are more oscillatory and sensitive to noise.

Multi-frequency mapping: Consider the affinity in (8) for $k = 1, \dots, k_{\max}$, if i and j are connected by multiple paths

with consistent transformations, the affinity $|\hat{S}_k^{2t}(i, j)|^2$ should be large for all k . Then we can combine multiple representations (i.e., combine multiple k) to evaluate the consistencies of the group transformations along connected paths. Therefore, a straightforward way is to concatenate the truncated mappings $\hat{V}_t^{(k)}$ for all $k = 1, 2, \dots, k_{\max}$ as:

$$\hat{V}_t(i) : i \mapsto \left(\hat{V}_t^{(1)}(i); \hat{V}_t^{(2)}(i); \dots; \hat{V}_t^{(k_{\max})}(i) \right), \quad (13)$$

called *multi-frequency vector diffusion maps (MFVDM)*. We define the new affinity of i and j as the inner product of $\hat{V}_t(i)$ and $\hat{V}_t(j)$:

$$\begin{aligned} |\hat{S}^{2t}(i, j)|^2 &:= \sum_{k=1}^{k_{\max}} |\hat{S}_k^{2t}(i, j)|^2 = \sum_{k=1}^{k_{\max}} \left\langle \hat{V}_t^{(k)}(i), \hat{V}_t^{(k)}(j) \right\rangle \\ &= \left\langle \hat{V}_t(i), \hat{V}_t(j) \right\rangle. \end{aligned} \quad (14)$$

MFVDM systematically incorporates the cycle consistencies on the geometric graph across multiple irreducible representations of the transformation group elements (in-plane rotational alignments in this case, see Fig. 1). Using information from multiple irreducible group representations leads to a more robust measure of rotationally invariant similarity.

Remark 2: Empirically, we find the normalized mapping $i \mapsto \frac{\hat{V}_t(i)}{\|\hat{V}_t(i)\|}$ to be more robust to noise than $\hat{V}_t(i)$. A similar phenomenon was discussed in VDM (Singer & Wu, 2012). The normalized affinity is defined as,

$$N_t(i, j) = \left\langle \frac{\hat{V}_t(i)}{\|\hat{V}_t(i)\|}, \frac{\hat{V}_t(j)}{\|\hat{V}_t(j)\|} \right\rangle. \quad (15)$$

Comparison with DM and VDM: Diffusion maps (DM) only consider scalar weights over the edges and the vector diffusion maps (VDM) only take into account consistencies of the transformations along connected edges using only one representation of $\text{SO}(2)$, i.e. $e^{i\alpha_{ij}}$. In this paper, we generalize VDM and use not only one irreducible representation, i.e. $k = 1$, but also higher order k up to k_{\max} .

3.2. Nearest neighbor search and rotational alignment

In this section we introduce our method for joint nearest neighbor search and rotational alignment.

Nearest neighbor search: Based on the extended and normalized mapping $i \mapsto \frac{\hat{V}_t(i)}{\|\hat{V}_t(i)\|}$, we define the multi-frequency vector diffusion distance $d_{\text{MFVDM},t}(i, j)$ between node i and j as

$$\begin{aligned} d_{\text{MFVDM},t}^2(i, j) &= \left\| \frac{\hat{V}_t(i)}{\|\hat{V}_t(i)\|} - \frac{\hat{V}_t(j)}{\|\hat{V}_t(j)\|} \right\|_2^2 \\ &= 2 - 2 \left\langle \frac{\hat{V}_t(i)}{\|\hat{V}_t(i)\|}, \frac{\hat{V}_t(j)}{\|\hat{V}_t(j)\|} \right\rangle = 2 - 2N_t(i, j), \end{aligned} \quad (16)$$

which is the Euclidean distance between mappings of i and




Figure 2. Illustration of MFVDM rotational alignment. Solid lines indicate the local frames at node i and dashed lines at node j .

j . We define the nearest neighbor for a node i to be the node j with smallest $d_{\text{MFVDM},t}^2(i, j)$. Similarly, for VDM and DM, we define the distances $d_{\text{VDM},t}$ and $d_{\text{DM},t}$, and perform the nearest neighbor search accordingly.

Rotational alignment: We notice that the eigenvectors of S_k encode the alignment information between neighboring nodes, as illustrated in Fig. 2. Assume that two nodes i and j are located at the same base manifold point, for example, the same point on S^2 , but their tangent bundle frames are oriented differently, with an in-plane rotational angle α_{ij} . Then the corresponding entries of the eigenvectors are vectors in the complex plane and the following holds,

$$u_l^{(k)}(i) = e^{ik\alpha_{ij}} u_l^{(k)}(j), \quad \forall l = 1, 2, \dots, n. \quad (17)$$

When i and j are close but not identical, (17) holds approximately. Recalling Remark 1, due to the existence of noise, for each frequency k we approximate the alignment $e^{ik\alpha_{ij}}$ using only top m_k eigenvectors. We then use weighted least squares to estimate α_{ij} , which can be written as the following optimization problem:

$$\begin{aligned} \hat{\alpha}_{ij} &= \arg \min_{\alpha} \sum_{k=1}^{k_{\max}} \sum_{l=1}^{m_k} \left(\lambda_l^{(k)} \right)^{2t} \left| u_l^{(k)}(i) - e^{ik\alpha} u_l^{(k)}(j) \right|^2 \\ &= \arg \max_{\alpha} \sum_{k=1}^{k_{\max}} \left(\sum_{l=1}^{m_k} \left(\lambda_l^{(k)} \right)^{2t} u_l^{(k)}(i) \overline{u_l^{(k)}(j)} \right) e^{-ik\alpha} \\ &= \arg \max_{\alpha} \sum_{k=1}^{k_{\max}} S_k^{2t}(i, j) e^{-ik\alpha}. \end{aligned} \quad (18)$$

To solve this, we define a sequence z and set $z(k)$ for $k = 1, 2, \dots, k_{\max}$ to be

$$z(k) = S_k^{2t}(i, j) = \sum_{l=1}^{m_k} \left(\lambda_l^{(k)} \right)^{2t} u_l^{(k)}(i) \overline{u_l^{(k)}(j)}. \quad (19)$$

According to (19) and (18), the alignment angles $\hat{\alpha}_{ij}$ can be efficiently estimated by using an FFT on zero-padded z and identifying its peak. Due to usage of multiple unitary irreducible representations of $\text{SO}(2)$, this approximation is more accurate and robust to noise than VDM. The improvement of the alignment estimation using higher order trigonometric moments is also observed in phase synchronization (Gao & Zhao, 2019).

Algorithm 1 Joint nearest neighbor search and alignment

Input: Initial noisy nearest neighbor graph $G = (V, E)$ and the corresponding edge weights $w_{ij} e^{ik\alpha_{ij}}$ defined on the edges, truncation cutoff m_k for $k = 1, \dots, k_{\max}$

Output: κ -nearest neighbors for each data point and the corresponding alignments $\hat{\alpha}_{ij}$

- 1 **for** $k = 1, \dots, k_{\max}$ **do**
 - 2 Construct the normalized affinity matrix W_k and S_k according to (4) and (6)
 - Compute the largest m_k eigenvalues $\lambda_1^{(k)} \geq \lambda_2^{(k)} \geq \dots \geq \lambda_{m_k}^{(k)}$ of S_k and the corresponding eigenvectors $\{u_l^{(k)}\}_{l=1}^{m_k}$
 - Compute the truncated frequency- k embedding $\hat{V}_t^{(k)}$ according to (11)
 - 3 **end**
 - 4 Concatenate the truncated embedding $\{\hat{V}_t^{(k)}\}_{k=1}^{k_{\max}}$, compute the normalized affinity by (15)
 - Identify κ nearest neighbors for each data point
 - Compute $\hat{\alpha}_{ij}$ for nearest neighbor pairs using (18).
-

Computational complexity: Our joint nearest neighbor search and alignment algorithm is summarized in Alg. 1. The computational complexity is dominated by the eigen-decomposition: Computing the top m_k eigenvectors of the sparse Hermitian matrices S_k , for $k = 1, \dots, k_{\max}$ requires $O(\sum_{k=1}^{k_{\max}} n(m_k^2 + m_k l))$, where l is the average number of non-zero elements in each row of S_k (e.g. number of nearest neighbors). If we assume to use an identical truncation m (i.e., $m_k = m$ for all k), and express the above in terms of the mapping dimension $d = k_{\max} m^2$, then the complexity is $O(n(d + l\sqrt{k_{\max} d}))$. For large d and moderate k_{\max} , the dominant term is $O(nd)$, therefore MFVDM and VDM ($k_{\max} = 1$) could have similar computational complexity for generating the mapping. Moreover, MFVDM can be faster by parallelizing for each frequency k . Next, searching for κ -nearest neighbors takes $O(n\kappa d \log n)$ flops. The alignment step requires FFT of zero-padded z of length T , therefore identifying the alignments takes $O(n\kappa(k_{\max} m + T \log T))$ or $O(n\kappa(\sqrt{k_{\max} d} + T \log T))$.

4. Analysis

We use a probabilistic model to illustrate the noise robustness of our embedding using the top eigenvectors and eigenvalues of W_k 's. We start with the clean neighborhood graph, i.e. $(i, j) \in E$ if i is among j 's κ -nearest neighbors or j is among i 's κ -nearest neighbors according to the \mathcal{G} -invariant distances. We construct a noisy graph based on the following process starting from the existing clean graph edges: with probability p , the distance d_{ij} is still small and we keep the edge between i and j . With probability $1 - p$ we remove the edge (i, j) and link i to a random vertex, drawn uniformly at random from the remaining vertices that are not already connected to i . We assume that if the link between i and j is a random link, then the optimal alignment α_{ij} is uniformly distributed over $[0, 2\pi)$. Our model assumes that the underlying graph of links between noisy data points is a small-world graph (Watts & Strogatz, 1998) on the mani-

fold, with edges being randomly rewired with probability $1 - p$. The alignments take their correct values for true links and random values for the rewired edges. The parameter p controls the signal to noise ratio of the graph connection where $p = 1$ indicates the clean graph.

The matrix W_k is a random matrix under this model. Since the expected value of the random variable $e^{ik\theta}$ vanishes for $\theta \sim \text{Uniform}[0, 2\pi)$, the expected value of the matrix W_k is

$$\mathbb{E}W_k = pW_k^{\text{clean}}, \quad (20)$$

where W_k^{clean} is the clean matrix that corresponds to $p = 1$ obtained in the case that all links and angles are set up correctly. At a single frequency k , the matrix W_k can be decomposed into

$$W_k = pW_k^{\text{clean}} + R_k, \quad (21)$$

where R_k is a random matrix whose elements are independent and identically distributed (i.i.d) zero mean random variables with finite moments, since the elements of R_k are bounded for $1 \leq k \leq k_{\max}$. The top eigenvectors of W_k approximate the top eigenvectors of W_k^{clean} as long as the 2-norm of R_k is not too large. Various bounds on the spectral norm of random sparse matrices are proven in (Khorunzhy, 2001; Khorunzhiy, 2003). This ensures the noise robustness for each frequency- k -VDM. Combining an ensemble of classifiers is able to boost the performance (Zhou, 2012). Across different frequencies, the entries R_k are dependent through the relations of the irreducible representations. We will provide detailed analysis across frequency channels in the future.

Spectral properties for SO(3): Related to the application in cryo-EM image analysis, we assume that the data points x_i are uniformly distributed over SO(3) according to the Haar measure. The base manifold characterized by the viewing directions v_i 's is a unit two sphere S^2 and the pairwise alignment group is SO(2). Then $e^{ik\alpha_{ij}}$ approximates the local parallel transport operator from $T_{v_j}S^2$ to $T_{v_i}S^2$, whenever x_i and x_j have similar viewing directions v_i and v_j that satisfy $\langle v_i, v_j \rangle \geq 1 - h$, where h characterizes the size of the small spherical cap of the neighborhood. The matrices W_k^{clean} approximate the local parallel transport operators $P_h^{(k)}$, which are integral operators over SO(3). We have the following spectral properties for the integral operators,

Theorem 1 *The operator $P_h^{(k)}$ has a discrete spectrum $\lambda_l^{(k)}(h)$, $l \in \mathbb{N}$, with multiplicities equal to $2(l+k) - 1$, for every $h \in (0, 2]$. Moreover, in the regime $h \ll 1$, the eigenvalue $\lambda_l^{(k)}(h)$ has the asymptotic expansion*

$$\lambda_l^{(k)}(h) = \frac{1}{2}h - \frac{k + (l-1)(l+2k)}{8}h^2 + O(h^3). \quad (22)$$

The proof of Theorem 1 is detailed in the Appendix A.1 of (Gao et al., 2019b). Each eigenvalue $\lambda_l^{(k)}(h)$, as a function of h , is a polynomial of degree $l+k$. This extends Theorem 3 in (Hadani & Singer, 2011) to frequencies $k > 1$.

The multiplicities of the eigenvalues can be seen in the last column of Fig. 3 and Fig. 11. A direct consequence of Theorem 1 is that the top spectral gap of $P_h^{(k)}$ for small $h > 0$ can be explicitly obtained. When $h \ll 1$, the top spectral gap is $G^{(k)}(h) \approx \frac{1+k}{4}h^2$, which increases with the angular frequency. If we use top $m_k = 2k + 1$ eigenvectors for the frequency- k -VDM, then from a perturbation analysis perspective, it is well known (see e.g. (Rohe et al., 2011; Eldridge et al., 2018; Fan et al., 2018) and the references therein) that the stability of the eigenmaps essentially depends on the top spectral gap. Therefore, we are able to jointly achieve more robust embedding and nearest neighbor search under high level of noise or a large number of outliers. Moreover, we are not restricted to use only top $2k + 1$ eigenvectors and incorporating more eigenvectors can improve the results (Singer et al., 2011).

5. Experiments

5.1. Synthetic examples on 2 dimensional sphere and torus

We test MFVDM on two synthetic examples: 2-D sphere S^2 and torus T^2 . For the first example, we simulate $n = 10^4$ points x_i uniformly distributed over SO(3) according to the Haar measure. Each x_i can be represented by a 3×3 orthogonal matrix R_i whose determinant is equal to 1. The third column of the rotation matrices R_i (denoted as v_i) forms a point on the manifold S^2 ,

$$S^2 = \{v \in \mathbb{R}^3 : \|v\| = 1\}. \quad (23)$$

The pairwise alignment α_{ij} is computed based on (2). The hairy ball theorem (Milnor, 1978) says that a continuous tangent vector field to the two dimensional sphere must vanish at some points on the sphere, therefore, we cannot identify $\alpha_i \in [0, 2\pi)$ for $i = 1, \dots, n$, such that $\alpha_{ij} = \alpha_i - \alpha_j$, for all i and j . As a result, we cannot globally align the tangent vectors. For the torus, we sample $n = 10^4$ points uniformly distributed on the manifold, which are embedded in three dimensional space according to,

$$T^2 = \begin{cases} x = (R + r \cos u) \cos v, \\ y = (R + r \cos u) \sin v, \\ z = r \sin u, \end{cases} \quad (24)$$

where $R = 1$, $r = 0.2$ and $(u, v) \in [0, 2\pi) \cup [0, 2\pi)$, and for each node i we assign an angle α_i that is uniformly distributed in $[0, 2\pi)$, due to the existence of a continuous vector field, we set the pairwise alignment $\alpha_{ij} = \alpha_i - \alpha_j$. For both examples, we connect each node with its top 150 nearest neighbors based on their geodesic distances on the base manifold, then noise is added on edges following the random graph model described in Sec. 4 with parameter p . Finally, we build the affinity matrix W_k by setting weights $w_{ij} \equiv 1 \forall (i, j) \in E$, with $k = 1, 2, \dots, k_{\max}$.

Parameter setting: For MFVDM, we set the maximum frequency $k_{\max} = 50$ and for each k , we select top $m_k = 50$ eigenvectors. For VDM and DM, we set the number of eigenvectors to be $m = 50$. In addition, we set random walk




Figure 3. S^2 case: Bar plots of the 30 smallest eigenvalues $1 - \lambda^{(k)}$ of the graph connection Laplacian $I - S_k$ on S^2 for different p 's and k 's.




Figure 4. S^2 case: The normalized $d_{MFVDM,t}$, $d_{VDM,t}$, and $d_{DM,t}$ between a reference point (marked in red) and other points, with $t = 1, 10, \text{ and } 100, p = 1$.

step size $t = 1$.

Spectral property on S^2 : We numerically verify the spectrum of graph connection Laplacian $I - S_k$ on S^2 for different k and random rewiring parameter p . Smaller p indicates more edges are corrupted by noise. Fig. 3 shows that the multiplicities of S_k (normalized W_k matrix) agree with Theorem 1. The spectral gaps persist even when 80% of the edges are corrupted (see the right column of Fig. 3).

Multi-frequency vector diffusion distances on S^2 : Based on (16), Fig. 4 displays the normalized and truncated multi-frequency vector diffusion distances $d_{MFVDM,t}^2(i, j)$, vector diffusion distances $d_{VDM,t}^2(i, j)$, and diffusion distances $d_{DM,t}^2(i, j)$ between a reference point (marked in red) and others, on S^2 at $p = 1$ (clean graph). Moreover, we increase the diffusion step size t from $t = 1$ to $t = 10$ and 100 . In this clean case, all three distances are highly correlated to the geodesic distance. Specifically, MFVDM and VDM perform similarly.

To demonstrate the robustness to noise of $d_{MFVDM,t}$, we compare $d_{MFVDM,t}$, $d_{VDM,t}$, and $d_{DM,t}$ against the geodesic distance on S^2 in Fig. 5 at different noise levels. When $p = 1$, all the distances are highly correlated with the geodesic distance, e.g., small $d_{MFVDM,t}$, $d_{VDM,t}$, and $d_{DM,t}$ all corre-




Figure 5. S^2 case: Scatter plots comparing the normalized $d_{MFVDM,t}$, $d_{VDM,t}$, and $d_{DM,t}$ at $p = 0.2, 0.4, \text{ and } 1$.

spond to small geodesic distance. However at high noise level as $p = 0.4$ or 0.2 , both $d_{VDM,t}$ and $d_{DM,t}$ become more scattered, while $d_{MFVDM,t}$ remains correlated with the geodesic distance. Here the random walk steps $t = 10$ and the results are similar for $t = 1$ or 100 .

Nearest neighbor search and rotational alignment: We test the nearest neighbor search (NN search) and rotational alignment results on both sphere and torus, with different noise levels p . As mentioned, one advantage of MFVDM is its robustness to noise. Even at a high noise level, the true affinity between nearest neighbors can still be preserved. In our experiments, for each node we identify its $\kappa = 50$ nearest neighbors.

We evaluate the NN search by the geodesic distance between each node and its nearest neighbors. A better method should find more neighbors with geodesic distance close to 0. In the top rows of Fig. 6 and Fig. 7 we show the histograms of such geodesic distance. Note that in the low noise regime ($p \geq 0.2$), MFVDM, VDM and DM all perform well and MFVDM is slightly better. When the noise level increases to $p = 0.1$, both VDM and DM have poor result while MFVDM still works well. These comparisons show MFVDM, which benefits from multiple irreducible representations, is very robust to noise.

We evaluate the rotational alignment estimation by computing the alignment errors $\alpha_{ij} - \hat{\alpha}_{ij}$ for all pairs of nearest neighbors (i, j) , where α_{ij} is the ground truth and $\hat{\alpha}_{ij}$ is the estimation. In the bottom rows of Fig. 6 and Fig. 7, we show the histograms of such alignment errors. The results demonstrate that for a wide range of p , i.e., $p \geq 0.1$, the MFVDM alignment errors are closer to 0 than the baseline VDM. At $p = 0.08$, the VDM errors disperse between 0 to 180 degrees, whereas a large number of the alignment errors of MFVDM are still close to 0.

At each frequency k , we individually perform NN search based on frequency- k -VDM and the corresponding affinity in (12). For the S^2 example, we find that all single frequency mappings achieve similar accuracies when m_k 's are identical (see Fig. 8). MFVDM combines those weak single

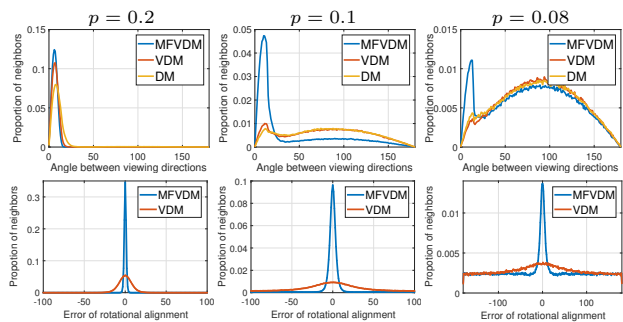


Figure 6. S^2 case: *Top*: histograms of the viewing direction difference between nearest neighbors found by MFVDM, VDM and DM; *Bottom*: the accuracy of the rotational alignment estimated by MFVDM and VDM.

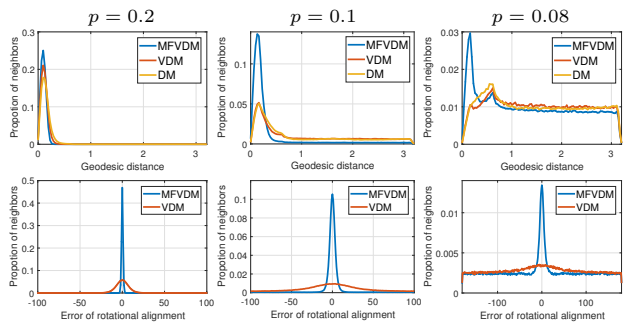


Figure 7. T^2 case: *Top*: histograms of the geodesic distances between nearest neighbors identified by MFVDM, VDM and DM; *Bottom*: the accuracy of the rotational alignment estimated by MFVDM and VDM.




Figure 8. S^2 case: *Weak classifier versus strong classifier*: histograms of the angles between nearest neighbors found by using single frequency- k -VDM (weak classifier) and MFVDM (all) with $k = 1, \dots, k_{\max}$ (strong classifier, shown as ‘all’). Here $k_{\max} = 10$.

frequency classifiers into a strong classifier to boost the accuracy of nearest neighbor search

Choice of parameters: The performance of MFVDM depends on two parameters: the maximum frequency cutoff k_{\max} and the number of top eigenvectors m_k . We assume that m_k ’s are the same for all frequencies, that is $m_1 = m_2 = \dots = m_{k_{\max}} = m_k$. In the top row of Fig. 9, we show the average geodesic distances between the nearest neighbor pairs identified by MFVDM, with different values of k_{\max} and m_k . First, we fix $m_k = 50$ and vary k_{\max} . The performance of MFVDM improves with increasing k_{\max} and plateaus when k_{\max} approaches 50 (see the upper left panel of Fig. 9). Then we fix $k_{\max} = 10$ and vary m_k . The upper right panel of Fig. 9 shows that choosing $m_k = 50$ achieves the best performance. Using a larger number of eigenvectors, i.e. $m_k = 100$, does not lead to higher accu-




Figure 9. S^2 case: Nearest neighbor search by MFVDM, VDM and DM under varying parameters: maximum frequency k_{\max} and the number of eigenvectors m_k . *Upper left*: MFVDM with varying k_{\max} and $m_k = 50$; *Upper right*: MFVDM with varying m_k and $k_{\max} = 10$; *Lower left*: VDM with varying m ($k_{\max} = 1$); *Lower right*: DM with varying m . Horizontal axis: the value of the parameter p in the random graph model, lower p means larger number of outliers in the edge connections. Vertical axis: the average geodesic distances of nearest neighbors pairs (lower is better).

racy in nearest neighbor search, because the eigenvectors of S_k with small eigenvalues are more sensitive to noise and including them will reduce the robustness to noise of the mappings. In addition, we evaluate the performance of VDM and DM under varying number of eigenvectors m in the bottom row of Fig. 9. VDM and DM also achieve the best performance at $m = 50$. Comparing the upper left and lower left panels of Fig. 9, we find that MFVDM greatly improves the nearest neighbor search accuracy of VDM when 90% of the true edges are rewired. Note that the solid blue line in the upper left panel of Fig. 9 corresponds to the best performance curve in the lower left panel of Fig. 9 (green line with $m = 50$).

5.2. Application: Cryo-EM 2-D image analysis

MFVDM is motivated by the cryo-EM 2-D class averaging problem. In the experiments, protein samples are frozen in a very thin ice layer. Each image is a tomographic projection of the protein density map at an unknown random orientation. It is associated with a 3×3 rotation matrix R_i , where the third column of R_i indicates the projection direction v_i , which can be realized by a point on S^2 . Projection images I_i and I_j that share the same views look the same up to some in-plane rotation. The goal is to identify images with similar views, then perform local rotational alignment and averaging to denoise the image. Therefore, MFVDM is suitable to perform the nearest neighbor search and rotational alignment estimation.

In our experiment, we simulate $n = 10^4$ projection images from a 3-D electron density map of the 70S ribosome (see Fig. 10), the orientations for the projection images are uniformly distributed over $SO(3)$ and the images are




Figure 10. Cryo-EM 2-D image analysis: *Left*: Reference volume of 70S ribosome; *Mid*: Clean projection images; *Right*: Noisy projection images at SNR= 0.05.




Figure 11. Bar plots of the 30 smallest eigenvalues of the graph connection Laplacian $I - S_k$ that build upon the initial NN search and alignment results on cryo-EM images (MFVDM) and the corresponding eigenvalues of the steerable graph Laplacian (SGL) in (Landa & Shkolnisky, 2018).

contaminated by additive white Gaussian noise at signal to noise ratio (SNR) equal to 0.05. Note such high noise level is commonly observed in real experiments. In Fig. 10, we display samples of such clean and noisy images. As pre-processing, we use fast steerable PCA (sPCA) (Zhao et al., 2016) and rotationally invariant features (Zhao & Singer, 2014) to initially identify the images of similar views and the in-plane rotational alignment angles according to (Zhao & Singer, 2014). Then we take the initial graph structure and estimation of the optimal alignments as the input of MFVDM.

As a comparison to (4), we introduce another kernel introduced in steerable graph Laplacian (SGL) (Landa & Shkolnisky, 2018), which is defined on images including all rotated versions. Then similar to the synthetic examples, in Fig. 11 we present the spectrum of the graph connection Laplacian $I - S_k$. The spectral gap clearly exists at both clean and noisy cases. Fig. 12 shows the NN search and rotational alignment results. We set $t = 10$, $k_{\max} = 10$, $m_k = 10$, and $m = 10$ for MFVDM, VDM, and DM respectively. Although using SGL kernel achieves slightly better NN search, its performance on alignment estimation is worse than MFVDM.

6. Discussion

In the current probabilistic model, we only consider independent edge noise, i.e., the entries in R_k for a fixed k are independent. This does not cover the measurement scenarios in some applications. For example, in cryo-EM 2-D image analysis, each image is corrupted by independent noise. Therefore, the entries in R_k become dependent since the edge connections and alignments are affected by the noise in each image node. Empirically, our new algorithm is still applicable and results in the improved nearest




Figure 12. Nearest neighbor search and rotational alignment results for simulated cryo-EM images of 70S ribosome with SNR = 0.05. *Left*: distribution of the angles between nearest neighbors; *Right*: Rotational alignment accuracy. See text for the methods description.

neighbor search and rotational alignment estimation compared to the state-of-the-art VDM. We leave the analysis of node level noise to future work. In addition, there can be other approaches to define the multi-frequency mapping, such as weighted average among different frequencies or majority voting. We will explore other ways to integrate multi-frequency information in the future.

The current analysis focuses on data points that are uniformly distributed on the manifold. For non-uniformly distributed points, different normalization techniques introduced in DM(Coifman & Lafon, 2006) and (Zelnik-Manor & Perona, 2005) are needed to compensate the sampling density.

Since our framework is motivated by the cryo-EM nearest neighbor image search and alignment, we have so far only considered the compact manifold \mathcal{M} where the intrinsic dimension is 2 and the local parallel transport operator can be well approximated by the in-plane rotational alignment of the images or the alignment of the local tangent bundles as discussed in VDM (Singer & Wu, 2012). In the future, we will extend the current algorithm to manifolds with higher intrinsic dimension and other compact group alignments $g \in \mathcal{G}$ with their corresponding irreducible representations $\rho_k(g)$, for example, the symmetric group which is widely used in computer vision (Bajaj et al., 2018).

7. Conclusion

In this paper, we have introduced MFVDM for joint nearest neighbor search and rotational alignment estimation. The key idea is to extend VDM using multiple irreducible representations of the compact Lie group. Enforcing the consistency of the rotational transformations at different frequencies allows us to achieve better nearest neighbor identification and accurately estimate the alignments between the updated nearest neighbor pairs. The approach is based on spectral decomposition of multiple kernel matrices. We use the random matrix theory and the rationale of ensemble methods to justify the robustness of MFVDM. Experimental results show efficacy of our approach compared to the state-of-the-art methods. This general framework can be applied to many other problems, such as joint synchronization and clustering (Gao et al., 2019a) and multi-frame alignment in computer vision.

References

- Bajaj, C., Gao, T., He, Z., Huang, Q., and Liang, Z. SMAC: Simultaneous mapping and clustering using spectral decompositions. In *International Conference on Machine Learning*, pp. 334–343, 2018.
- Belkin, M. and Niyogi, P. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *NIPS*, 2002.
- Belkin, M. and Niyogi, P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 2003.
- Belkin, M., Niyogi, P., and Sindhvani, V. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of machine learning research*, 7(Nov):2399–2434, 2006.
- Coifman, R. R. and Lafon, S. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.
- Dashti, A., Schwander, P., Langlois, R., Fung, R., Li, W., Hosseinizadeh, A., Liao, H. Y., Pallesen, J., Sharma, G., Stupina, V. A., et al. Trajectories of the ribosome as a brownian nanomachine. *Proceedings of the National Academy of Sciences*, 111(49):17492–17497, 2014.
- Donoho, D. L. and Grimes, C. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003.
- Eldridge, J., Belkin, M., and Wang, Y. Unperturbed: spectral analysis beyond Davis-Kahan. In *Algorithmic Learning Theory*, pp. 321–358, 2018.
- Fan, J., Wang, W., and Zhong, Y. An ℓ_∞ eigenvector perturbation bound and its application. *Journal of Machine Learning Research*, 18(207):1–42, 2018.
- Gao, T. and Zhao, Z. Multi-frequency phase synchronization. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2132–2141, 2019.
- Gao, T., Brodzki, J., and Mukherjee, S. The geometry of synchronization problems and learning group actions. *Discrete & Computational Geometry*, 2019a.
- Gao, T., Fan, Y., and Zhao, Z. Representation theoretic patterns in multi-frequency class averaging for three-dimensional cryo-electron microscopy. *arXiv preprint arXiv:1906.01082*, 2019b.
- Giannakis, D., Schwander, P., and Ourmazd, A. The symmetries of image formation by scattering. I. theoretical framework. *Optics express*, 20(12):12799–12826, 2012.
- Goldberg, A., Zhu, X., Singh, A., Xu, Z., and Nowak, R. Multi-manifold semi-supervised learning. In *Artificial Intelligence and Statistics*, pp. 169–176, 2009.
- Gong, D., Sha, F., and Medioni, G. Locally linear denoising on image manifolds. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 265–272, 2010.
- Hadani, R. and Singer, A. Representation theoretic patterns in three-dimensional cryo-electron microscopy II: the class averaging problem. *Foundations of Computational Mathematics*, 11(5):589–616, 2011.
- Khorunzhiy, O. Rooted trees and moments of large sparse random matrices. In *Discrete Mathematics and Theoretical Computer Science*, pp. 145–154. Discrete Mathematics and Theoretical Computer Science, 2003.
- Khorunzhy, A. Sparse random matrices: spectral edge and statistics of rooted trees. *Advances in Applied Probability*, 33(1):124–140, 2001.
- Landa, B. and Shkolnisky, Y. The steerable graph laplacian and its application to filtering image datasets. *SIAM Journal on Imaging Sciences*, 11(4):2254–2304, 2018.
- Milnor, J. Analytic proofs of the “hairy ball theorem” and the brouwer fixed point theorem. *The American Mathematical Monthly*, 85(7):521–524, 1978.
- Nadler, B., Lafon, S., Kevrekidis, I., and Coifman, R. R. Diffusion maps, spectral clustering and eigenfunctions of Fokker-Planck operators. In *Advances in neural information processing systems*, pp. 955–962, 2006.
- Rohe, K., Chatterjee, S., and Yu, B. Spectral clustering and the high-dimensional stochastic block model. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- Roweis, S. T. and Saul, L. K. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- Schwander, P., Giannakis, D., Yoon, C. H., and Ourmazd, A. The symmetries of image formation by scattering. II. applications. *Optics express*, 20(12):12827–12849, 2012.
- Singer, A. and Wu, H.-T. Vector Diffusion Maps and the Connection Laplacian. *Communications on Pure and Applied Mathematics*, 65(8):1067–1144, 2012.
- Singer, A., Shkolnisky, Y., and Nadler, B. Diffusion interpretation of nonlocal neighborhood filters for signal denoising. *SIAM Journal on Imaging Sciences*, 2(1):118–139, 2009.
- Singer, A., Zhao, Z., Shkolnisky, Y., and Hadani, R. Viewing angle classification of cryo-electron microscopy images using eigenvectors. *SIAM Journal on Imaging Sciences*, 4(2):723–759, 2011.
- Tenenbaum, J. B., De Silva, V., and Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *Science*, 2000.

- Von Luxburg, U. A tutorial on spectral clustering. *Statistics and computing*, 2007.
- Watts, D. J. and Strogatz, S. H. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440, 1998.
- Yang, Z., Cohen, W., and Salakhudinov, R. Revisiting semi-supervised learning with graph embeddings. In *International Conference on Machine Learning*, pp. 40–48, 2016.
- Zelnik-Manor, L. and Perona, P. Self-tuning spectral clustering. In *Advances in neural information processing systems*, pp. 1601–1608, 2005.
- Zhao, Z. and Singer, A. Rotationally invariant image representation for viewing direction classification in cryo-EM. *Journal of structural biology*, 186(1):153–166, 2014.
- Zhao, Z., Shkolnisky, Y., and Singer, A. Fast steerable principal component analysis. *IEEE transactions on computational imaging*, 2(1):1–12, 2016.
- Zhou, Z.-H. *Ensemble methods: foundations and algorithms*. Chapman and Hall/CRC, 2012.
- Zhu, X. Semi-supervised learning literature survey. *Computer Science, University of Wisconsin-Madison*, 2(3):4, 2006.