Connection graph Laplacian methods can be made robust to noise

Noured dine El Karoui * and Hau-tieng Wu †

May 18th, 2014

Abstract

Recently, several data analytic techniques based on connection graph laplacian (CGL) ideas have appeared in the literature. At this point, the properties of these methods are starting to be understood in the setting where the data is observed without noise. We study the impact of additive noise on these methods, and show that they are remarkably robust. As a by-product of our analysis, we propose modifications of the standard algorithms that increase their robustness to noise. We illustrate our results in numerical simulations.

1 Introduction

In the last few years, several interesting variants of kernel-based spectral methods have arisen in the applied mathematics literature. These ideas appeared in connection with new types of data, where pairs of objects or measurements of interest have a relationship that is "blurred" by the action of a nuisance parameter. More specifically, we can find this type of data in a wide range of problems, for instance in the class averaging algorithm for the cryo-electron microscope (cryo-EM) problem [62, 71], in a modern light source imaging technique known as ptychography [45], in graph realization problems [24, 25], in vectored PageRank [20], in multi-channels image processing [5], etc...

Before we give further details about the cryo-EM problem, let us present the main building blocks of the methods we will study. They depend on the following three components:

- 1. an undirected graph G = (V, E) which describes all observations. The observations are the vertices of the graph G, denoted as $\{V_i\}_{i=1}^n$.
- 2. an affinity function $w : \mathbf{E} \to \mathbb{R}_+$, satisfying $w_{i,j} = w_{j,i}$, which describes how close two observations are (*i* and *j* index our observations). One common choice of $w_{i,j} = w(V_i, V_j)$ is of the form $w_{i,j} = \exp(-m(V_i, V_j)^2/\epsilon)$, where m(x, y) is a metric measuring how far *x* and *y* are.
- 3. a connection function $r : E \to G$, where G is a Lie group, which describes how two samples are related. In its application to the cryo-EM problem, $r_{i,j}$'s can be thought of estimates of our nuisance parameters, which are orthogonal matrices.

These three components form the connection graph associated with the data, which is denoted as (G, w, r). They can be either given to the data analyst or have to be estimated from the data, depending on the application.

This fact leads to different connection graph models and their associated noise models. For example, in the cryo-EM problem, all components of the connection graph (G, w, r) are determined from the given projection images, where each vertex represents an image [31, Appendix A]; in the ptychography problem [45], G is given by the experimenter, r is established from the experimental setup, and w is built up from the diffractive images collected in the experiment. Depending on applications, different metrics, deformations

Keywords: concentration of measure, random matrices, vector diffusion maps, spectral geometry, kernel methods; AMS MSC 2010 Classification: 60F99, 53A99

^{*}UC Berkeley, Department of Statistics; **Contact:** nkaroui@berkeley.edu ; Support from NSF grant DMS-0847647 (CA-REER) is gratefully acknowledged.

[†]Stanford University, Department of Mathematics; **Contact:** hauwu@stanford.edu ; Support from AFOSR grant FA9550-09-1-0643 is gratefully acknowledged.

or connections among pairs of observations are considered, or estimated from the dataset, to present the local information among data (see, for example, [1, 4, 15, 23, 47, 64, 66, 68]).

In this paper, since we focus on the connection graph Laplacian (CGL), we take the Lie Group $G = O(k)^1$, where $k \in \mathbb{N}$, and assume that r satisfies $r_{i,j} = r_{j,i}^{-1}$. Our primary focus in this paper is on k = 1 and k = 2.

We now give more specifics about one of the problems motivating our investigation.

Cryo-EM problem In the cryo-EM problem, the experimenter collects 2-dimensional projection images of a 3-dimensional macro-molecular object of interest, and the goal is to reconstruct the 3-dimensional geometric structure of the macro-molecular object from these projection images. Mathematically, the collected images $\mathcal{X}_{cryoEM} := \{I_i\}_{i=1}^N \in \mathbb{R}^{m^2}$ can be modeled as the X-ray transform of the potential of the macro-molecular object of interest, denoted as $\psi : \mathbb{R}^3 \to \mathbb{R}_+$. More precisely, in the setting that is usually studied, we have $I_i = X_{\psi}(R_i)$, where $R_i \in SO(3)$, SO(3) is the 3-dimensional special orthogonal group, X_{ψ} is the X-ray transform of ψ . The X-ray transform $X_{\psi}(R_i)$ is a function from \mathbb{R}^2 to \mathbb{R}_+ and hence can be treated by the data analyst as an image. We refer the reader to [31, Appendix A] for precise mathematical details. (For the rest of the discussion, we write $R_i = [R_i^1 R_i^2 R_i^3]$ in the canonical basis, where R_i^k are three dimensional unit vectors.)

The experimental process produces data with high level of noise. Therefore, to solve this inverse problem, it is a common consensus to preprocess the images to increase the signal-to-noise ratio (SNR) before sending them to the cryo-EM data analytic pipeline. An efficient way to do so is to estimate the projection directions of these images, i.e R_i^3 . This direction plays a particular role in the X-ray transform, which is different from the other two directions. If R_i^3 's were known, we would cluster the images according to these vectors and for instance take the mean of all properly rotationally aligned images to increase the SNR of the projection images - more on this below - in a cluster as a starting point for data-analysis. With these "improved" images, we can proceed to estimate R_i for the *i*-th image by applying, for example, the common line algorithm [37], so that the 3-D image can be reconstructed by the inverse X-ray transform[33]. We note that R_i^3 is a unit vector in \mathbb{R}^3 and hence lives on the standard sphere S^2 .

Conceptually, the problem is rendered difficult by the fact that the X-ray transform $X_{\psi}(R_i)$ is equivariant under the action of rotations that leave R_i^3 unchanged. In other words, if r_{θ} is an in-plane rotation, i.e a rotation that leaves R_i^3 unchanged but rotates R_i^1 and R_i^2 by an angle θ , the image $X_{\psi}(r_{\theta}R_i)$ is $X_{\psi}(R_i)$ rotated by the angle θ . In other words, $X_{\psi}(r_{\theta}R_i) = r_2(\theta)X_{\psi}(R_i)$, where $r_2(\theta)$ stands for the 2-dimensional rotation by the angle θ . These in-plane rotations are clearly nuisance parameters if we want to evaluate the projection direction R_i^3 .

To measure the distance between R_i^3 and R_j^3 , we hence use a rotationally invariant distance, i.e $d_{i,j}^2 = \inf_{\theta \in [0,2\pi]} \|P_i - r_2(\theta)P_j\|_2^2$. In other words, we look at the Euclidian distance between our two X-ray transforms/images after we have "aligned" them as best as possible. We now think of R_i^3 's - the vectors we would like to estimate - as elements of the manifold S^2 , equipped with a metric \mathbf{g}_{ψ} , which depends on the macro-molecular object of interest. It turns out that the Vector Diffusion Maps algorithm (VDM), which is based on CGL and which we study in this paper, is effective in producing a good approximation of \mathbf{g}_{ψ} from the local information $d_{i,j}$'s and the rotations we obtain by aligning the various X-ray transforms. This in turns imply better clustering of the R_i^3 's and improvement in the data-analytic pipeline for cryo-EM problems [62, 71].

The point of this paper is to understand how the CGL algorithms perform when the input data is corrupted by noise. The relationship between this method and the connection concept in differential geometry is the following: the projection images P_i form a graph, and we can define the affinity and connection among a pair of images so that the topological structure of the 2-dimensional sphere (S^2, \mathbf{g}_{ψ}) is encoded in the graph. This amounts to using the local geometric information derived from our data to estimate the global information - including topology - of (S^2, \mathbf{g}_{ψ}) .

Impact of noise on these problems What is missing from these considerations and the current literature is an understanding of how noise impact the procedures which are currently used and have mathematical backing in the noise-free context. The aim of our paper is to shed light on the issue of the

¹We may also consider U(k). But in this paper we focus on O(k) to simplify the discussion.

impact of noise on these interesting and practically useful procedures. We will be concerned in this paper with the impact of adding noise on the observations - collected for instance in the way described above.

Note that additive noise may have impact in all three building blocks of the connection graph associated with the data. First, it might make the graph noisy. For example, in the cryo-EM problem, the standard algorithm builds up the graph from a given noisy data set $\{P_i\}_{i=1}^n = \{I_i + \xi_i\}_{i=1}^n - \xi_i$ is our additive noise - using the nearest neighbors determined by a pre-assigned metric, that is, we add an edge between two vertices when they are close enough in that metric. Then, clearly, the existence of the noise ξ_i will likely create a different nearest neighbor graph from the the one that would be built up from the (clean) projection images $\{I_i\}_{i=1}^n$. As we will see in this paper, in some applications, it might be beneficial to consider a complete graph instead of a nearest neighbor graph.

The second noise source is how w and r are provided or determined from the samples. For example, in the cryo-EM problem, although $\{P_i\}$ are points located in a high dimensional Euclidean space, we determine the affinity and connection between two images by evaluating their rotationally invariant distance. It is clear that when P_i is noisy, the derived affinity and connection will be noisy and likely quite different from the affinity and connection we would compute from the clean dataset $\{I_i\}_{i=1}^n$. On the other hand, in the ptychography problem, the connection is directly determined from the experimental setup, so that it is a noise-free even when our observations are corrupted by additive noise.

In summary, corrupting the observations by additive noise might impact the following elements of the data analysis:

- 1. which scheme and metric we choose to construct the graph;
- 2. how we build up the affinity function;
- 3. how we build up the connection function.

More details on CGL methods

At a high-level, connection graph Laplacian (CGL) methods create a block matrix from the connection graph. The spectral properties of this matrix are then used to estimate properties of the intrinsic structure from which we posit the data is drawn from. This in turns lead to good estimation methods for, for instance, geodesic distance on the manifold, if the underlying intrinsic structure is a manifold. We refer the reader to Appendix B and references [3, 19, 20, 59, 61] for more information.

Given a $n \times n$ matrix W, with scalar entries denoted by $w_{i,j}$ and a $nk \times nk$ block matrix G with $k \times k$ block entries denoted by $G_{i,j}$, we define a $nk \times nk$ matrix S with (i, j)-block entries

$$S_{i,j} = w_{i,j}G_{i,j}$$

and a $nk \times nk$ block diagonal matrix D with (i, i)-block entries

$$D_{i,i} = \sum_{j \neq i} w_{i,j} \mathrm{Id}_k$$

which is assumed to be invertible. Let us call

$$L(W,G) := D^{-1}S$$
 and $L_0(W,G) := L(W \circ 1_{i \neq j},G).$ (1)

In other words, $L_0(W, G)$ is the matrix L(W, G) computed from the weight matrix W where the diagonal weights have been replaced by 0.

Suppose we are given a connection graph (\mathbf{G}, w, r) , and construct the $n \times n$ affinity matrix W so that $w_{i,j} = w(i,j)$ and the connection matrix G, the $nk \times nk$ block matrix with $k \times k$ block entries $G_{i,j} = r(i,j)$, the CGL associated with the connection graph (\mathbf{G}, w, r) is defined as $\mathrm{Id}_{nk} - L(W, G)$ and the modified CGL associated with the connection graph (\mathbf{G}, w, r) is defined as $\mathrm{Id}_{nk} - L_0(W, G)$. We note that under our assumptions on r, i.e $r_{i,j} = r_{i,j}^{-1} = r_{i,j}^*$ the connection matrix G is Hermitian.

We are interested in the large eigenvalues of L(W, G) (or, equivalently, the small eigenvalues of the CGL $Id_{nk} - L(W, G)$), as well as the corresponding eigenvectors. In the case where the data is not corrupted by noise, the CGL's asymptotic properties have been studied in [59, 61], when the underlying intrinsic structure is a manifold. Its so-called synchronization properties have been studied in [3, 20].

The aim of our study is to understand the impact of additive noise on CGL algorithms. Two main results are Proposition 2.1, which explains the effect of noise on the affinity, and Theorem 2.2, which explains the effect of noise on the connection. These two results lead to suggestions for modifying the standard CGL algorithms: the methods are more robust when we use a complete graph than when we use a nearestneighbor graph, the latter being commonly used in practice. One should also use the matrix $L_0(W, G)$ instead of L(W, G) to make the method more robust to noise. After we suggest these modifications, our main result is Proposition 2.3, which shows that even when the signal-to-noise-ratio (SNR) is very small, i.e going to 0 asymptotically, our modifications of the standard algorithm will approximately yield the same spectral results as if we had been working on the CGL matrix computed from noiseless data. We develop in Section 2 a theory for the impact of noise on CGL algorithms and show that our proposed modifications to the standard algorithms render them more robust to noise. We present in Section 3 some numerical results.

Notation: Here is a set of notations we use repeatedly. \mathcal{T} is a set of linear transforms. Id_k stands for the $k \times k$ identity matrix. If $v \in \mathbb{R}^n$, $D(\{v\})$ is a $nk \times nk$ block diagonal matrix with the (i, i)-th block equal to $v_i \operatorname{Id}_k$. We denote by $A \circ B$ the Hadamard, i.e entry-wise, product of the matrices A and B. $|||M|||_2$ is the largest singular value (a.k.a operator norm) of the matrix M. $||M||_F$ is its Frobenius norm.

2 Theory

Our aim in this section is to develop a theory that explains the behavior of CGL algorithms in the presence of noise. In particular, it will apply to algorithms of the cryo-EM type. We give in Subsection 2.1 approximation results that apply to general CGL problems. In Subsection 2.2, we study in details the impact of noise on both the affinity and the connection used in the computation of the CGL when using the rotationally invariant distance (this is particularly relevant for the cryo-EM problem). We put these results together for a detailed study of CGL algorithms in Subsection 2.3. We also propose in Subsection 2.3 modifications to the standard algorithms.

2.1 General approximation results

We first present a result that applies generally to CGL algorithms.

Lemma 2.1. Suppose W and \widetilde{W} are $n \times n$ matrices, with scalar entries denoted by $w_{i,j}$ and $\tilde{w}_{i,j}$ and G and \widetilde{G} are $nd \times nd$ block matrices, with $d \times d$ blocks denoted by $G_{i,j}$ and $\tilde{G}_{i,j}$. Suppose that

$$\sup_{i,j} |\tilde{w}_{i,j} - w_{i,j}| \le \epsilon , \text{ and } \sup_{i,j} ||\widetilde{G}_{i,j} - G_{i,j}||_F \le \eta ,$$

where $\epsilon, \eta \geq 0$. Suppose furthermore that there exists C > 0 such that $0 \leq w_{i,j} \leq C$, $\sup_{i,j} ||G_{i,j}||_F \leq C$ and $\sup_{i,j} ||\widetilde{G}_{i,j}||_F \leq C$. Then, if $\inf_i \sum_{j \neq i} w_{i,j}/n > \gamma$ and $\gamma > \epsilon$, we have

$$|||L(W,G) - L(\widetilde{W},\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^2$$

The proof of this lemma is given in Appendix A-2. This lemma says that if we can approximate the matrix W well entrywise and each of the individual matrices $G_{i,j}$ well, too, data analytic techniques working on the CGL matrix $L(\widetilde{W}, \widetilde{G})$ will do essentially as well as those working on the corresponding matrix for L(W, G) in the spectral sense.

This result is useful because many methods rely on these connection graph ideas, with different input in terms of affinity and connection functions [1, 15, 16, 23, 24, 25, 47, 62, 64, 66, 71]. However, it will often be the case that we can approximate $w_{i,j}$ - which we think of as measurements we would get if our signals were not corrupted by noise - only up to a constant. The following result shows that in certain situations, this will not affect dramatically the spectral properties of the CGL matrix. **Lemma 2.2.** We work under the same setup as in Lemma 2.1 and with the same notations. However, we now assume that

$$\exists \{f_i\}_{i=1}^n , f_i > 0 : \sup_{i,j} \left| \frac{\tilde{w}_{i,j}}{f_i} - w_{i,j} \right| \le \epsilon , \text{ and } \sup_{i,j} \|\widetilde{G}_{i,j} - G_{i,j}\|_F \le \eta .$$

Suppose furthermore that there exists C > 0 such that $0 \le w_{i,j} \le C$, $\sup_{i,j} ||G_{i,j}||_F \le C$ and $\sup_{i,j} ||\widetilde{G}_{i,j}||_F \le C$. Then, if $\inf_i \sum_{j \ne i} w_{i,j}/n > \gamma$ and $\gamma > \epsilon$, we have

$$|||L(W,G) - L(\widetilde{W},\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^2$$

We note that quite remarkably, there are essentially no conditions on f_i 's: in particular, $\tilde{w}_{i,j}$ and $w_{i,j}$ could be of completely different magnitudes. The previous lemma also shows that, for the purpose of understanding the large eigenvalues and eigenvectors of L(W, G), we do not need to estimate f_i 's: we can simply use $L(\widetilde{W}, \widetilde{G})$, i.e just work with the noisy data.

Proof. Let us call \widetilde{W}_f the matrix with scalar entries $\tilde{w}_{i,j}/f_i$. We note simply that

$$L(\widetilde{W}_f, \widetilde{G}) = L(\widetilde{W}, \widetilde{G})$$
.

The assumptions of Lemma 2.1 apply to $(\widetilde{W}_f, \widetilde{G})$ and hence we have

$$|||L(W,G) - L(\widetilde{W}_f,\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^2.$$

But since $L(\widetilde{W_f}, \widetilde{G}) = L(\widetilde{W}, \widetilde{G})$, we also have

$$|||L(W,G) - L(\widetilde{W},\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta+\epsilon) + \frac{\epsilon}{\gamma(\gamma-\epsilon)}C^2.$$

In some situations that will be of interest to us below, it is however, not the case that we can find f_i 's such that

$$\exists \{f_i\}_{i=1}^n , f_i > 0 : \sup_{i,j} \left| \frac{\tilde{w}_{i,j}}{f_i} - w_{i,j} \right| \le \epsilon .$$

Rather, this approximation is possible only when $i \neq j$, yielding the condition

$$\left| \forall i, \exists f_i > 0 : \sup_{i \neq j} \left| \frac{\tilde{w}_{i,j}}{f_i} - w_{i,j} \right| \le \epsilon$$

This apparently minor difference turns out to have significant consequences, both practical and theoretical. We propose in the following lemma to modify the standard way of the computing the CGL matrix to handle this more general case.

Lemma 2.3. We work under the same setup as in Lemma 2.1 and with the same notations. We now assume that multiplicative approximations of the weights is possible only on the off-diagonal elements of our weight matrix:

$$\exists \{f_i\}_{i=1}^n, f_i > 0: \sup_{i \neq j} \left| \frac{\tilde{w}_{i,j}}{f_i} - w_{i,j} \right| \le \epsilon, \text{ and } \sup_{i,j} \|\tilde{G}_{i,j} - G_{i,j}\|_F \le \eta.$$

Suppose furthermore that there exists C > 0 such that $0 \le w_{i,j} \le C$, $\sup_{i,j} ||G_{i,j}||_F \le C$, and $\sup_{i,j} ||\widetilde{G}_{i,j}||_F \le C$. Then, if $\inf_i \sum_{j \ne i} w_{i,j}/n > \gamma$ and $\gamma > \epsilon$, we have

$$|||L_0(W,G) - L_0(\widetilde{W},\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^2 ,$$

and

$$|||L(W,G) - L_0(\widetilde{W},\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^2 + \frac{C^2}{n\gamma}$$

Comment: Concretely, this lemma means that if we do not include the block diagonal terms in the computation of the CGL obtained from our "noisy data", i.e $(\widetilde{W}, \widetilde{G})$, we will get a matrix that is very close in spectral norm to the CGL computed from the "clean data", i.e (W, G). The significance of this result lies in the fact that recent work in applied mathematics has proposed to use the large eigenvalues and eigenvectors of L(W, G) for various data analytic tasks, such as the estimation of local geodesic distances when the data is thought to be sampled from an unknown manifold.

What our result shows is that even when f_i are arbitrarily large, which we can think of as the situation where the signal to noise ratio in \widetilde{W} is basically 0, working with $L_0(\widetilde{W}, \widetilde{G})$ will allow us to harness the power of these recently developed tools. Naturally, working with $(\widetilde{W}, \widetilde{G})$ is a much more realistic assumption than working with (W, G) since we expect all our measurements to be somewhat noisy whereas results based on (W, G) essentially assume that there is no measurement error in the dataset.

Proof. Recall also that the computation of the D matrix does not involve the diagonal weights. Therefore,

$$L(W,G) = L_0(W,G) + D^{-1}\Delta(\{w_{i,i}\},G_{i,i})$$

where $\Delta(\{w_{i,i}\}, G_{i,i})$ is the block diagonal matrix with (i, i) block diagonal $w_{i,i}G_{i,i}$, and $D^{-1}\Delta(\{w_{i,i}\}, G_{i,i})$ is a block diagonal matrix with (i, i) block diagonal

$$\frac{w_{i,i}}{\sum_{j \neq i} w_{i,j}} G_{i,i}$$

This implies that

$$|||L(W,G) - L_0(W,G)|||_2 \le \sup_i \frac{w_{i,i}}{\sum_{j \ne i} w_{i,j}} |||G_{i,i}|||_2$$

Our assumptions imply that $\sum_{j\neq i} w_{i,j} > \gamma n$, $\sup_i w_{i,i} \leq C$ and $|||G_{i,i}|||_2 \leq ||G_{i,i}||_F \leq C$. Hence,

$$|||L(W,G) - L_0(W,G)|||_2 \le \frac{C^2}{n\gamma}$$
.

We still have $L(\widetilde{W}_f, \widetilde{G}) = L(\widetilde{W}, \widetilde{G})$ and of course

$$L_0(\widetilde{W}_f, \widetilde{G}) = L_0(\widetilde{W}, \widetilde{G})$$
.

Note further that the assumptions of Lemma 2.2 apply to the matrices $(W \circ 1_{i \neq j}, G)$ and $(W \circ 1_{i \neq j}, G)$. Indeed, the off-diagonal conditions on the weights are the assumptions we are working under. The diagonal conditions on the weights in Lemma 2.2 are trivially satisfied here since both $W \circ 1_{i \neq j}$ and $\widetilde{W} \circ 1_{i \neq j}$ have diagonal entries equal to 0. Hence, Lemma 2.2 gives

$$|||L_0(W,G) - L_0(\widetilde{W},\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^2$$

Using Weyl's inequality (see [13]) and our bound on $|||L(W,G) - L_0(W,G)|||_2$, we therefore get

$$|||L(W,G) - L_0(\widetilde{W},\widetilde{G})|||_2 \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^2 + \frac{C^2}{n\gamma}.$$

We now turn to the analysis of a specific algorithm, the class averaging algorithm in the cryo-EM problem, with a broadly accepted model of noise contamination to demonstrate the robustness of CGL-like algorithms.

2.2 Impact of noise on the rotationally invariant distance

We assume that we observe noisy versions of the k-dimensional images/objects, $k \ge 2$, we are interested in. If the images in the - unobserved - clean dataset are called $\{S_i\}_{i=1}^n$, we observe

$$I_i = S_i + N_i \; .$$

Here $\{N_i\}_{i=1}^n$ are pure-noise images/objects. Naturally, after discretization, the images/objects we consider are just data vectors of dimension p – we view S_i and N_i as vectors in \mathbb{R}^p . In other words, for a k-dim image, we sample p points from the domain \mathbb{R}^k , which is denoted as $\mathfrak{X} := \{\mathsf{x}_i\}_{i=1}^p \subset \mathbb{R}^k$ and called the sampling grid, and the image is discretized accordingly on these points. We also assume that the random variables N_i 's, $i = 1, \ldots, n$, are independent.

2.2.1 Distance measurement between pairs of images

We start from a general definition. Take a set of linear transforms $\mathcal{T}^{(k)} \subset O(k)$. Consider the following measurement between two objects/images, $d_{ij} \geq 0$, with

$$d_{ij}^2 = \inf_{\mathbf{0} \in \mathcal{T}^{(k)}} \|I_i - \mathbf{0} \circ I_j\|_2^2 \,,$$

where \circ means that the transform is acting on the pixels. For example, in the continuous setup where I_j is replaced by $f_j \in L^2(\mathbb{R}^k)$, given $\mathbf{0} \in SO(k)$, we have $\mathbf{0} \circ f_j(x) := f_j(\mathbf{0}^{-1}x)$ for all $x \in \mathbb{R}^k$. When $\mathcal{T}^{(k)} = SO(k)$, d_{ij} is called *the rotationally invariant distance (RID)*.

In the discrete setup of interest in this paper, we assume that $\mathfrak{X} = 0^{-1}\mathfrak{X}$ for all $0 \in \mathcal{T}^{(k)}$; that is, the linear transform is *exact* (with respect to the grid \mathfrak{X}), in that it maps the sampling grid onto itself. For concreteness, here is an example of sampling grid and associated exact linear transforms. Let k = 2 and take the sampling grid to be the polar coordinates grid. Since we are in dimension 2, we pick m rays of length 1 at angles $2\pi k/m$, $k = 0, \ldots, m-1$ and have l equally spaced points on each of those rays. We consider I_i to be the discretization of the function $f_i \in L^2(\mathbb{R}^2)$ which is compactly supported inside the unit disk, at the polar coordinate grid. The set $\mathcal{T}^{(2)}$ consisting of elements of SO(2) with angles $\theta_k = 2\pi \frac{k}{m}$, where $k = 1, \ldots, m$, is thus exact and associated to the polar coordinate grid.

The discretization and notation merit further discussion. As a linear transform of the domain \mathbb{R}^k , $\mathsf{O} \in \mathcal{T}^{(k)}$ can be represented by a $k \times k$ matrix. On the other hand, in the discretized setup we consider here, we can map $\mathcal{T}^{(k)}$ to a set \mathcal{T} of $p \times p$ matrices O which acts on the discretized images I_j . These images are viewed as a set of p-dim vectors, denoted as I_j^{\vee} , and O acts on a "flattened" or "vectorized" (i.e. 1-dimensional) version of the k-dimensional object of interest. Note that to each transform O there corresponds a unique $p \times p$ matrix O . In the following, we will use O to denote the transform acting on the pixels, and use O to mean its companion matrix acting on the vectorized version of the object we are interested in. A simple but very important observation is that

$$(\mathbf{O} \circ I_i)^{\vee} = OI_i^{\vee}$$
.

In other words, we will have $\inf_{\mathbf{0}\in\mathcal{T}^{(k)}} ||I_i - \mathbf{0}\circ I_j|| = \inf_{\mathbf{0}\in\mathcal{T}} ||I_i^{\vee} - \mathbf{0}I_j^{\vee}||$. To simplify the notation, when it is clear from the context, we will use I_j to mean both the discretized object of interest and its vectorized version.

In what follows, we assume that \mathcal{T} always contains Id_p . We study the impact of noise on d_{ij} through a uniform approximation argument. Let us call for $\mathsf{O} \in \mathcal{T}$,

$$d_{ij,\text{noisy}}^2(\mathsf{O}) := \|I_i^{\vee} - \mathsf{O}I_j^{\vee}\|^2$$
, and $d_{ij,\text{clean}}^2(\mathsf{O}) := \|S_i^{\vee} - \mathsf{O}S_j^{\vee}\|^2$.

Essentially we will show that, when \mathcal{T} contains only orthogonal matrices and is not "too large",

$$\sup_{\mathbf{O}\in\mathcal{T}}\sup_{i\neq j}|d_{ij,\text{noisy}}^2(\mathbf{O}) - d_{ij,\text{clean}}^2(\mathbf{O}) - f(i,j)| = o_P(1) \, .$$

where f(i, j) does not depend on O. Our approximations will in fact be much more precise than this. But we will be able to conclude that in these circumstances,

$$\sup_{i \neq j} \left| \inf_{\mathbf{O} \in \mathcal{T}} d_{ij,\text{noisy}}^2(\mathbf{O}) - \inf_{\mathbf{O} \in \mathcal{T}} d_{ij,\text{clean}}^2(\mathbf{O}) - f(i,j) \right| = o_P(1) \ .$$

We have the following theorem for any given set of transforms \mathcal{T} .

Theorem 2.1. Suppose that for $1 \leq i \leq n$, N_i are independent, with $N_i^{\vee} \sim \mathcal{N}(0, \Sigma_i)$. Call $\mathsf{t}_{\mathsf{p}} := \sup_i \sup_{0 \in \mathcal{T}} \sqrt{\operatorname{trace}\left((\mathsf{O}\Sigma_i \mathsf{O}')^2\right)}$ and $\mathsf{s}_p := \sup_{1 \leq i \leq n} \sup_{0 \in \mathcal{T}} \sqrt{|||\mathsf{O}\Sigma_i \mathsf{O}'|||_2}$. Then, we have

$$\sup_{\mathbf{O}\in\mathcal{T}}\sup_{i\neq j} \sup_{i\neq j} |d_{ij,noisy}^{2}(\mathbf{O}) - d_{ij,clean}^{2}(\mathbf{O}) - trace\left(\Sigma_{i} + \mathbf{O}\Sigma_{j}\mathbf{O}'\right)|$$
$$= O_{P}\left(\sqrt{\log[Card\left\{\mathcal{T}\right\}n^{2}]}\left(\mathsf{t}_{\mathsf{p}} + \mathsf{s}_{p}\sup_{i,\mathsf{O}\in\mathcal{T}} ||\mathsf{O}S_{i}\mathsf{O}'||\right) + \log[Card\left\{\mathcal{T}\right\}n^{2}]\mathsf{s}_{p}^{2}\right)$$

•

Proof. We first note that $N_i^{\vee} - \mathsf{O}N_j^{\vee} \sim \mathcal{N}(0, \Sigma_i + \mathsf{O}\Sigma_j\mathsf{O}')$. Applying Lemma A-2 to $||N_i^{\vee} - \mathsf{O}N_j^{\vee}||^2$ with $Q_i = \mathrm{Id}$, we get

$$\sup_{\mathbf{O}\in\mathcal{T}} \sup_{i\neq j} \left| \|N_i^{\vee} - \mathbf{O}N_j^{\vee}\|^2 - \operatorname{trace}\left(\Sigma_i + \mathbf{O}\Sigma_j\mathbf{O}'\right) \right| = O_P\left(\sqrt{\log[\operatorname{Card}\left\{\mathcal{T}\right\}n^2]} \sup_{i,j,\mathbf{O}} \left[\sqrt{\operatorname{trace}\left((\Sigma_i + \mathbf{O}\Sigma_j\mathbf{O}')^2\right)}\right] + \log[\operatorname{Card}\left\{\mathcal{T}\right\}n^2] \sup_{i,j,\mathbf{O}} |||(\Sigma_i + \mathbf{O}\Sigma_j\mathbf{O}')|||_2\right).$$

Of course, using the fact that for positive semi-definite matrices, $(A + B)^2 \leq 2(A^2 + B^2)$ in the positivesemidefinite order, we have

trace
$$((\Sigma_i + \mathsf{O}\Sigma_j\mathsf{O}')^2) \le 2$$
trace $(\Sigma_i^2 + [\mathsf{O}\Sigma_j\mathsf{O}']^2)$.

Hence,

$$\sup_{\mathbf{O}\in\mathcal{T}} \sup_{i\neq j} \left| \|N_i^{\vee} - \mathbf{O}N_j^{\vee}\|^2 - \operatorname{trace}\left(\Sigma_i + \mathbf{O}\Sigma_j\mathbf{O}'\right) \right| = O_P\left(\sqrt{\log[\operatorname{Card}\left\{\mathcal{T}\right\}n^2]} \sup_{i} \sup_{\mathbf{O}\in\mathcal{T}} \sqrt{\operatorname{trace}\left(\left[\mathbf{O}\Sigma_i\mathbf{O}'\right]^2\right)} + \log[\operatorname{Card}\left\{\mathcal{T}\right\}n^2] \sup_{i} \sup_{\mathbf{O}\in\mathcal{T}} |||\mathbf{O}\Sigma_i\mathbf{O}'|||_2\right) .$$

We also note that

$$(S_i^{\vee} - \mathsf{O}S_j^{\vee})'(N_i^{\vee} - \mathsf{O}N_j^{\vee}) \sim \mathcal{N}(0, \gamma_{i,j,\mathsf{O}}^2) ,$$

where

$$\gamma_{i,j,\mathbf{O}}^2 = (S_i^{\vee} - \mathbf{O}S_j^{\vee})'(\Sigma_i - \mathbf{O}\Sigma_j\mathbf{O}')(S_i^{\vee} - \mathbf{O}S_j^{\vee})$$

We note that

$$\gamma_{i,j,\mathsf{O}}^2 \le 2\left(\|S_i^{\vee}\|^2 + \|\mathsf{O}S_j^{\vee}\mathsf{O}'\|^2\right)\left(|||\Sigma_i|||_2 + |||\mathsf{O}\Sigma_j\mathsf{O}'|||_2\right) \le 8\sup_{i,\mathsf{O}\in\mathcal{T}}|||\mathsf{O}\Sigma_i\mathsf{O}'|||_2\sup_{i,\mathsf{O}\in\mathcal{T}}\|\mathsf{O}S_i^{\vee}\mathsf{O}'\|^2$$

Recall also that it is well known that if Z_1, \ldots, Z_N are $\mathcal{N}(0, \gamma_i^2)$ random variables,

$$\sup_{1 \le k \le N} |Z_k| = \mathcal{O}_P(\sqrt{\log N} \sup_k \gamma_k) +$$

This result can be obtained by a simple union bound argument. In our case, it means that

$$\sup_{\mathsf{O}\in\mathcal{T}} \sup_{i\neq j} |(S_i^{\vee} - \mathsf{O}S_j^{\vee})'(N_i^{\vee} - \mathsf{O}N_j^{\vee})| = \mathcal{O}_P\left(\sqrt{\log[\operatorname{Card}\left\{\mathcal{T}\right\}n^2]} \sup_{i,\mathsf{O}\in\mathcal{T}} \sqrt{|||\mathsf{O}\Sigma_i\mathsf{O}'|||_2} \sup_{i,\mathsf{O}\in\mathcal{T}} ||\mathsf{O}S_i^{\vee}\mathsf{O}'||\right).$$

In light of the previous theorem, we have the following proposition.

Proposition 2.1. Suppose that for all $1 \le i \le n$ and $\mathbf{O} \in \mathcal{T}$, $|||\mathbf{O}\Sigma_i\mathbf{O}'|||_2 \le \sigma_p^2$, $\sqrt{trace([\mathbf{O}\Sigma_i\mathbf{O}']^2)/p} \le s_p^2$, and $||\mathbf{O}S_i^{\vee}\mathbf{O}'|| \le K$, where K is a constant independent of p. Then,

$$\sup_{\mathbf{O}\in\mathcal{T}}\sup_{i\neq j}|d_{ij,noisy}^2(\mathbf{O}) - d_{ij,clean}^2(\mathbf{O}) - trace\left(\Sigma_i + \mathbf{O}\Sigma_j\mathbf{O}'\right)| = \mathcal{O}_P(u_{n,p}).$$
(2)

where $u_{n,p} := \sqrt{\log[Card\{\mathcal{T}\}n^2]}(\sqrt{p}s_p^2 + K\sigma_p) + \log[Card\{\mathcal{T}\}n^2]\sigma_p^2$. It follows that, if $\sqrt{\log[Card\{\mathcal{T}\}n^2]}\max(\sqrt{p}s_p^2,\sigma_p) \to 0$, and \mathcal{T} contains only orthogonal matrices,

$$\sup_{\mathbf{O}\in\mathcal{T}}\sup_{i\neq j}|d_{ij,noisy}^{2}(\mathbf{O})-d_{ij,clean}^{2}(\mathbf{O})-trace(\Sigma_{i}+\Sigma_{j})|=O_{P}(u_{n,p})=o_{P}(1)$$

Furthermore, in this case,

$$\sup_{i \neq j} \left| d_{ij,noisy}^2 - d_{ij,clean}^2 - trace(\Sigma_i + \Sigma_j) \right| = o_P(1)$$

where

$$d_{ij,noisy}^2 := \inf_{\mathsf{O}\in\mathcal{T}} \|I_i^{\vee} - \mathsf{O}I_j^{\vee}\|^2 , \quad d_{ij,clean}^2 := \inf_{\mathsf{O}\in\mathcal{T}} \|S_i^{\vee} - \mathsf{O}S_j^{\vee}\|^2 .$$

In light of the previous proposition, the following set of assumptions is natural:

Assumption G1: $\forall i, \mathbf{O} \in \mathcal{T}, |||\mathbf{O}\Sigma_i\mathbf{O}'|||_2 \leq \sigma_p^2, \sqrt{\operatorname{trace}\left([\mathbf{O}\Sigma_i\mathbf{O}']^2\right)/p} \leq s_p^2, \text{ and } ||\mathbf{O}S_i^{\vee}\mathbf{O}'|| \leq K, \text{ where } K \text{ is a constant independent of } p. \text{ Furthermore, } \sqrt{\log\left[\operatorname{Card}\left\{\mathcal{T}\right\}n^2\right]}\max(\sqrt{p}s_p^2, \sigma_p) \to 0 \text{ and hence } u_{n,p} \to 0.$

We refer the reader to Proposition A.1 on page 26 for a bound on Card $\{\mathcal{T}\}$ that is relevant to the class averaging algorithm in the cryo-EM problem.

Proof of Proposition 2.1. The first two statements are immediate consequences of Theorem 2.1. For the second one, we use the fact that since $O \in \mathcal{T}$ is orthogonal, trace $(O\Sigma_j O') = \text{trace}(\Sigma_j)$.

Now, if F and G are two functions, we clearly have $|\inf F(x) - \inf G(x)| \le \sup |G(x) - F(x)|$. Indeed, $\forall x, F(x) \le G(x) + \sup |G(x) - F(x)|$. Hence, for all x,

$$\inf_{x} F(x) \le F(x) \le G(x) + \sup |G(x) - F(x)| ,$$

and we conclude by taking inf in the right-hand side. The inequality is proved similarly in the other direction. The results of Theorem 2.1 therefore show that

$$\sup_{i \neq j} \left| d_{ij,\text{noisy}}^2 - d_{ij,\text{clean}}^2 - \text{trace}\left(\Sigma_i + \Sigma_j\right) \right| = \mathcal{O}_P(\sqrt{\log[\text{Card}\left\{\mathcal{T}\right\}n^2]}(\sqrt{p}s_p^2 + K\sigma_p) + \log[\text{Card}\left\{\mathcal{T}\right\}n^2]\sigma_p^2)$$

and we get the announced conclusions under our assumptions.

We now present two examples to show that our assumptions are quite weak and prove that the algorithms we are studying can tolerate very large amount of noise.

Magnitude of noise: First example Assume that $N_i^{\vee} \sim p^{-(1/4+\epsilon)} \mathcal{N}(0, \mathrm{Id}_p)$, where $\epsilon > 0$. In this case, $||N_i^{\vee}|| \sim p^{1/4-\epsilon} \gg \sup_i ||S_i^{\vee}||$ if $\epsilon < 1/4$. In other words, the norm of the error vector is much larger than the norm of the signal vector. Indeed, asymptotically, the signal to noise ratio $||S_i^{\vee}|| / ||N_i^{\vee}||$ is 0. Furthermore, $\sigma_p = p^{-(1/4+\epsilon)}$ and $\sqrt{p}s_p^2 = p^{-2\epsilon}$. Hence, if Card $\{\mathcal{T}\} = O(p^{\gamma})$ for some γ , our conditions translate into $\sqrt{\log(np)} \max(p^{-(1/4+\epsilon)}, p^{-2\epsilon}) \to 0$. This is of course satisfied provided n is subexponential in p.

Magnitude of noise: Second example We now consider the case where Σ_i has one eigenvalue equal to $p^{-\epsilon}$ and all the others are equal to $p^{-(1/2+\eta)}$, $\epsilon, \eta > 0$. In other words, the noise is much larger in one direction than in all the others. In this case, $\sigma_p^2 = p^{-\epsilon}$ and trace $(\Sigma_i^2) = p^{-2\epsilon} + (p-1)*p^{-(1+2\eta)} \le p^{-2\epsilon} + p^{-2\eta}$. So if once again, Card $\{\mathcal{T}\} = O(p^{\gamma})$, our conditions translate into $\sqrt{\log(np)} \max(p^{-\epsilon} + p^{-\eta}, p^{-\epsilon/2}) \to 0$. This example would also work if the number of eigenvalues equal to $p^{-\epsilon}$ were $o(p^{2\epsilon}/[\log(np)])$, provided $\sqrt{\log(np)} \max(p^{-\eta}, p^{-\epsilon/2}) \to 0$.

Comment on the conditions on the signal in Assumption G1 At first glance, it might look like the condition $\sup_{i,O\in\mathcal{T}} ||OS_i^{\vee}O'|| \leq K$ is very restrictive due to the fact that, after discretization, S_i has p pixels. However, it is typically the case that if we start from a function in $L^2(\mathbb{R}^k)$, the discretized and vectorized image S_i^{\vee} is normalized by the number of pixels p, so that $||S_i^{\vee}||$ is roughly equal to the L^2 -norm of the corresponding function. Hence, our condition $\sup_{i,O\in\mathcal{T}} ||OS_i^{\vee}O'|| \leq K$ is very reasonable.

2.2.2 The case of "exact rotations"

We now focus on the most interesting case for our problem, namely the situation where $\mathbf{0}$ leaves our sampling grid invariant. We call $\mathcal{T}_{\text{exact}}^{(k)} \subset SO(k)$ the corresponding matrices $\mathbf{0}$ and $\mathcal{T}_{\text{exact}}$ the companion $p \times p$ matrices. We note that $\mathcal{T}_{\text{exact}}^{(k)}$ depends on p, but since this is evident, we do not index $\mathcal{T}_{\text{exact}}^{(k)}$ by p to avoid cumbersome notations. From the standpoint of statistical applications, our focus in this paper is mostly on the case k = 1 (which corresponds to "standard" kernel methods commonly used in statistical learning) and k = 2.

We show in Proposition A.1 that if $O \in \mathcal{T}_{exact}$, O is an orthogonal $p \times p$ matrix. Furthermore, we show in Proposition A.1 that Card $\{\mathcal{T}_{exact}\}$ is polynomial in p. We therefore have the following proposition.

Proposition 2.2. Let

$$d_{ij,noisy}^2 := \inf_{\mathbf{0} \in \mathcal{T}_{exact}^{(k)}} \|I_i - \mathbf{0} \circ I_j\|^2 , \quad d_{ij,clean}^2 := \inf_{\mathbf{0} \in \mathcal{T}_{exact}^{(k)}} \|S_i - \mathbf{0} \circ S_j\|^2 .$$

Suppose N_i are independent with $N_i^{\vee} \sim \mathcal{N}(0, \Sigma_i)$. When Assumption G1 holds with \mathcal{T}_{exact} being the set of companion matrices of $\mathcal{T}_{exact}^{(k)}$, we have

$$\sup_{i \neq j} \left| d_{ij,noisy}^2 - d_{ij,clean}^2 - trace \left(\Sigma_i + \Sigma_j \right) \right| = o_P(1) ,$$

and

$$\sup_{\substack{\in \mathcal{T}_{exact}^{(k)} \ i \neq j}} \sup_{\substack{i \neq j}} |d_{ij,noisy}^2(\mathbf{0}) - d_{ij,clean}^2(\mathbf{0}) - trace(\Sigma_i + \Sigma_j)| = \mathcal{O}_P(u_{n,p}) = \mathcal{O}_P(1) .$$

2.2.3 On the transform $O_{ii,noisy}^*$

We now use the notations

0

$$d_{ij,\text{noisy}}(\mathbf{0}) = \|I_i - \mathbf{0} \circ I_j\|$$
, and $d_{ij,\text{clean}}(\mathbf{0}) = \|S_i - \mathbf{0} \circ S_j\|$.

Naturally, the study of

$$\mathbf{0}_{ij,\text{noisy}}^* = \operatorname{argmin}_{\mathbf{0} \in \mathcal{T}_{\text{exact}}^{(k)}} d_{ij,\text{noisy}}(\mathbf{0})$$
(3)

is more complicated than the study of $\inf_{\mathbf{0}\in\mathcal{T}_{exact}^{(k)}} d_{ij,noisy}(\mathbf{0})$. We will assume that the clean images are nicely behaved when it comes to the $d_{ij,clean}(\mathbf{0})$ minimization, in that rotations that are near minimizers of $d_{ij,clean}(\mathbf{0})$ are close to one another. More formally, we assume the following.

Assumption A0 : $\mathcal{T}_{\text{exact}}^{(k)}$ is a subset of SO(k) and contains only exact rotations. Call $\mathbb{O}_{ij,\text{clean}}^* := \arg\min_{\mathbf{0}\in\mathcal{T}_{\text{exact}}^{(k)}} d_{ij,\text{clean}}^2(\mathbf{0})$ and call $\mathcal{T}_{ij,\epsilon}^{(k)} := \left\{\mathbf{0}\in\mathcal{T}_{\text{exact}}^{(k)}: d_{ij,\text{clean}}^2(\mathbf{0}) \le d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{clean}}^*) + \epsilon\right\}$. We assume that

$$\exists \delta_{ij,p} > 0 : \forall \epsilon < \delta_{ij,p} \forall \mathbf{0} \in \mathcal{T}_{ij,\epsilon}^{(k)}, \ d(\mathbf{0}, \mathbf{0}_{ij,\text{clean}}^*) \le g_{ij,p}(\epsilon) ,$$

for d the canonical metric on the orthogonal group and some positive $g_{ij,p}(\epsilon)$.

Assumption A1 : $\delta_{ij,p}$ can be chosen independently of i, j and p. Furthermore, there exists a function g such that $g(\epsilon) \to 0$ as $\epsilon \to 0$ and $g_{ij,p}(x) \le g(x)$, if $x \le \delta_{ij,p} \le \delta$.

We discuss the meaning of these assumptions after the statement and proof of the following theorem.

Theorem 2.2. Suppose that the assumptions underlying Theorem 2.1 hold and that Assumptions G1, A0 and A1 hold. Suppose further that $\mathcal{T}_{exact}^{(k)}$ is the set of exact rotations for our discretization. Then, for any η given, where $0 < \eta < 1$, as p and n go to infinity,

$$\sup_{i \neq j} d(\mathbf{O}_{ij,noisy}^*, \mathbf{O}_{ij,clean}^*) = \mathcal{O}_P(g(u_{n,p}^{1-\eta})) , \qquad (4)$$

where $u_{n,p}$ is defined in (2). (Under Assumption G1, $u_{n,p} \to 0$ as n and p tend to infinity.)

The informal meaning of this theorem is that under regularity assumptions on the set of clean images, the optimal rotation computed from the set of noisy images is close to the optimal rotation computed from the set of clean images. In other words, this step of the CGL procedure is robust to noise.

Proof. Clearly, $\mathsf{O}_{ij,clean}^*$ is a minimizer of $L_{ij}(\mathsf{O}) := d_{ij,clean}^2(\mathsf{O}) + \operatorname{trace}(\Sigma_i + \Sigma_j)$, since the second term does not depend on O . Naturally, if Assumptions A0 and A1 apply to $d_{ij,clean}^2(\mathsf{O})$, they apply to $d_{ij,clean}^2(\mathsf{O}) + C$, for C any constant. In particular, taking $C = \operatorname{trace}(\Sigma_i + \Sigma_j)$, we see that Assumptions A0 and A1 apply to the function $L_{ij}(\mathsf{O})$.

The approximation results of Proposition 2.2 guarantee that, under Assumptions G1, A0 and A1, $O_{ij,\text{noisy}}^*$ is a near minimizer of $d_{ij,\text{clean}}^2(\mathbf{0})$. Indeed, we have by definition,

$$d_{ij,\text{noisy}}^2(\mathsf{O}_{ij,\text{noisy}}^*) \le d_{ij,\text{noisy}}^2(\mathsf{O}_{ij,\text{clean}}^*).$$
(5)

But under assumption G1, Proposition 2.2 and the fact that the elements of $\mathcal{T}_{\text{exact}}$ are orthogonal matrices imply that

$$\forall \mathbf{0} \in \mathcal{T}, \forall i \neq j \ d_{ij,\text{noisy}}^2(\mathbf{0}) = L_{ij}(\mathbf{0}) + \mathcal{O}_P(u_{n,p}) .$$
(6)

Hence, we can rephrase Equation (5) as

$$L_{ij}(\mathbf{0}_{ij,\text{noisy}}^*) \le L_{ij}(\mathbf{0}_{ij,\text{clean}}^*) + \mathcal{O}_P(u_{n,p}) .$$

$$\tag{7}$$

Indeed, by plugging (6) into (5), we have

$$\forall i \neq j, \ d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{noisy}}^*) + \text{trace}\left(\Sigma_i + \Sigma_j\right) \le d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{clean}}^*) + \text{trace}\left(\Sigma_i + \Sigma_j\right) + \mathcal{O}_P(u_{n,p})$$

Now, by definition of $d_{ij,\text{clean}}^2$, we have

$$d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{clean}}^*) \le d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{noisy}}^*) .$$
(8)

So by (7) and (8), we have shown that

$$\begin{aligned} \forall i \neq j, \ d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{clean}}^*) + \text{trace}\left(\Sigma_i + \Sigma_j\right) &\leq d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{noisy}}^*) + \text{trace}\left(\Sigma_i + \Sigma_j\right) \ ,\\ &\leq d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{clean}}^*) + \text{trace}\left(\Sigma_i + \Sigma_j\right) + \mathcal{O}_P(u_{n,p}) \ .\end{aligned}$$

This clearly implies that

$$\forall i \neq j, \ d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{clean}}^*) \leq d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{noisy}}^*) \leq d_{ij,\text{clean}}^2(\mathbf{0}_{ij,\text{clean}}^*) + \mathcal{O}_P(u_{n,p}) \ .$$

Since $u_{n,p} \to 0$ as n and p grow, this means that, for any given η , with $0 < \eta < 1$, with very high probability,

$$\forall 1 \le i \ne j \le n , \mathbf{0}_{ij,\text{noisy}}^* \in \mathcal{T}_{ij,u_{n,p}^{1-\eta}}^{(k)}.$$

We conclude, using Assumption A0, that with very high-probability,

ł

$$\forall 1 \leq i \neq j \leq n , d(\mathbf{0}^*_{ij,\text{noisy}},\mathbf{0}^*_{ij,\text{clean}}) \leq g(u^{1-\eta}_{n,p}) .$$

Interpretation of Assumptions A0-A1 Assumption A0 guarantees that all near minimizers of $d_{ij,\text{clean}}(0)$ are close to one another and hence the optimum. Our uniform bounds in Proposition 2.2 only guarantee that $O^*_{ij,\text{noisy}}$ is a near minimizer of $d_{ij,\text{clean}}(0)$ and nothing more. If $d_{ij,\text{clean}}(0)$ had near minimizers that were far from the optimum $O^*_{ij,\text{clean}}$, it could very well happen that $O^*_{ij,\text{noisy}}$ end up being close to one of these near minimizers but far from $O^*_{ij,\text{clean}}$, and we would not have the consistency result of Theorem 2.2. Hence, the robustness to noise of this part of the CGL algorithm is clearly tied to some regularity or "niceness" property for the set of clean images.

In the cryo-EM problem, these assumptions reflect a fundamental property of a manifold dataset – *its condition number* [50]. Conceptually, the condition number reflects "how difficult it is to reconstruct the manifold" from a finite sample of points from the manifold. Precisely, it is the inverse of the reach of the manifold, which is defined to be the radius of the smallest normal bundle that is homotopic to the manifold. This also highlights the fact that even if we were to run the CGL algorithm on the clean dataset, without these assumptions, the results might not be stable and reliable since intrinsically distant points (i.e. distant in the geodesic distance) might be identified as neighbors.

About $\mathcal{T}_{exact}^{(k)}$ and extensions We are chiefly interested in this paper about 2-dimensional images and hence about the case k = 2 (see the cryoEM example). It is then clear that when our polar coordinate grid is fine, $\mathcal{T}_{exact}^{(k)}$ is also a fine discretization of SO(2) and contains many elements. (More details are given in Subsection A-3.) The situation is more intricate when $k \geq 3$, but since it is a bit tangential to the main purpose of the current paper, we do not discuss it further here. We refer the interested reader to Subsection A-3 for more details about the case $k \geq 3$.

We also note that our arguments are not tied to using a standard polar coordinate grid for the discretization of the images. For another sampling grid, we would possibly get another $\mathcal{T}_{\text{exact}}^{(k)}$. Our arguments go through when : a) if $\mathbf{0} \in \mathcal{T}^{(k)}$, the operation $\mathbf{0}$ ° maps our sampling grid of points onto itself; b) Card $\{\mathcal{T}^{(k)}\}$ grows polynomially in p.

2.2.4 Extensions and different approaches

At the gist of our arguments are strong concentration results for quadratic forms in Gaussian random variables. Naturally, our results extend to other types of random variables for which these concentration properties hold. We refer to [43] and [29] for examples. A natural example in our context would be a situation where $N_i = \sum_i^{1/2} X_i$, and X_i has i.i.d uniformly bounded entries. This is particularly relevant in the case where Σ_i is diagonal for instance - the interpretation being then that the noise contamination is through the corruption of each individual pixel by independent random variables with possibly different standard deviations. The arguments in Lemma A-1 handle this case, though the bound is slightly worse than the one in Lemma A-2 when a few eigenvalues of Σ_i are larger than most of the others. Indeed, the only thing that matters in this more general analysis is the largest eigenvalue of Σ_i , so that in the notation of Assumption G1, $\sqrt{p}s_p^2$ is replaced by $\sqrt{p}\sigma_p^2$. Hence, our approximation will require in this more general setting that $\sigma_p = o(p^{-1/4})$, whereas we have seen in the Gaussian case that we can tolerate a much larger largest eigenvalue.

We also note that we could of course settle for weaker results on concentration of quadratic forms, which would apply to more distributions. For instance, using bounds on $\mathbf{E}\left(||N_i||^2 - \mathbf{E}\left(||N_i||^2\right)|^k\right)$ would change the dependence of results such as Proposition 2.1 on Card $\{\mathcal{T}\}$ n^2 from powers of logarithm to powers of 1/k. This is in turn would mean that our results would become tolerant to lower levels of noise but apply to more noise distributions.

2.3 Consequences for CGL algorithm and other kernel-based methods

2.3.1 Reminders and preliminaries

Recall that in CGL methods performed with the rotationally invariance distance - henceforth RID induced by SO(k), we mostly care about the spectral properties - especially large eigenvalues and corresponding eigenvectors - of the CGL matrix $L(\widetilde{W}, \widetilde{G})$, where \widetilde{W} is a $n \times n$ matrix and \widetilde{G} is a $nk \times nk$ block-matrix with $k \times k$ blocks defined through

$$\widetilde{W}_{i,j} = \exp(-d_{ij,\text{noisy}}^2/\epsilon), \quad \widetilde{G}_{i,j} = \mathbf{O}_{ij,\text{noisy}}^*,$$

where $O_{ij,\text{noisy}}^*$ is defined in Equation (3).

The "good" properties of CGL stem from the fact that the matrix L(W, G), the CGL matrix associated with the clean images, has "good" spectral properties. For example, when a manifold structure is assumed, the theoretical work of [59, 61] relates the properties of L(W, G) - the matrix obtained in the same manner as above when we replace $d_{ij,noisy}$ by $d_{ij,clean}$ and $O^*_{ij,noisy}$ by $O^*_{ij,clean}$ - to the geometric and topological properties of the manifold from which the data is sampled. The natural approximate "sparsity" of the spectrum of this kind of matrices is discussed in Section B.

In practice, the data analyst has to work with $L(\widetilde{W}, \widetilde{G})$. Hence, it could potentially be the case that $L(\widetilde{W}, \widetilde{G})$ does not share many of the good properties of L(W, G). Indeed, we explain below that this is in general the case and propose a modification to the standard algorithm to make the results of CGL methods more robust to noise. All these arguments suggest that it is natural to study the properties of the standard CGL algorithm applied to noisy data.

We mention that CGL algorithms may apply beyond the case of the rotational invariance distance and O(k) and we explain in Subsubsection 2.3.3 how our results apply in this more general context.

2.3.2 Modified CGL algorithm and rotationally invariant distance

We now show that our modification to the standard algorithm is robust to noise. More precisely, we show that the modified CGL matrix $L_0(\widetilde{W}, \widetilde{G})$ is spectrally close to the CGL matrix computed from the noise-free data, L(W, G).

Proposition 2.3. Consider the modified CGL matrix $L_0(\widetilde{W}, \widetilde{G})$ computed from the noisy data and the CGL matrix L(W, G) computed from the noise-free data. Under Assumptions **G1** and **A0-A1**, we have, if trace $(\Sigma_i) = trace(\Sigma_j) = trace(\Sigma)$ for all (i, j),

$$|||L_0(\widetilde{W}, \widetilde{G}) - L(W, G)|||_2 = o_P(1)$$
,

provided there exists $\gamma > 0$, independent of n and p such that

$$\inf_{i} \sum_{j \neq i} \frac{\exp(-d_{ij,clean}^2/\epsilon)}{n} \ge \gamma > 0 \; .$$

Note that the previous result means that $L_0(\widetilde{W}, \widetilde{G})$ and L(W, G) are essentially spectrally equivalent: indeed we can use the Davis-Kahan theorem or Weyl's inequality to relate eigenvectors and eigenvalues of $L_0(\widetilde{W}, \widetilde{G})$ to those of L(W, G) (see [63], [13] or [28] for a brief discussion putting all the needed results together). In particular, if the large eigenvalues of L(W, G) are separated from the rest of the spectrum, the eigenvalues of $L_0(\widetilde{W}, \widetilde{G})$ and corresponding eigenspaces will be close to those of L(W, G).

Proof. The proposition is a simple consequence of our previous results and Lemma 2.3 above. Indeed, in the notation of Lemma 2.3, we call

$$w_{i,j} = \begin{cases} \exp(-d_{ij,\text{clean}}^2/\epsilon) & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases} \quad \text{and} \quad \tilde{w}_{i,j} = \begin{cases} \exp(-d_{ij,\text{noisy}}^2/\epsilon) & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}$$

Similarly, we call

$$G_{i,j} = \begin{cases} \mathsf{O}_{ij,clean}^* & \text{if } i \neq j \\ \mathrm{Id}_d & \text{if } i = j \end{cases} \quad \text{and} \quad \widetilde{G}_{i,j} = \begin{cases} \mathsf{O}_{ij,noisy}^* & \text{if } i \neq j \\ \mathrm{Id}_d & \text{if } i = j \end{cases}$$

Under Assumption G1, we know that, if $f_i = \exp(-2\operatorname{trace}(\Sigma)/\epsilon)$, $\sup_{i\neq j} |w_{i,j} - \tilde{w}_{i,j}/f_i| = o_P(1)$. Similarly, under Assumptions G1, A0 and A1, we know, using Theorem 2.2 that

$$\sup_{i,j} d(G_{i,j}, \widetilde{G}_{i,j}) = o_P(1)$$

and therefore, since k, the parameter of SO(k), is held fixed in our asymptotics,

$$\sup_{i,j} \|G_{i,j} - \widetilde{G}_{i,j}\|_F = \mathrm{o}_P(1) \; .$$

Since we assumed that

$$\inf_{i} \sum_{j \neq i} \frac{\exp(-d_{ij,\text{clean}}^2/\epsilon)}{n} \ge \gamma > 0 ,$$

i.e, in the notations of Lemma 2.3

$$\inf_{i} \frac{\sum_{j \neq i} w_{i,j}}{n} \ge \gamma > 0 ,$$

where γ is independent of n and p, all the assumptions of Lemma 2.3 are satisfied when n and p are large enough, and we conclude that, in the notations of this lemma,

$$|||L_0(W,G) - L_0(\widetilde{W},\widetilde{G})|||_2 = o_P(1)$$

Furthermore, we have $0 \le w_{i,j}$, $\tilde{w}_{i,j} \le 1$, $\|G_{i,j}\|_F \le \sqrt{k}$ and $\|\widetilde{G}_{i,j}\|_F \le \sqrt{k}$, the latter two results coming from the fact that the columns of $G_{i,j}$ and $\widetilde{G}_{i,j}$ have unit norm. So we conclude that

$$|||L(W,G) - L_0(\widetilde{W},\widetilde{G})|||_2 = o_P(1)$$
.

Is the modification of the algorithm really needed? It is natural to ask what would have happened if we had not modified the standard algorithm, i.e if we had worked with $L(\widetilde{W}, \widetilde{G})$ instead of $L_0(\widetilde{W}, \widetilde{G})$. It is easy to see that

$$L(\widetilde{W},\widetilde{G}) = L_0(\widetilde{W},\widetilde{G}) + \mathsf{D}$$

where D is a block diagonal matrix with

$$\mathsf{D}(i,i) = \frac{\tilde{w}_{i,i}}{\sum_{j \neq i} \tilde{w}_{i,j}} \mathrm{Id}_k = \frac{1}{\sum_{j \neq i} \tilde{w}_{i,j}} \mathrm{Id}_k \; .$$

Under our assumptions,

$$|||n\exp(-2\operatorname{trace}(\Sigma)/\epsilon)\mathsf{D} - D\left(\left\{\frac{\sum_{j\neq i}\exp(-d_{ij,\operatorname{clean}}^2/\epsilon)}{n}\right\}_{i=1}^n\right)|||_2 = o_P(1).$$

We also recall that under Assumption **G1**, trace (Σ) can be as large as $p^{1/2-\eta}$ - a very large number in our asymptotics. So in particular, if *n* is polynomial in *p*, we have then $n^{-1} \exp(2 \operatorname{trace}(\Sigma) / \epsilon) \to \infty$. This implies that

$$L(\widetilde{W},\widetilde{G}) = L_0(\widetilde{W},\widetilde{G}) + \mathsf{D}$$

is then dominated in spectral terms by D. So it is clear that in the high-noise regime, if we had used the standard CGL algorithm, the spectrum of $L(\widetilde{W}, \widetilde{G})$ would have mirrored that of D - which has little to do with the spectrum of L(W, G), which we are trying to estimate - and the noise would have rendered the algorithm ineffective.

By contrast, by using the modification we propose, we guarantee that even in the high-noise regime, the spectral properties of $L_0(\widetilde{W}, \widetilde{G})$ mirror those of L(W, G). We have hence made the CGL algorithm more robust to noise.

On the use of nearest neighbor graphs In practice, variants of the CGL algorithms we have described use nearest neighbor information to replace $w_{i,j}$ by 0 if $w_{i,j}$ is not among the k largest elements of $\{w_{i,j}\}_{j=1}^n$. In the high-noise setting, the nearest-neighbor information is typically not robust to noise, which is why we proposed to use all the $w_{i,j}$'s and avoid the nearest neighbor variant of the CGL algorithm, even though the latter probably makes more intuitive sense in the noise-free context. A systematic study of the difference between these two variants is postponed to future work.

Comparison with previous results in the literature As far as we know, the study of the impact of high-dimensional additive noise on kernel methods was started in [30]. Compared to this paper, our extension is two-fold: 1) the noise level (i.e trace (Σ)) that is studied in the current paper is much higher than what was studied in [30]. This is partly a result of the fact that the current paper focuses on the Gaussian kernel whereas [30] studied many more kernels. 2) [30] focused on standard kernel methods based on the graph Laplacian, such as diffusion maps - where the connection information is not included in the data analysis. Incorporating this new element creates new difficulties. In other respects, we also refer to [57] for another study of the influence of noise in a different setup.

2.3.3 CGL beyond the rotational invariance distance

The previous analysis has been carried out for the RID and corresponding rotations for whom we studied the impact of additive noise in Subsection 2.2. However, it is clear that our results apply much more broadly. We have the following proposition.

Proposition 2.4. Suppose we are given a collection $d_{i,j,noisy}$ of (scalar-valued) dissimilarities between noisy versions of objects i and j, $1 \le i, j \le n$. Suppose objects i and j have (scalar-valued) dissimilarity $d_{i,j,clean}$. Consider the asymptotic regime where $n \to \infty$ and suppose that there exists $\xi_n \in \mathbb{R}$ such that

$$\sup_{i\neq j} |\mathsf{d}_{i,j,noisy}^2 - \mathsf{d}_{i,j,clean}^2 - \xi_n| = o_P(1) \; .$$

Call $\tilde{w}_{i,j} = \exp(-\mathsf{d}_{i,j,noisy}^2/\nu)$ and $w_{i,j} = \exp(-\mathsf{d}_{i,j,clean}^2/\nu)$ the corresponding affinities. ν is held fixed in our asymptotics, though the way affinities are computed may change with n.

Suppose $G_{i,j}$ is the connection between noisy versions of objects i and j and $G_{i,j}$ is the connection between the clean version of objects i and j. Suppose that $w_{i,j}$, $G_{i,j}$ and $\tilde{G}_{i,j}$ satisfy the assumptions of Lemma 2.3, with ϵ and η possibly random but $o_P(1)$ and γ bounded below as $n \to \infty$. Then

$$|||L(W,G) - L_0(\widetilde{W},\widetilde{G})|||_2 = o_P(1)$$
.

Proof. This proposition is just a consequence of Lemma 2.3. Indeed, the affinities are all bounded by 1. Furthermore, we can use $f_i = \exp(-\xi_n/\nu)$ and all the approximation results needed in Lemma 2.3 are true, so the result follows.

2.3.4 A situation without robustness to noise

So far, our work has been quite general and has shown that when the noise is Gaussian (or Gaussianlike) and its covariance Σ_i is such that trace $(\Sigma_i) = \text{trace}(\Sigma_j)$ for all i, j, CGL algorithms can be made robust to noise.

It has been recognized [26, 29, 30, 32] that to study the robustness of various statistical procedures in high-dimension, it is essential to move beyond the Gaussian-like situation and study for instance elliptical/scale mixture of Gaussian models. This largely due to the peculiar geometry of high-dimensional Gaussian and Gaussian-like vectors (see above references and [38]).

If we now write down a model for the noise where $N_i = \lambda_i Z_i$, where Z_i are i.i.d $\mathcal{N}(0, \Sigma)$, λ_i 's are i.i.d with $\mathbf{E}(\lambda_i^2) = 1$ and $\lambda_i \in \mathbb{R}$ is independent of Z_i , it is easy to modify our analysis (assuming for instance that λ_i^2 are bounded, though this condition could easily be relaxed) and to realize that our main approximation result in Proposition 2.1 is replaced by

$$\sup_{i \neq j} \left| d_{ij,\text{noisy}}^2 - d_{ij,\text{clean}}^2 - [\lambda_i^2 + \lambda_j^2] \text{trace}\left(\Sigma\right) \right| = o_P(1) \; .$$

In this situation, Theorem 2.2 is still valid. However, Proposition 2.3 is not valid anymore. The matrix $L_0(\widetilde{W}, \widetilde{G})$ can be approximated by a matrix that depends both on the signal and the distribution of the λ_i^2 's. And there is no guarantee in general that this matrix will have approximately the same spectral properties as L(W, G) or $L_0(W, G)$, the CGL matrix generated from the noise-free signals. This suggests that even our modification of the original CGL algorithm will not be robust to this "elliptical"-noise contamination.

3 Numerical work

Although the robustness properties of CGL methods were not well studied in the past, these methods have been successfully applied to different problems; for example, [2, 25, 45, 59, 62, 71]. In this section, we show simulated examples to illustrate the practical performance of our theoretical findings about CGL methods. We refer interested readers to the aforementioned papers for details and results of its applications.

To demonstrate the main finding of this paper - that CGL methods are robust to high-levels of noise in the spectral sense - we take the noise to be a random Gaussian vector $Z \sim \mathcal{N}(0, cI_p/p^{\alpha})$, where $\alpha \leq 1$ and c > 0. Note that the amount of noise, or the trace of the covariance matrix of Z, is $cp^{1-\alpha}$ and will blow up when $p \to \infty$ and $\alpha < 1$.

3.1 1-dim manifold

Our first example is a dataset sampled from a low dimensional manifold, which is embedded in a high dimensional space. This dataset can be viewed as a collection of high dimensional points which is (locally) parametrized by only few parameters², but in a nonlinear way.

As a concrete example, we take the twisted bell-shaped simple and closed curve, denoted as M, embedded in the first 3 axes of \mathbb{R}^p , where $p \gg 2$, via $\iota : [0, 2\pi) \to \mathbb{R}^p$:

$$\iota: t \mapsto [\cos(t), (1 - 0.8e^{-8\cos^2 t})\cos(\pi(\cos(t) + 1)/4), (1 - 0.8e^{-8\cos^2 t})\sin(\pi(\cos(t) + 1)/4), 0, \dots, 0] \in \mathbb{R}^p.$$

M is a 1-dim smooth manifold without boundary; that is, no matter how big p is, locally the points on M can be parametrized by only 1 parameter. See Figure 1 (A) for an illustration. We mention that one interesting dataset of this kind is the 2-D tomography from noisy projections taken at unknown random directions [60].

For our numerical work, we independently sample n points uniformly at random from $[0, 2\pi)$. Due to the non-linear nature of ι , it is equivalent to non-uniformly sampling n points from M independently. Denote the clean data as $\mathcal{Y} = \{y_i\}_{i=1}^n \subset M$. The data $\mathcal{X} = \{x_i\}_{i=1}^n$ we analyze is the clean data contaminated by noise, i.e $x_i = y_i + Z_i$, with Z_i i.i.d with the same distribution as Z. We measure the signal-to-noise ratio of the dataset by the quantity snrdb := $20 \log \frac{\sqrt{\mathbb{E}X^T X}}{\sqrt{\mathbb{E}Z^T Z}}$. We take n = p = 1000 and $\alpha = 1/4$. Note that $\alpha = 1/4$ is the critical value in our analysis beyond which our results do not apply. It corresponds to a high-noise level; for example, the snrdb will be -9.25 and -18.73 respectively when c = 0.25, 0.4.

Then, we build up the connection graph by setting $V := \mathcal{X}$ and $E := \{(i, j); i, j \in V\}$. Note that in practice, it is common to use a nearest-neighbor scheme to build up the graph, denoted as G^{NN} , for the sake of computational efficiency. However, since the sets of nearest neighbors are not stable under the action of the noise, we also consider here the complete graph scheme, denoted as G. Next we assign the weight function as $w : (i, j) \mapsto e^{-\|x_i - x_j\|_{\mathbb{R}^p}^2/m}$, where m is the 25% quantile of all Euclidean distances between pairs of $(x_i, x_j) \in E$, and the connection function is defined to be a trivial one, that is, r(i, j) = 1 for all $(x_i, x_j) \in E$.

With the connection graph, we build up the CGL matrix (in this 1-dim manifold with the trivial connection, it is equivalent to the graph Laplacian (GL)) from \mathbb{G}^{NN} and \mathbb{G} for comparison, denoted as $L^{NN}(\widetilde{W},\widetilde{G})$ and $L(\widetilde{W},\widetilde{G})$ respectively (see (1)). We have seen in the analysis described earlier in the paper that, when $\alpha < 1$, we have to remove the diagonal terms of the CGL matrix in order to preserve spectral properties. So, we also consider the matrix $L_0(\widetilde{W},\widetilde{G})$ for the comparison.

We then evaluate the eigenvalues and eigenvectors of the above three different CGL's. To simplify the notation, we use the same notations to denote the eigenvectors $u_1, u_2, u_3 \ldots \in \mathbb{R}^n$ associated with the eigenvalues $1 = \lambda_1 > \lambda_2 \ge \lambda_3 \ge \ldots \ge 0$. We now show two sets of results to demonstrate the robustness of the CGL methods studied in this paper.

<u>Dimension Reduction and Data Visualization</u>: To achieve this, we may embed the sampled points into \mathbb{R}^m by the truncated diffusion maps (tDM) with time t > 0 and precision $\delta > 0$:

$$\Phi_{t,m}: x_i \mapsto (\lambda_2^t u_2(i), \lambda_3^t u_3(i), \dots, \lambda_{m+1}^t u_{m+1}(i)) \in \mathbb{R}^m,$$

 $^{^{2}}$ By definition, although locally the manifold resembles Euclidean space near a point, globally it might not. Thus, in general we can only parametrize the manifold locally. This feature captures the possible nonlinear structure in the data.



Figure 1: Clean samples from the twisted bell-shaped manifold. (A): the clean samples. Here we only plot the first 3 axes of the high dimensional data \mathcal{Y} . The color of each point is a surrogate of the norm of each embedded point – blue means a relative small norm and dark red means a relative large norm; the scale above the figure refers to $\{\|x_i\|_2\}_{i=1}^n$, i.e the norm of the data vectors in \mathbb{R}^p .(B): the results of the truncated diffusion maps (tDM), $\Phi_{1,3}$, when the connection graph is \mathbf{G}^{NN} and the diagonal entries are not removed, where the number of nearest neighbors is chosen to be 100; (C): the result of tDM, $\Phi_{1,3}$, when the connection graph is **G** and the diagonal entries are not removed; (D): the result of tDM, $\Phi_{1,3}$, when the connection graph is **G** and the diagonal entries are removed. Note that without surprise, the "parametrization" of the bell shaped manifold is recovered in (B), (C) and (D). For (B), (C), and (D), the scales above the figures refer to the norm of $\{\Phi_{1,3}(x_i)\}_{i=1}^n$; those vectors are of course 3-dimensional, which explains the difference in magnitude of our scales.

where $\lambda_{m+1} > \delta$ and $\lambda_{m+2} \leq \delta$; that is, we map the *i*-th data point to \mathbb{R}^m using the first *m* non-trivial eigenvectors of the CGL. We choose $\delta = 0.2$ in this simulation. The embedding results of \mathcal{Y} , $\Phi_{1,3}$, based on the above different CGL's are shown in Figure 1, and the results from \mathcal{X} with c = 0.4 are shown in Figure 2. Ideally, we would expect to recover the "parametrization" of the dataset by the idea that the eigenvectors of the CGL represent a set of new coordinates for the data points, so the high dimensional dataset can be visualized in this new set of coordinates or its dimension can be reduced. In this specific example, we would expect to find a simple and closed curve out of the noisy dataset which represents the dataset in \mathbb{R}^3 . Clearly when the dataset is clean, we succeed in the task no matter which CGL we use. However, if the dataset is noisy, at high-noise levels, the embedding might not be that meaningful if we use $L^{\text{NN}}(\widetilde{W}, \widetilde{G})$ or $L(\widetilde{W}, \widetilde{G})$. Indeed, as shown in Figure 2, with $L^{\text{NN}}(\widetilde{W}, \widetilde{G})$ the structure of the dataset is barely recovered; with $L(\widetilde{W}, \widetilde{G})$, even though we can get the simple closed curve³ back, there are several outliers which might deteriorate the interpretation. In this noisy case, we can only succeed in the task if we choose $L_0(\widetilde{W}, \widetilde{G})$, as is discussed in this paper.

<u>Nearest Neighbors Estimation</u> Estimating nearest neighbors of a given data point from a noisy dataset is not only important but also challenging in practice (for example, it is essential in the class averaging algorithm for the cryo-EM problem). This problem is directly related to local geodesic distance estimation when the dataset is modeled by the manifold. Their theoretical properties make diffusion maps and vector diffusion maps particularly well-suited for these tasks. To determine the neighbors, we need the notion of distance. In addition to the naive L^2 distance between points, we consider the *diffusion distance* between two points $x_i, x_i \in \mathcal{X}$ by

$$d_{\rm DD}(x_i, x_j) := \|\Phi_{t,m}(x_i) - \Phi_{t,m}(x_j)\|_{\mathbb{R}^m}$$

Then, we determine the nearest neighbors of each data point based on these distances, where we choose t = 1 and $\delta = 0.2$ for the diffusion distance. More precisely, we first determine 10 nearest neighbors of x_i , denoted as x_{i_j} , $j = 1, \ldots, 10$, from the noisy dataset \mathcal{X} , for all i. Then, since we know the ground truth, we may check the true relationship between y_i and y_{i_j} , $j = 1, \ldots, 10$, i.e. $d_{\text{DD}}(y_i, y_{i_j})$ for various CGL methods, or $||y_i - y_{i_j}||$ if we use L^2 distance. Clearly, if the method preserves nearest neighbor information, at least approximately, the ranks of the y_{i_j} 's measured in terms of distances to y_i should be small. To quantify the estimation accuracy, we collect the ranks of all estimated nearest neighbors, and plot the cumulative distribution results in Figure 3. In other words, if we call R_{i_j} the rank of y_{i_j} in terms of distance to y_i , we plot the cdf of $\{\{R_{i_j}\}_{i=1}^n\}_{i=1}^n$ for the various distances we use. (There are many other methods one could

 $^{^{3}}$ The main idea behind tDM is embedding the dataset to a lower dimensional Euclidean space so that the structure underlying the data can be extracted. Please see Section B-5 for details.



Figure 2: Noisy samples from the twisted bell-shaped manifold with $\alpha = 1/4$ and c = 0.4. (A): the noisy samples. Note that we only plot the first 3 axes of the data \mathcal{X} , so the "big" noise seems small, since 997 out of 1000 coordinates are not plotted. To emphasize the relationship among data points, the color of each point is a surrogate of the norm of each embedded point – blue means a relative small norm and dark red means a relative large norm; the scale above the figure refers to $\{||x_i||\}_{i=1}^n$, i.e the norm of our 1000-dimensional vectors. (B): the results of the truncated diffusion maps (tDM), $\Phi_{1,3}$, when the connection graph is \mathbf{G}^{NN} and the diagonal entries are not removed, where the number of nearest neighbors is chosen to be 100. We can barely see the circle structure in the middle, and there are several big outliers; (C): the result of tDM, $\Phi_{1,3}$, when the connection graph is **G** and the diagonal entries are not removed. Note that when compared with (B), the embedding is better in the sense that the "parametrization", the simple and close curve, is better recovered. But we can still observe several outliers; (D): the result of tDM, $\Phi_{1,3}$, when the connection graph is **G** and the diagonal entries are removed. Note that compared with (C), the embedding is yet better in the sense that the number of outliers is reduced and the parametrization of the manifold is recovered. Note that for (B), (C),(D), the scale above the figures refer to $\{||\Phi_{1,3}(x_i)||\}_{i=1}^n$, which are 3-dimensional vectors. The different scales indicate the presence of outliers. Compare also with the scales in Figure 1, (B), (C), (D).

use to do these comparisons, such as using Kendall's τ and variants (see [34]). The one we use here has the benefit of simplicity.) When the dataset is clean, all methods perform the same, as is predicted in Theorem B.7. It is clear from the results that when the noise is large, the result based on the L^2 distance is much worse than the others. The performance based on the diffusion distance from $L^{\text{NN}}(\widetilde{W}, \widetilde{G})$ is better when the noise level is not big, but still a non-negligible portion of error exists; the results based on $L(\widetilde{W}, \widetilde{G})$ and $L_0(\widetilde{W}, \widetilde{G})$ are much better, while the result based on $L_0(\widetilde{W}, \widetilde{G})$ is the best.

In conclusion, in addition to showing the robustness of CGL to noise, we have demonstrated the spectrally close relationship between $L_0(\widetilde{W}, \widetilde{G})$ and L(W, G), which is proved in Proposition 2.3.



Figure 3: The result of nearest neighbors estimation. In all subfigures, the x-axis is the true rank of an estimated nearest neighbor and the y-axis is its cumulative distribution. To emphasize the difference, we only show the area ranging from 90% to 100% in the y-axis. The gray dashed (gray, black dashed and black respectively) curve is the cumulative distribution of the true ranks of the estimated nearest neighbors estimated from the ordinary Euclidean distance (diffusion distance based on $L^{NN}(\widetilde{W}, \widetilde{G})$, $L(\widetilde{W}, \widetilde{G})$ and $L_0(\widetilde{W}, \widetilde{G})$ respectively). From left to right: clean samples from the bell shaped manifold, noisy samples with $\alpha = 1/4$ and c = 0.25, 0.4, 0.5 respectively. It is clear that when the noise is large, the result based on the L^2 distance is much worse than the others. The result based on $L^{NN}(\widetilde{W}, \widetilde{G})$ is slightly better, but not that ideal, $L(\widetilde{W}, \widetilde{G})$ is even better and $L_0(\widetilde{W}, \widetilde{G})$ is the best.

3.2 2-dim images

In Subsection 3.1, we investigated numerically the influence of noise on CGL methods when the connection function is trivial. In this subsection, we discuss an example where the connection function plays an essential role in the analysis. We consider a dataset which contains randomly rotated versions of a set of objects, and the task is to align these objects in addition to classifying them. We encounter this kind of datasets and problems in, for example, image processing [59, 62, 71], shape analysis [40], phase retrieval problems [2, 45], etc. In [2, 45, 59, 62, 71] and others, the CGL methods have been applied to solve the problem.

To focus specifically on demonstrating the influence of noise on this problem, we work with 2-dimensional images observed in polar coordinates. If an image is defined in Cartesian coordinates, then in general a numerical rotation will introduce numerical artifacts and errors since resampling or interpolation procedures are then involved to compare two rotated images. These numerical issues are alleviated if we work with polar coordinates. To further minimize these numerical artifacts, we use surrogate images as our dataset – by a surrogate image, we mean a function defined on the circle S^1 , which is discretized into p equally spaced points. In other words, we consider images defined in polar coordinates, where we only have one sample on the radial axis.

Now we discretize the 2×2 rotational group, SO(2), which is the same as the circle S^1 , into p equally spaced points, that is, $\mathcal{T}^{(2)} := \{e^{i2\pi k/p}\}_{k=1}^p$ - the sample $e^{i2\pi k/p} \in \mathcal{T}^{(2)}$ simply rotates vectors in \mathbb{R}^2 by an angle $2\pi k/p$. Note that since the surrogate images are defined on p equally spaced points on S^1 , the rotations in $\mathcal{T}^{(2)}$ act exactly on the images without introducing any numerical error. We choose n_K different surrogate images, denoted as $\{f_i\}_{i=1}^{n_K} \subset \mathbb{R}^p$. Then we randomly and independently rotate each of them by n_R angles; that is, for all $k = 1, \ldots, n_K$ and $l = 1, \ldots, n_R$, we have $S_i := R_{k,l} \circ f_k$, where $R_{k,l} \in \mathcal{T}^{(2)}$, $R_{k,l} \circ f_k$ means rotating f_k by $R_{k,l}$ and $i = (k-1)n_R + l$. We assume that $\operatorname{argmin}_{R \in \mathcal{T}^{(2)}} ||f_i - R \circ f_j|| > 0$, for all $i, j = 1, \ldots, n_K$; that is, the image f_i is not a rotated version of another one f_j . In the end we get $n = n_K n_R$ randomly rotated images $\{S_i\}_{i=1}^n \subset \mathbb{R}^p$. Denote by σ the standard deviation of all pixels of all images in $\{f_i\}_{i=1}^{n_K}$. The data $\mathcal{X} = \{I_i\}_{i=1}^n$ we analyze is the clean data contaminated by the noise which is i.i.d. sampled from Z, that is, we have $I_i = S_i + Z_i$.

We now build up the connection graph by setting $\mathbf{V} := \{I_i\}_{i=1}^n$ and $\mathbf{E} := \{(i,j); I_i, I_j \in \mathbf{V}\}$; that is, we take the complete graph scheme. Next we assign the weight function as $w : (i,j) \mapsto e^{-d_{\mathrm{RID}}^2(I_i,I_j)/m}$, where m is the 25% quantile of all non-zero RID distances defined on \mathbf{E} , and the connection function as $r : (i,j) \mapsto \operatorname{argmin}_{R \in \mathcal{T}^{(2)}} ||I_i - R \circ I_j||$. For comparison purposes, we also take the nearest neighbor scheme to construct the connection graph, denoted by $(\mathbf{G}^{\mathrm{NN}}, w^{\mathrm{NN}}, r^{\mathrm{NN}})$, where we choose 100 nearest neighbors as defined by the RID distance - to construct edges. When the images are noise-free, due to the connection function, we can recover $R_{k,l}$'s up to a rotation from the top eigenvector v_1 of different CGL's built up from different connection graphs, (\mathbf{G}, w, r) or $(\mathbf{G}^{\mathrm{NN}}, w^{\mathrm{NN}}, r^{\mathrm{NN}})$, with or without removing the diagonal entries. To simplify the notation, we will use the same notation v_1 to denote the top eigenvector of the different CGL's. Precisely, the estimated rotation is built up from v_1 , denoted as $v \in \mathbb{C}^n$, by setting $v(i) = \frac{v_1(i)}{|v_1(i)|}$ when $|v_1(i)| > 0$ and v(i) = 1 when $|v_1(i)| = 0$. (In a slight departure from the descriptions given earlier in the paper, the $r_{i,j}$'s are not 2×2 matrices here, but complex numbers describing the corresponding rotations. Hence, v_1 is in \mathbb{C}^n . If we had used 2×2 matrices, v_1 would have been in \mathbb{R}^{2n} and we could have computed the vector v by using pairs of entries of v_1 .)

To evaluate the performance of the estimated rotation when noise exists, we construct a complex vector $u \in \mathbb{C}^n$ whose *i*-th entry - where $i = (k-1)n_R + l$, $k = 1, \ldots, n_K$ and $l = 1, \ldots, n_R$ - is the complex form of $R_{k,l}$. We then evaluate the difference between the estimated rotation of the *l*-th object and the ground truth by observing the angle of $u(i)^*v(i)$. In other words, this quantity shows the discrepancy between the true rotation and the estimated rotation. To visualize this result, we plot the vector $z \in \mathbb{R}^n$ where z(i) is the angle of the complex number $u(i)^*v(i)$. In Figure 4, the resulting z's with p = 1000, $n_K = 5$, $n_R = 200$, $\alpha = 1/4$ and $c = 6\sigma$ are illustrated. Note that since there are 5 different images, we see a piecewise function with 5 different values when the images are clean. When noise exists, we can see clearly the benefit of removing the diagonal entries (see Figure 4, (H)).

We mention that depending on the problem, the affinity function and the connection function are constructed in different ways (see, for example [2, 40, 45, 59, 62, 71]). Also, the CGL is only one of several



Figure 4: (A): a clean surrogate image. (B)-(D): alignment vectors z computed from clean images; (E): a noisy surrogate image. (F)-(H): alignment vectors z computed from noisy images with $c = 6\sigma$. (A) and (E): the black curve is a clean surrogate image, and the gray curve is its noisy version; (B) and (F): the result from the CGL built up from G^{NN} ; (C) and (G): the result from the CGL built from G and the diagonal entries are not removed; (D) and (H): the result from the CGL built from G with the diagonal entries removed. It is clear that when the images are clean, all different CGL's give equivalent results. But in the presence of noise, the CGL built up from G^{NN} is obviously worse.

techniques we could use to analyze datasets described by the connection graph. We might consider other techniques, such as, semidefinite programming relaxation [69], nonlinear independent component analysis [66], and other methods, to obtain the information we are interested in, reorganize the data, etc...

References

- R. Al-Aifari, I. Daubechies, and Y. Lipman. Continuous Procrustes Distance Between Two Surfaces. Comm. Pure Appl. Math., 66:934–964, 2013.
- [2] B. Alexeev, A. S. Bandeira, M. Fickus, and D. G. Mixon. Phase retrieval with polarization. SIAM J. Imaging Sci., 2013.
- [3] A. S. Bandeira, A. Singer, and D. A. Spielman. A Cheeger Inequality for the Graph Connection Laplacian. SIAM Journal on Matrix Analysis and Applications, to appear, 2013. arXiv:1204.3873 [math.SP].
- [4] T. Batard and N. Sochen. Polyakov action on (\rho, G) -equivariant functions Application to color image regularization. In Scale Space and Variational Methods in Computer Vision, pages 483–494. 2012.
- [5] T. Batard and N. Sochen. A class of generalized laplacians on vector bundles devoted to multi-channel image processing. J. Math. Imaging Vis., 48(3):517–543, 2014.
- [6] J. Bates. The embedding dimension of laplacian eigenfunction maps. Appl. Comput. Harmon. Anal., (0):-, 2014.
- [7] M. Belkin and P. Niyogi. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural. Comput.*, 15(6):1373–1396, June 2003.
- [8] M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. In Proceedings of the 18th Conference on Learning Theory (COLT), pages 486–500, 2005.
- [9] M. Belkin and P. Niyogi. Convergence of laplacian eigenmaps. In Advances in Neural Information Processing Systems 19: Proceedings of the 2006 Conference, volume 19, page 129. The MIT Press, 2007.
- [10] P. Bérard. Spectral Geometry: Direct and Inverse Problems. Springer, 1986.
- [11] P. Bérard, G. Besson, and S. Gallot. Embedding riemannian manifolds by their heat kernel. Geom. Funct. Anal., 4:373–398, 1994. 10.1007/BF01896401.
- [12] N. Berline, E. Getzler, and M. Vergne. Heat Kernels and Dirac Operators. Springer, 2004.

- [13] Rajendra Bhatia. Matrix analysis, volume 169 of Graduate Texts in Mathematics. Springer-Verlag, New York, 1997.
- [14] R. L. Bishop and R. J. Crittenden. Geometry of Manifolds. Amer Mathematical Society, 2001.
- [15] D. M. Boyer, Y. Lipman, E. St. Clair, J. Puente, B. A. Patel, T. Funkhouser, J. Jernvall, and I. Daubechies. Algorithms to automatically quantify the geometric similarity of anatomical surfaces. *P. Natl. Acad. Sci. USA*, 108(45):18221–18226, 2011.
- [16] P. Chen, C. Lin, and I. Chern. A perfect match condition for point-set matching problems using the optimal mass transport approach. SIAM J. Imaging Sci., 6(2):730–764, 2013.
- [17] M.-Y. Cheng and H.-T. Wu. Local linear regression on manifolds and its geometric interpretation. J. Am. Stat. Assoc., 108:1421–1434, 2013.
- [18] F. Chung. Spectral Graph Theory. American Mathematical Society, 1996.
- [19] F. Chung and M. Kempton. A local clustering algorithm for connection graphs. In Anthony Bonato, Michael Mitzenmacher, and Pawel Pralat, editors, *Algorithms and Models for the Web Graph*, volume 8305 of *Lecture Notes in Computer Science*, pages 26–43. Springer International Publishing, 2013.
- [20] F. Chung, W. Zhao, and M. Kempton. Ranking and sparsifying a connection graph. In Anthony Bonato and Jeannette Janssen, editors, *Algorithms and Models for the Web Graph*, volume 7323 of *Lecture Notes in Computer Science*, pages 66–77. Springer Berlin Heidelberg, 2012.
- [21] R. R. Coifman and S. Lafon. Diffusion maps. Appl. Comput. Harmon. Anal., 21(1):5–30, 2006.
- [22] R. R. Coifman and M. Maggioni. Diffusion wavelets. Appl. Comput. Harmon. Anal., 21(1):53 94, 2006.
- [23] A. Collins, A. Zomorodian, G. Carlsson, and L. J. Guibas. A barcode shape descriptor for curve point cloud data. *Computers & Graphics*, 28(6):881 – 894, 2004.
- [24] M. Cucuringu, Y. Lipman, and A. Singer. Sensor network localization by eigenvector synchronization over the euclidean group. ACM Transactions on Sensor Networks, 8(3):19:1–19:42, 2012.
- [25] M. Cucuringu, A. Singer, and D. Cowburn. Eigenvector synchronization, graph rigidity and the molecule problem. *Information and Inference: A Journal of the IMA*, 1:21–67, 2012.
- [26] Persi Diaconis and David Freedman. Asymptotics of graphical projection pursuit. Ann. Statist., 12(3):793-815, 1984.
- [27] M.P. do Carmo and F. Flaherty. *Riemannian Geometry*. Birkhauser Boston, 1992.
- [28] N. El Karoui. Operator norm consistent estimation of large dimensional sparse covariance matrices. The Annals of Statistics, 36(6):2717–2756, December 2008.
- [29] N. El Karoui. Concentration of measure and spectra of random matrices: Applications to correlation matrices, elliptical distributions and beyond. *The Annals of Applied Probability*, 19(6):2362–2405, December 2009.
- [30] N. El Karoui. On information plus noise kernel random matrices. Ann. Statist., 38(5):3191–3216, 2010.
- [31] N. El Karoui and H.-T. Wu. Vector diffusion maps and random matrices with random blocks. *ArXiv e-prints*, 2013. arXiv:1310.0188 [math.PR].
- [32] Noureddine El Karoui, Derek Bean, Peter J. Bickel, Chinghway Lim, and Bin Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 2013.

- [33] C. Epstein. Introduction to the Mathematics of Medical Imaging. SIAM, second edition, 2007.
- [34] R. Fagin, R. Kumar, and D. Sivakumar. Comparing top k lists. SIAM Journal on Discrete Mathematics, 17(1):134–160, 2003.
- [35] D. Giannakis, P. Schwander, and A. Ourmazd. The symmetries of image formation by scattering. i. theoretical framework. *Opt. Express*, 20(12):12799–12826, Jun 2012.
- [36] P. Gilkey. The Index Theorem and the Heat Equation. Princeton, 1974.
- [37] R. Hadani and A. Singer. Representation theoretic patterns in three dimensional Cryo-Electron Microscopy I: The intrinsic reconstitution algorithm. Annals of Mathematics, 174(2):1219–1241, 2011.
- [38] P. Hall, J. S. Marron, and A. Neeman. Geometric representation of high dimension, low sample size data. J. R. Stat. Soc. Ser. B Stat. Methodol., 67(3):427–444, 2005.
- [39] M. Hein, J. Audibert, and U. von Luxburg. From graphs to manifolds weak and strong pointwise consistency of graph Laplacians. In *Proceedings of the 18th Conference on Learning Theory (COLT)*, pages 470–485, 2005.
- [40] Q.-X. Huang, H. Su, and L. Guibas. Fine-grained semi-supervised labeling of large shape collections. ACM Transactions on Graphics (TOG), 32(6):190, 2013.
- [41] P. W. Jones, M. Maggioni, and R. Schul. Manifold parametrizations by eigenfunctions of the Laplacian and heat kernels. P. Natl. Acad. Sci. USA, 105(6):1803–8, February 2008.
- [42] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. Ann. Statist., 28(5):1302–1338, 2000.
- [43] M. Ledoux. The concentration of measure phenomenon, volume 89 of Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2001.
- [44] R. Lee, J. S. Oveis Gharan, and L. Trevisan. Multi-way spectral partitioning and higher-order cheeger inequalities. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 1117–1130, 2012.
- [45] S. Marchesini, Y.-C. Tu, and H.-T. Wu. Alternating Projection, Ptychographic Imaging and Phase Synchronization. ArXiv e-prints, 2014. arXiv:1402.0550 [math.OC].
- [46] F. Mémoli. A spectral notion of gromov-wasserstein distance and related methods. Appl. Comput. Harmon. Anal., 30(3):363 – 401, 2011.
- [47] F. Mémoli and G. Sapiro. A Theoretical and Computational Framework for Isometry Invariant Recognition of Point Cloud Data. Found. Comput. Math., 5:313–347, 2005.
- [48] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis. Diffusion maps, spectral clustering and eigenfunctions of Fokker-Planck operators. In Y. Weiss, B. Schölkopf, and J. Platt, editors, Adv. Neur. In., volume 18, pages 955–962, Cambridge, MA, 2006. MIT Press.
- [49] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis. Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. Appl. Comput. Harmon. Anal., 21(1):113–127, 2006.
- [50] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. In *Twentieth Anniversary Volume:*, pages 1–23. Springer New York, 2009.
- [51] M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. In Proceedings of the Symposium on Geometry Processing, SGP '08, pages 1341–1348. Eurographics Association, 2008.
- [52] J. W. Portegies. Embeddings of Riemannian manifolds with heat kernels and eigenfunctions. ArXiv e-prints, 2013. arXiv:1311.7568 [math.DG].

- [53] H. Qiu and E.R. Hancock. Clustering and embedding using commute times. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 29(11):1873–1890, Nov 2007.
- [54] M. Reuter. Hierarchical Shape Segmentation and Registration via Topological Features of Laplace-Beltrami Eigenfunctions. Int. J. Comput. Vision, 89:287–308, 2010.
- [55] R. M. Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In Proceedings of the fifth Eurographics symposium on Geometry processing, SGP '07, pages 225–233. Eurographics Association, 2007.
- [56] A. Singer. From graph to manifold Laplacian: The convergence rate. Appl. Comput. Harmon. Anal., 21(1):128–134, 2006.
- [57] A. Singer. Angular Synchronization by Eigenvectors and Semidefinite Programming. Appl. Comput. Harmon. Anal., 30(1):20–36, 2011.
- [58] A. Singer and R. R. Coifman. Non-linear independent component analysis with diffusion maps. *Appl. Comput. Harmon. Anal.*, 25(2):226 239, 2008.
- [59] A. Singer and H.-T. Wu. Vector diffusion maps and the connection Laplacian. Comm. Pure Appl. Math., 65(8):1067–1144, 2012.
- [60] A. Singer and H.-T. Wu. 2-d tomography from noisy projections taken at unknown random directions. SIAM J. Imaging Sci., 6(1):136–175, 2013.
- [61] A. Singer and H.-T. Wu. Spectral convergence of the connection laplacian from random samples. *submitted*, 2013.
- [62] A. Singer, Zhao Z., Shkolnisky Y., and Hadani R. Viewing angle classification of cryo-electron microscopy images using eigenvectors. SIAM J. Imaging Sci., 4(2):723–759, 2011.
- [63] G. W. Stewart and Ji Guang Sun. Matrix perturbation theory. Computer Science and Scientific Computing. Academic Press Inc., Boston, MA, 1990.
- [64] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proceedings of the Symposium on Geometry Processing*, SGP '09, pages 1383–1392. Eurographics Association, 2009.
- [65] A. D. Szlam, M. Maggioni, and R. R. Coifman. Regularization on Graphs with Function-adapted Diffusion Processes. The Journal of Machine Learning Research, 9:1711–1739, 2008.
- [66] R. Talmon, I. Cohen, S. Gannot, and R. Coifman. Diffusion maps for signal processing: A deeper look at manifold-learning techniques based on kernels and graphs. *Signal Processing Magazine*, *IEEE*, 30(4):75–86, July 2013.
- [67] U. von Luxburg, M. Belkin, and O. Bousquet. Consistency of spectral clustering. Ann. Stat., 36(2):555– 586, 2008.
- [68] F. Wang, Q. Huang, and L. J. Guibas. Image co-segmentation via consistent functional maps. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2013.
- [69] L. Wang and A. Singer. Exact and Stable Recovery of Rotations for Robust Synchronization. Information and Inference: A Journal of the IMA, 2013. accepted for publication.
- [70] H.-T. Wu. Embedding riemannian manifolds by the heat kernel of the connection laplacian. *submitted*, 2012.
- [71] Z. Zhao and A. Singer. Rotationally Invariant Image Representation for Viewing Direction Classification in Cryo-EM. ArXiv e-prints, 2013. arXiv:1309.7643 [math.CV].

APPENDIX

to "Connection graph Laplacian methods can be made robust to noise"

A Technical results

A-1 On quadratic forms

Lemma A-1. Suppose Z_1, \ldots, Z_n are random vectors in \mathbb{R}^p , with $Z_i = \sum_i^{1/2} X_i$, where X_i has mean 0 and covariance Id_p. We further assume that for every convex 1-Lipschitz function f, if $m_{f(X_i)}$ is a median of $f(X_i)$, $P(|f(X_i) - m_{f(X_i)}| > t) \leq 2 \exp(-c_i t^2)$. Z_i 's are possibly dependent. Let $\{Q_i\}_{i=1}^n$ be $p \times p$ positive definite matrices. Call $|||Q_i|||_2$ the largest eigenvalue of Q_i . Then we have

$$\sup_{1 \le i \le n} \left| \sqrt{Z_i' Q_i Z_i} - \mathbf{E} \left(\sqrt{Z_i' Q_i Z_i} \right) \right| = \mathcal{O}_P(\sup_i \sqrt{||Q_i \Sigma_i / c_i|||_2} \sqrt{\log n})$$

This implies that, when $\sup_i \sqrt{|||Q_i \Sigma_i / c_i|||_2} \sqrt{\log n} \to 0$,

$$\sup_{1 \le i \le n} \left| Z_i' Q_i Z_i - trace(\Sigma_i Q_i) \right| = \mathcal{O}_P(\sup_i \sqrt{||Q_i \Sigma_i / c_i||_2} \sqrt{\log n} \left| \sup_i \sqrt{trace(\Sigma_i Q_i)} \lor 1 \right|).$$

As explained in [43], the condition we require on X_i is satisfied by many distributions. We refer also to [29] for many examples. Here are two examples. The Gaussian distribution in dimension p satisfies the previous assumptions with $c_i = 1/2$, independently of the dimension. When X_i 's have independent coordinates supported on intervals of width at most B_i , c_i is proportional to $1/B_i$.

Proof. The map $X_i \to \sqrt{X_i' \Sigma_i^{1/2} Q_i \Sigma_i^{1/2} X_i}$ is convex and $\sqrt{|||Q_i \Sigma_i|||_2} = \sqrt{|||\Sigma_i^{1/2} Q_i \Sigma_i^{1/2}|||_2}$ -Lipschitz as a function of X_i . Indeed, it is a norm, which gives convexity. The Lipschitz-property comes from the triangle inequality. Hence, under our assumptions, since $Z_i = \Sigma_i^{1/2} X_i$, we have

$$P\left(\left|\sqrt{Z_i'Q_iZ_i} - \mathbf{E}\left(\sqrt{Z_i'Q_iZ_i}\right)\right| > t\right) \le 2\exp(-c_it^2/(|||Q_i\Sigma_i|||_2))$$

By a simple union bound, we get

$$P\left(\sup_{1\le i\le n} \left|\sqrt{Z_i'Q_iZ_i} - \mathbf{E}\left(\sqrt{Z_i'Q_iZ_i}\right)\right| > t\right) \le 2\sum_{i=1}^n \exp(-c_it^2/[||Q_i\Sigma_i||_2]) \le 2n\exp(-t^2/(\sup_i||Q_i\Sigma_i/c_i||_2))$$

Taking $t_K = K\sqrt{\log(n)||Q_i\Sigma_i/c_i||_2}$, for K a constant, gives the first result. The second result follows from remarking that $|a^2 - b^2| = |a - b||a + b| \le |a - b|^2 + 2|b||a - b|$. When $\sup_i \sqrt{||Q_i\Sigma_i/c_i||_2}\sqrt{\log n} \to 0$, this gives immediately

$$\sup_{1 \le i \le n} \left| Z_i' Q_i Z_i - \left[\mathbf{E} \left(\sqrt{Z_i' Q_i Z_i} \right) \right]^2 \right| = \mathcal{O}_P(\sup_i \sqrt{||Q_i \Sigma_i / c_i|||_2} \sqrt{\log n} \sup_i \left[\sqrt{\operatorname{trace}\left(\Sigma_i Q_i\right)} \lor 1 \right]) ,$$

after we notice that trace $(\Sigma_i Q_i) = \mathbf{E} (Z'_i Q_i Z_i) \ge [\mathbf{E} (\sqrt{Z'_i Q_i Z_i})]^2$. Finally, using the variance bound in Proposition 1.9 of [43], we see that,

$$\mathbf{E}\left(Z_i'Q_iZ_i\right) - \left[\mathbf{E}\left(\sqrt{Z_i'Q_iZ_i}\right)\right]^2 \le 2|||\Sigma_iQ_i/c_i|||_2$$

Under our assumption that $\sup_i \sqrt{||Q_i \Sigma_i / c_i||_2} \sqrt{\log n} \to 0$, we have $\sup_i \sqrt{||Q_i \Sigma_i / c_i||_2} \to 0$ and therefore

$$\sup_{i} |||\Sigma_{i}Q_{i}/c_{i}|||_{2} = o(\sup_{i} \sqrt{|||\Sigma_{i}Q_{i}/c_{i}|||_{2}})$$

This gives the second bound.

In the case of the Gaussian distribution, the previous bounds can be improved, using an observation found in [42].

Lemma A-2. Suppose Z_1, \ldots, Z_n are random vectors in \mathbb{R}^p , with $Z_i \sim \mathcal{N}(0, \Sigma_i)$. Z_i 's are possibly dependent. Let $\{Q_i\}_{i=1}^n$ be $p \times p$ positive definite matrices. Then we have, if $S_i = \Sigma_i^{1/2} Q_i \Sigma_i^{1/2}$

$$\sup_{1 \le i \le n} \left| Z_i' Q_i Z_i - trace(S_i) \right| = \mathcal{O}_P\left(\sup_{1 \le i \le n} \sqrt{\log(n)} \sqrt{trace(S_i^2)} + |||S_i|||_2 \log(n) \right) \ .$$

Proof. By rotational invariance of the Gaussian distribution, we have

$$W_i \triangleq Z'_i Q_i Z_i - \operatorname{trace} \left(S_i \right) \stackrel{\mathcal{L}}{=} \sum_{k=1}^p \lambda_k(S_i) (X_k^2 - 1) ,$$

where X_k 's are i.i.d $\mathcal{N}(0, 1)$. Using Lemma 1, p, 1325 in [42], we see that

$$P\left(\frac{|W_i|}{2} > \sqrt{\operatorname{trace}\left(S_i^2\right)}\sqrt{x} + |||S_i|||_2x\right) \le \exp(-x) \ .$$

Taking $x = K \log(n)$ in the previous inequality and a simple union bound gives the announced result. \Box

A-2 Proof of Lemma 2.1

Proof. We have $L(W,G) = D^{-1}S = (D/n)^{-1}(S/n)$. If we call $d_{i,i} = \sum_{j \neq i} \tilde{w}_{i,j}$ and $\tilde{d}_{i,i} = \sum_{j \neq i} \tilde{w}_{i,j}$, we see that

$$|d_{i,i}/n - d_{i,i}/n| \le \sup_{j \ne i} |w_{i,j} - \tilde{w}_{i,j}|.$$

Hence,

$$\sup_{1 \le i \le n} |d_{i,i}/n - \tilde{d}_{i,i}/n| \le \sup_{1 \le i \le n} \sup_{j \ne i} |w_{i,j} - \tilde{w}_{i,j}| \le \sup_{i,j} |w_{i,j} - \tilde{w}_{i,j}| \le \epsilon .$$

We conclude that

$$|||D/n - \widetilde{D}/n|||_2 \le \epsilon$$
.

Under our assumptions, it is clear that $|||(D/n)^{-1}|||_2 \leq 1/\gamma$. The previous display also implies that $|||(\widetilde{D}/n)^{-1}|||_2 \leq 1/(\gamma - \epsilon)$.

Furthermore, since

$$(D/n)^{-1} - (\widetilde{D}/n)^{-1} = (D/n)^{-1} [D/n - \widetilde{D}/n] (\widetilde{D}/n)^{-1}$$

we see that

$$|||(D/n)^{-1} - (\widetilde{D}/n)^{-1}|||_2 \le \frac{\epsilon}{\gamma(\gamma - \epsilon)}$$

Also,

$$||S/n - \widetilde{S}/n||_F^2 \le \sup_{i,j} ||S_{i,j} - \widetilde{S}_{i,j}||_F^2$$
.

Naturally, since $S_{i,j} = w_{i,j}G_{i,j}$ and $\widetilde{S}_{i,j} = \widetilde{w}_{i,j}\widetilde{G}_{i,j}$,

$$\|S_{i,j} - \widetilde{S}_{i,j}\|_F^2 \le |w_{i,j}|^2 \|G_{i,j} - \widetilde{G}_{i,j}\|_F^2 + |w_{i,j} - \widetilde{w}_{i,j}|^2 \|\widetilde{G}_{i,j}\|_F^2 \le C^2(\eta^2 + \epsilon^2)$$

Hence,

$$|||S/n - \tilde{S}/n|||_2 \le ||S/n - \tilde{S}/n||_F \le C(\eta + \epsilon).$$

We also note that $\|\widetilde{S}/n\|_F \leq C^2$. So we can conclude that

$$|||D^{-1}S - \widetilde{D}^{-1}\widetilde{S}|||_{2} \le |||D^{-1}(S - \widetilde{S}) + (D^{-1} - \widetilde{D}^{-1})\widetilde{S}|||_{2} \le \frac{1}{\gamma}C(\eta + \epsilon) + \frac{\epsilon}{\gamma(\gamma - \epsilon)}C^{2}.$$

A-3 Card $\{\mathcal{T}\}$: an example when $\mathcal{T}_{exact}^{(k)} \subset SO(k)$

Naturally, when working with discretized images/objects with p pixels/voxels, we need to also discretize SO(k). In light of results like Proposition 2.1, one natural question we have to deal with concerns the cardinality of the discretized set of transformations, $\mathcal{T}^{(k)}$, and the corresponding set for companion matrices, \mathcal{T} . The following proposition answers this question.

The images/objects are discretized in polar coordinates. In other words, each point on our grid can be identified by its location on a ray emanating from the origin and reaching a point p on the sphere of radius r_0 centered at the origin. The discretization of each ray does not have to be uniform. But this discretization is the same for all rays.

In the case of SO(2), this simply means that we discretize the circle of radius r_0 , and our points lay on the corresponding radii. In this situation, it is natural to represent each point on our grid through (r, θ) . To give a concrete example, we assume that $\theta \in \{2\pi \frac{k}{M}\}_{k=0}^{M-1}$ and $r \in \{\mathbf{r}_1, \ldots, \mathbf{r}_{\alpha}\}$ with $\mathbf{r}_{\alpha} = r_0$. The discretization of SO(2) corresponds simply to rotations by an angle θ_k , where $\theta_k = 2\pi \frac{k}{M}$. These rotations clearly map our grid onto itself.

We assume that our images or objects, after having been uniformly discretized in polar coordinates, fit in a k-dimensional cube. Then we assume that the rotation group is properly discretized so that each rotation is exact in the sense that it commutes with discretization. In other words, the rotation does not change the pixel values - pixels are simply swapped and pixel values are not averaged or aggregated in other ways. Note that when the image or object is discretized in Cartesian coordinates, the discretization and rotation will not commute and a distortion is inevitable. We postpone the study of such a phenomenon to future work.

Proposition A.1. When $\mathcal{T}_{exact}^{(k)}$ is the discretized version of SO(k) we just discussed, we have

$$Card\left\{\mathcal{T}_{exact}^{(k)}\right\} = Card\left\{\mathcal{T}_{exact}\right\} = O(p^{k-1})$$
.

Furthermore, the elements of \mathcal{T}_{exact} are permutation matrices. In particular, they are orthogonal matrices.

Comment : Proposition A.1 shows that in checking Assumption **G1**, we can assume that $\operatorname{Card} \{\mathcal{T}\}$ is polynomial in p. This implies that Assumption G1 will be satisfied when $\max(\sigma_p, \sqrt{p}s_p^2) = o([\log(np)]^{-1/2})$. Hence the conditions we will have to check on σ_p and $\sqrt{p}s_p^2$ will be quite unrestrictive and we will see that this implies that CGL algorithms are robust to considerable amount of additive noise.

Proof. Our polar-coordinate discretization amounts to discretizing a sphere of radius r_0 in \mathbb{R}^k with M points and discretizing each ray linking a point on that sphere to the origin along α points. We naturally have the relationship $M\alpha = p$.

Now elements of $\mathcal{T}_{\text{exact}}^{(k)}$ are orthogonal matrices with determinant 1, hence they can be characterized by their action on k-1 vectors in \mathbb{R}^k which span a subspace of dimension k-1.

Let us pick k-1 elements among our M points on the sphere of radius r_0 . We require that these k-1 elements span a subspace of dimension k-1 in \mathbb{R}^k . We call the corresponding vectors v_1, \ldots, v_{k-1} .

Suppose now that $\mathbf{0} \in \mathcal{T}_{exact}^{(k)}$. Then, for each i, $\mathbf{0}v_i$ has to be one of the elements of our discretized sphere. Therefore,

Card
$$\left\{ \mathcal{T}_{\text{exact}}^{(k)} \right\} \le M^{k-1} = \left(\frac{p}{\alpha}\right)^{k-1} \le p^{k-1}$$
.

Now let $\mathbf{O} \in \mathcal{T}_{exact}^{(k)}$ and let O be the companion matrix of \mathbf{O} . Note that if I_i is our image and $I_i^{\vee} \in \mathbb{R}^p$ is its discretized version, OI_i^{\vee} swaps the position of the entries of the vector I_i^{\vee} , since \mathbf{O} maps our grid onto itself. Hence O is a permutation matrix and it is therefore orthogonal.

Another approach We note that another approach can be employed to generate a sampling grid \mathfrak{X} and an associated set of exact transformations for $k \geq 3$. Take *m* points in \mathbb{R}^k denoted as $X = \{x_i\}_{i=1}^m$. Take a finite subgroup, *T* of SO(k). Now consider the sampling grid $\mathfrak{X} \triangleq \{Rx_i; R \in T, x_i \in X\}$. Since *T* is a subgroup, we know that the sampling grid \mathfrak{X} is of finite size; that is, $|\mathfrak{X}| \leq m|T|$. It is also clear that *T* is an exact set of transforms for \mathfrak{X} , by simply using the fact that *T* is a group. Note that the standard polar coordinate grid in \mathbb{R}^2 described above can be viewed as an instance of the method we just discussed, with T consisting of powers of the rotation by the angle $\frac{2\pi}{M}$.

Note however that the classification of finite subgroups of SO(k), for $k \ge 3$, imposes strong constraints on the sampling grids obtained by such a construction.

B Background on CGL methods

In this section, we discuss the *noise-free* connection graph Laplacian (CGL) L(W, G) defined in (1), understand its asymptotical behavior under the assumption that the point clouds we collect are distributed on a manifold, and show that the CGL matrix built up under this assumption enjoy a sparsity property which allows the robustness result shown in this paper. In addition, we will discuss the fact that the CGL matrix can be viewed as a generalization of the graph Laplacian (GL) [61]. We will see that although GL and CGL share several similar properties but are fundamentally different.

We would assume the background knowledge of differential geometry in the following discussion. For a reader who is not familiar with the subject, we refer him to [10, 12, 14, 27, 36] for the topics we will encounter.

We start from some notations. Denote M to be a d-dimensional compact, connected and smooth Riemannian manifold embedded in \mathbb{R}^p via ι , where $d \leq p$. Denote the tangent bundle as TM. The tangent plane at $y \in M$ is denoted as T_yM . Introduce the metric g on M induced from the canonical metric of the ambient space \mathbb{R}^p . Denote d(y, y') to be the geodesic distance between $y, y' \in \mathcal{M}$. Denote by ∇ the covariant derivative of the vector field, Δ_g the Laplace-Beltrami operator, ∇^2 the connection Laplacian of the tangent bundle associated with the Levi-Civita connection, and by Ric the Ricci curvature of (M, g). We denote the spectrum of ∇^2 (resp. Δ_g) by $\{-\lambda_l\}_{l=0}^{\infty}$ (resp. $\{-\gamma_l\}_{l=0}^{\infty}$), where $0 = \lambda_0 \leq \lambda_1 \leq \ldots$ (resp. $0 = \gamma_0 < \gamma_1 \leq \ldots$), and the corresponding eigenspaces by $F_l := \{X \in L^2(TM) : \nabla^2 X = -\lambda_l X\}$ (resp. $E_l := \{\phi \in L^2(M) : \Delta_g \phi = -\gamma_l \phi\}, l = 0, 1, \dots$ In general, while $\gamma_0 = 0$ always exists, λ_0 may not: a simple example is found considering S^2 with the standard metric. It is well known [36] that dim $(F_l) < \infty$, the eigen-vector-fields are smooth and form a basis for $L^2(TM)$ (resp. dim $(E_l) < \infty$, the eigenfunctions are smooth and form a basis for $L^2(M)$, that is, $L^2(TM) = \overline{\bigoplus_{l \in \mathbb{N} \cup \{0\}} F_l}$ (resp. $L^2(M) = \overline{\bigoplus_{l \in \mathbb{N} \cup \{0\}} E_l}$), the completion of $\bigoplus_{l \in \mathbb{N} \cup \{0\}} F_l$ with relative to the measure induced by g. To simplify the statement, we assume that λ_l (resp. γ_l) for each l are simple and X_l (resp. ϕ_l) is a normalized basis of F_l (resp. E_l). Denote $\mathcal{B}(F_k)$ (resp. $\mathcal{B}(E_k)$) the set of bases of F_k (resp. E_k), which is identical to the orthogonal group $O(\dim(F_k))$ (resp. $O(\dim(E_k))$). Denote the set of the corresponding orthonormal bases of $L^2(TM)$ by $\mathcal{B}(TM,g) = \prod_{k=1}^{\infty} \mathcal{B}(F_k) \text{ (resp. } \mathcal{B}(M,g) = \prod_{k=1}^{\infty} \mathcal{B}(E_k) \text{)}.$

Given the collected data $\mathcal{X} = \{x_i\}_{i=1}^n \subset \mathbb{R}^p$, where x_i are signal random vector i.i.d. sampled from a random vector X. We assume a manifold structure inside the signal random vector; that is, we view $X : \Omega \to \mathbb{R}^p$ as a measurable function with respect to the probability space (Ω, \mathcal{F}, P) , and assume that its range is $\iota(M)$. Note that we cannot define the probability density function (p.d.f.) of X on the ambient space \mathbb{R}^p when d < p since $\iota(M)$ is degenerate, but we can still discuss how we sample points from $\iota(M)$, which leads to the p.d.f. of X on M. Indeed, we employ the following definition which is based on the induced measure [17, Section 4]. Denote $\tilde{\mathcal{B}}$ to be the Borel sigma algebra on $\iota(M)$, and \tilde{P}_X the probability measure of X, defined on $\tilde{\mathcal{B}}$, induced from P. Assume that \tilde{P}_X is absolutely continuous with respect to the volume density on $\iota(M)$ associated with g, that is, $d\tilde{P}_X(x) = f(\iota^{-1}(x))\iota_* dV(x)$, where $f : M \to \mathbb{R}$ and $x \in \iota(M)$. We interpret f as the p.d.f. of X on M. To alleviate the notation, in the following we abuse the notation and will not distinguish between $\iota(M)$ and M.

Pick up the kernel function $K(x) = e^{-x^2}$. Note that the kernel function can be more general, for example, $K \in C^2(\mathbb{R})$, non-zero, non-negative and monotonic decreasing, but we focus ourselves on this kernel to make the explanation clear. Denote $\mu_l^{(k)} := \int_{\mathbb{R}^d} ||x||^l K^{(k)}(||x||) dx$, where $k = 0, 1, 2, l \in \mathbb{N} \cup \{0\}$, and $K^{(k)}$ means the k-th order derivative of K. We assume $\mu_0^{(0)} = 1$ and \sqrt{h} is small enough so that \sqrt{h} is smaller than the reach [50] and the injectivity radius [27] of the manifold M, inj(M).

B-1 Connection Graph and Affinity Graph

An affinity graph, denoted as (G, w), is actually a special case of the connection graph in the sense that the connection function is not defined on the affinity graph. If we take a constant function $r_0 : E \to 1$, the affinity graph becomes a connection graph (G, w, r_0) .

We mention that in practice, if we decide to construct the connection graph from a given dataset, there are several different ways depending on the application and goal. For example, in addition to the examples discussed in the main context, in the geometric approach to the signal processing [58, 66], the affinity is defined to be the Mahalanobis distance reflecting the intrinsic property of the underlying state space; in the chair synchronization problem [40], the affinity between two chair meshes is defined based on their Hausdorff distance. The quality of the chosen affinity might influence the analysis result directly.

B-2 Connection Graph Laplacian and its Applications

Now we discuss the CGL. Consider the symmetric matrix $L(W,G)_s := D^{-1/2}SD^{-1/2}$, which is similar to L(W,G). Since $L(W,G)_s$ is symmetric, it has a complete set of eigenvectors $v_{n,i}$, $i = 1, \ldots, nd$ and its associated eigenvalues $\mu_{n,i}$, where the eigenvalues are bounded by 1 [59]. We would order the eigenvalues in the decreasing order. Note that the eigenvectors of $L(W,G)_s$ is related to those of L(W,G) via $D^{-1/2}$.

First, note that L(W,G) is an operator acting on $\boldsymbol{v} \in \mathbb{R}^{nd}$ by

$$(L(W,G)\boldsymbol{v})[i] = \frac{\sum_{j:(i,j)\in E} w(i,j)r(i,j)\boldsymbol{v}[j]}{\sum_{k:(i,k)\in E} w(i,k)},$$
(B-1)

where \boldsymbol{v} can be viewed as a vector-valued function defined on V so that $\boldsymbol{v}[j] := (\boldsymbol{v}((j-1)d+1), \ldots, \boldsymbol{v}(jd)) \in \mathbb{R}^d$. We could interpret this formula as a generalized random walk on the affinity graph. Indeed, if we view the vector-valued function \boldsymbol{v} as the status of a particle defined on the vertices, when we move from one vertex to the other one, the status is modified according to the relationship between vertices *encoded* in r. We mention that depending on the connection function, the structure L(W, G) might be very different, which leads to different analysis results and conclusion. We will give a precise example regarding this statement later. Now we discuss an important property of the CGL – the synchronization, which has been studied in [3, 20] and applied to the following problems, for example,

- 1. a new imaging technique aiming to obtain the atomic scale resolution images of a macro-scale object called "ptychographic imaging problem" [45];
- 2. a frame design called "polarization" for the phase retrieval problem [2];
- 3. a spectral relaxation approach to solve the least squares solution of the rotational synchronization problem [69];
- 4. graph realization problem by synchronization over the euclidean group [24, 25].

Here we give an intuition about this synchronization notion. Suppose there exists a vector-valued status $\boldsymbol{v} \in \mathbb{R}^{nd}$ of norm 1 which is "synchronized" according to the encoded relationship r in the sense that $\boldsymbol{v}[j] = r(j,i)\boldsymbol{v}[i]$ for all $(i,j) \in \mathbf{E}$, then $L(W,G)\boldsymbol{v}[i]$ will be the same as $\boldsymbol{v}[i]$, and hence the functional associated with the eigenvalue problem

$$\max_{\boldsymbol{v}\in\mathbb{R}^{nd};\,\|\boldsymbol{v}\|=1}\boldsymbol{v}^T L(W,G)\boldsymbol{v}$$

is maximized with the eigenvalue 1, and its eigenvector is the synchronized vector v. Thus, the top eigenvector of L(W,G), when viewed as the vector-valued status on the vertex, is the "synchronized" status with respect to the connection function. We mention that the existence of the synchronized vector-valued function is equivalent to the notion of "consistency" studied in [20].

When the connection function is constant; that is, the connection matrix $G_0 := \mathbf{11}^T$, where $\mathbf{1}$ is a $n \times 1$ vector with all entries 1, the GL is defined as $L := \mathrm{Id} - L(W, G_0)$. Notice a natural interpretation of $L(W, G_0)$ – since the sum of each row of $L(W, G_0)$ is 1, $L(W, G_0)$ is the *transition matrix* associated with a random walk on \mathbf{G} . To avoid confusion, the eigenvectors and eigenvalues of $L(W, G_0)_s$ are denoted as $u_{n,i}$ and $\nu_{n,i}$, where $i = 1, \ldots, n$, and $0 \le \nu_{n,i} \le 1$ are ordered in the decreasing order. As a special case of CGL, the GL has several applications which deserves discussion, for example

- 1. in the spectral clustering algorithm, we only need to find the first k trivial eigenvectors [44, 67];
- 2. to evaluate the Cheeger ratio, we need the second eigenvalue [18];
- 3. to visualize the high dimensional data, we need the first 3 non-trivial eigenvectors;
- 4. in the cryo-EM problem, if we want to reconstruct the rotational position of each projection image, we need the first 9 non-trivial eigenvectors [35];
- 5. in the 2d random tomography problem [60], only the first 2 non-trivial eigenvectors are needed;
- 6. in the orientability detection problem [59], we need the first eigenvector.

More algorithms depending on the eigenstructure of the GL can be found, to mention but a few, in [22, 41, 46, 51, 54, 55, 64, 65]. We comment that if the p.d.f f is not uniform, then a specific normalization stated in [21] allows us to study the dynamics of the underlying dynamical system [48, 49].

B-3 Asymptotical behavior of CGL

To better understand the CGL, we focus on the frame bundle and its associated tangent bundle [59] here to simplify the exploration. For CGL associated with a more general principal bundle structure, we refer the reader to [61].

- **Assumption B.1.** (D1) M is a smooth and compact d-dim manifold. When the boundary ∂M is not empty, it is assumed to be smooth, and we denote $M_{\delta} := \{x \in M : d(x, \partial M) \leq \sqrt{\delta}\};$
- (D2) The p.d.f. $f \in C^3(M)$ is uniformly bounded from below and above, that is, $0 < p_m \le f(x) \le p_M < \infty$. However, to simplify the exploration, we assume here that f is uniform, that is, f is a constant function defined on M. When f is non-uniform, its theoretical results can be found in [21, 61].

Under Assumption B.1, we collect the data \mathcal{X} independently and identically sampled from M and build up the following graph. First define a graph $G_{\mathrm{M}} := (\mathbb{V}, \mathbb{E})$ by taking $\mathbb{V} = \mathcal{X}$, and $\mathbb{E} = \{(x_i, x_j); x_i \in \mathcal{X}\}$. Then define the affinity function w on \mathbb{E} by

$$w: (i,j) \mapsto K_h(x_i, x_j) := K(\|x_i - x_j\|_{\mathbb{R}^p}^2 / h),$$

where h > 0 is the chosen bandwidth. Note that we choose to use the Euclidean distance, instead of the geodesic distance, to build up w since in practice we have only an access to the Euclidean distance (or other metric, depending on the application). Asymptotically this discrepancy will disappear.

Assumption B.2. (D3) For each point $x_i \in \mathcal{X}$, we also have a sample on the frame bundle $b(i) \in O(M)$ so that the b(i) is the basis of the tangent space $T_{x_i}M$. In particular, we are given a group-valued function $b : \mathcal{X} \to O(d)$.

With Assumption B.2, we define the *connection function* on E as

$$r: (i,j) \mapsto b(i)^T P_{x_i,x_j} b(j) \in O(d)$$

where P_{x_i,x_j} presents the parallel transport of the vector field from x_j to x_i . As a result, we have a connection graph $(\mathbf{G}_{\mathrm{M}}, w, r)$. With $(\mathbf{G}_{\mathrm{M}}, w, r)$, we build up the the CGL by $\mathrm{Id} - L(W, G)$, where W and G are the weight matrix and connection matrix associated with w and r. Under this framework, the GL is when we work with the trivial line bundle associated with M.

The geometrical meaning of the connection function deserves some discussions. First, note that although all tangent planes $T_{x_i}M$ are isomorphic to \mathbb{R}^d [27], but they are different in the sense that we cannot "compare" $T_{x_i}M$ and $T_{x_j}M$ directly. Precisely, it makes sense the say u - v when $u, v \in \mathbb{R}^d$, but we can not evaluate $u_i - u_j$ when $u_i \in T_{x_i}M$ and $u_j \in T_{x_j}M$. To carry out the comparison between different tangent planes, we need a bit more work. Indeed, $b(i) \in O(d)$ is a basis of the tangent plane of $T_{x_i}M$, which practical meaning is mapping \mathbb{R}^d isomorphically to $T_{x_i}M$. In other words, given a vector field Y, $b(i)^T Y(x_i)$ evaluates its coordinate at x_i . The parallel transport P_{x_i,x_j} is a geometrical generalization of the notion "translation" in the Euclidean space – it is an isometric map mapping $T_{x_j}M$ to $T_{x_i}M$. As a result, $r(i, j) \in O(d)$ is an isometric map from \mathbb{R}^d to \mathbb{R}^d , and geometrically it maps the coordinate of a vector field at x_j , that is, v[j], to the vector field at x_i , that is, b(j)v[j], then parallelly transports b(j)v[j] to x_i , and then evaluate the coordinate of $P_{x_i,x_j}b(j)\boldsymbol{v}[j]$ with related to the basis b(i). We emphasize that the connection function in the connection graph associated with the frame bundle encodes not only the geometry but also the topology of the manifold. In practice, this constraint may lead to a better understanding of the underlying data structure. For example, in the cryo-EM problem, this viewpoint leads to a better angular classification result.

We now state the pointwise convergence and the spectral convergence of the L(W, G). These theorems apply to the GL, while we replace the vector fields by the functions with the same regularity and the connection Laplacian operator by the Laplace-Beltrami operator and b(i) = 1 (see [61] for details).

Theorem B.3 (CGL Pointwise Convergence [59, 61]). Suppose Assumption B.1 and Assumption B.2 hold and $X \in C^4(TM)$. For all $x_i \notin M_{\sqrt{h}}$ with high probability (w.h.p.)

$$b(i) \left((\mathrm{Id} - L(W,G))\bar{X} \right) [i] = h \frac{\mu_2}{2d} \nabla^2 X(x_i) + O(h^2) + O\left(\frac{1}{n^{1/2} h^{d/4 - 1/2}}\right)$$

where $\bar{X} \in \mathbb{R}^{nd}$ and $\bar{X}[i] = b(i)^{-1}X(x_i)$. For all $x_i \in M_{\sqrt{h}}$, we have w.h.p.

$$b(i) \left((\mathrm{Id} - L(W,G))\bar{X} \right) [i] = O(\sqrt{h}) P_{x_i,x_0} \nabla_{\partial_d} X(x_0) + O(h) + O\left(\frac{1}{n^{1/2} h^{d/4 - 1/2}}\right),$$

where $x_0 = \operatorname{argmin}_{u \in \partial M} d(x_i, y)$ and ∇_{∂_d} is the derivative in the normal direction.

To state the spectral convergence result, define an operator $T_{C,h}: C(TM) \to C(TM)$:

$$T_{C,h}X(y) := \frac{\sum_{j:(i,j)\in E} K_h(y,x_j) P_{y,x_j}X(x_j)}{\sum_{j:(i,j)\in E} K_h(y,x_j)}$$

where $X \in C(TM)$. To simplify the discussion, we assume that the eigenvalues of the heat kernel of the connection Laplacian $e^{t\nabla^2}$ are simple. When there exists an eigenvalue with multiplicity greater than 2, the theorem can be proved using the projection operators onto the eigenspaces. We mention that in the special case GL, the point convergence theorem was first established in [8] under the uniform sampling and boundary-free assumption, and then extended to a more general setup in [21, 39, 56, 61], and the spectral convergence was established in [9, 67].

Theorem B.4 (CGL Spectral Convergence [61]). Suppose Assumption B.1 and Assumption B.2 hold and fix t > 0. Denote $\mu_{C,t,h,i}$ to be the *i*-th eigenvalue of $T_{C,h}^{t/h}$ with the associated eigenvector $X_{C,t,h,i}$. Also denote $\mu_{t,i} > 0$ to be the *i*-th eigenvalue of $e^{t\nabla^2}$ with the associated eigen-vector field $X_{t,i}$. We assume that both $\mu_{C,t,h,i}$ and $\mu_{t,i}$ decrease as *i* increases, respecting the multiplicity. Fix $i \in \mathbb{N}$. Then there exists a sequence $h_n \to 0$ such that

$$\lim_{n \to \infty} \mu_{C,t,h_n,i} = \mu_{t,i}, \quad and \quad \lim_{n \to \infty} \|X_{C,t,h_n,i} - X_{t,i}\|_{L^2(TM)} = 0$$

in probability.

With these Theorems, we are able to discuss why different connection functions lead to different analysis results. Consider S^2 embedded in \mathbb{R}^3 with the standard metric. If we define the connection function according to the Levi-Civita connection, then the top eigenvalue of $e^{t\nabla^2}$ is strictly less than 1 due to the hairy-ball theorem [36]. In other words, asymptotically we are not able to find a synchronized vectorvalued status on it. On the other hand, if we take the trivial connection function, that is, $r(i, j) = I_2$, then asymptotically we obtain Δ_g acting on two independent functions. Since the dimension of the null space of Δ_g is the number of the connected components of the manifold, the top eigenvalue of the CGL with the trivial connection function is 1; that is, a synchronized vector-valued status exists. See Figure 5 for the result.

The main reason leading to this difference is rooted in the connection theory, and we refer the interested reader to [14].



Figure 5: The first 10 eigenvalues of the Id - L(W, G) with related to a non-trivial connection function determined from the Levi-Civita connection of the frame bundle of S^2 (left) and a trivial connection function (right). 1000 points are sampled uniformly from S^2 . Note that the eigenvalues on the right figure are the same as those of GL while the multiplicities of all the eigenvalues are 2.

B-4 The "Sparsity" of the CGL

We define the following "sparsity" condition.

Definition 1. Fix $\gamma > 0$. For a $n \times n$ matrix Q, we sort its eigenvalues $\nu_{Q,\ell}$, $\ell = 1, \ldots, n$, so that $|\nu_{Q,1}| \ge |\nu_{Q,2}| \ge \ldots \ge |\nu_{Q,n}|$. Then Q satisfies the γ -sparsity property if

$$|\nu_{Q,\ell}| \le e^{-C_Q \ell^{\gamma}}$$

for all ℓ , where $C_Q > 0$ depends on Q.

We now claim that the CGL under the manifold assumption satisfies the 2/d-sparsity property. Note that this theorem also holds for the GL.

Theorem B.5. Asymptotically when $n \to \infty$, for $\ell \in \mathbb{N}$, $\mu_{n,\ell} \leq e^{-C_L \ell^{2/d}}$, where the constants $C_L > 0$ depends d, the lower bound of the Ricci curvature k and the diameter D.

Proof. Note that the Weyl's theorem [12] holds for the connection Laplacian, that is,

$$\tilde{N}(\mu) \sim \frac{1}{(4\pi)^{d/2} \Gamma(d/2+1)} \mu^{d/2},$$

where $\tilde{N}(\mu)$ is the number of eigenvalues of ∇^2 less than $\mu > 0$, and we have the consequence that [70]

$$\lambda_j \ge c'(d,k,D)j^{2/d},\tag{B-2}$$

where $j \in \mathbb{N}$ and c'(d, k, D) is the universal constant depending only on d, the lower bound of the Ricci curvature k and the diameter D. Note that since $\mu_{n,\ell} \to e^{-\lambda_{\ell}}$ in probability, we have

$$\mu_{n,\ell} \le e^{-c'(d,k,D)\ell^{2/d}}.$$

Hence, combined with Theorem B.4, we conclude the claim with $C_{\rm L} := c'(d, k, D)$.

B-5 Vector Diffusion Maps and Diffusion Maps

In this subsection, we discuss a potential application of the CGL and GL – estimate the local geodesic distance. Fix t > 0, we define the vector diffusion maps (VDM) $V_{t,n} : \mathcal{X} \to \mathbb{R}^{(nd)^2}$ by

$$V_{t,n}: x_i \mapsto \left((\mu_{n,l}\mu_{n,r})^t \langle v_{n,l}[i], v_{n,r}[i] \rangle \right)_{l,r=1}^{nd}$$

where $v_{n,l}[i]$ is a *d*-dim vector containing the ((i-1)d+1)-th entry to the (id)-th entry of the eigenvector $v_{n,l}$. With this map, the Hilbert-Schmidt norm of the (i, j)-th block of $L(W, G)_s$ satisfies

$$||L(W,G)_{s}^{2t}(i,j)||_{HS}^{2} = \langle V_{t,n}(x_{i}), V_{t,n}(x_{j}) \rangle_{S}$$

that is, $||L(W,G)_s^{2t}(i,j)||_{HS}^2$ becomes an inner product for the finite dimensional Hilbert space. The reason we need to consider $L(W,G)_s^{2t}$ but not $L(W,G)_s^t$ is that all eigenvalues $\mu_{n,l}$ of $L(W,G)_s$ reside in the interval

[-1, 1], and we can not guarantee the positivity of $\mu_{n,l}$ when n is finite. We can then define the vector diffusion distance (VDD) to quantify the affinity between nodes i and j:

$$d_{\mathcal{C},t,n} := \|V_{t,n}(x_i) - V_{t,n}(x_j)\|^2.$$

The theoretical properties of VDM and VDD will be clear when $n \to \infty$. Fix $a \in \mathcal{B}(TM, g)$ and t > 0, define the following map embedding $x \in M$ to ℓ^2 :

$$V_t^a : x \mapsto \left(\frac{1}{\sqrt{d}(4\pi)^{d/2} t^{(d+1)/2}} e^{-(\lambda_k + \lambda_l)t/2} \langle X_k(x), X_l(x) \rangle \right)_{k,l=1}^{\infty}.$$
 (B-3)

With V_t^a , we define a new affinity between pairs of points by

$$d_{\mathcal{C},t}(x,y) := \|V_t^a(x) - V_t^a(y)\|_{\ell^2}.$$
(B-4)

Due to Theorem B.3 and Theorem B.4, the VDM (resp. VDD) is a discretization of V_t^a (resp. $d_{C,t}$), so we may abuse the notation and call V_t^a VDM and $d_{C,t}$ VDD. We have the following Theorem saying that locally $d_{C,t}$ approximates the geodesic distance:

Theorem B.6 ([59]). Take a Riemannian manifold (M, g). For all t > 0, the VDM V_t^a is diffeomorphic. Furthermore, suppose $x, y \in M$ so that $x = \exp_u v$, where $v \in T_y M$. When $||v||^2 \ll t \ll 1$ we have

$$d_{C,t}^2(x,y) = \|v\|^2 + O(t\|v\|^2).$$
(B-5)

Although GL is a special case of CGL, with GL we may define a different embedding which has different features. Given t > 0 and $0 \le \delta < 1$, the diffusion maps (DM) with diffusion time t^4 as

$$\Phi_{t,n}: x_i \mapsto \left(\nu_{n,l}^t u_{n,l}(i)\right)_{l=2}^n.$$
(B-6)

One similar but different algorithm is the Laplacian eigenmaps [7, 8], that is, x_i is mapped to $(u_{n,l}(i))_{l=2}^m$, which can be viewed as a special DM with diffusion time t = 0 and $1 < m \leq n$ is chosen by the user. Yet another similar quantity referred to as the global point signature proposed in [55], which maps x_i to $((-\ln \nu_{n,l})^{-1/2} u_{n,l}(i))_{l=2}^m$, where m is chosen by the user. Another variation is the commute time embedding [53]. We mention in the Laplacian eigenmaps, global point signature and commute time embedding, the notion "diffusion" does not exist. Although these mappings are diffeomorphic to each other when m = n via a linear transformation, asymptotically their behaviors are different. Furthermore, even if the connection function is trivial, the VDM and DM are different. With DM, we introduce a new metric between sampled points, which is referred to as diffusion distance (DD):

$$d_{\mathrm{DM},t,n}(x_i, x_j) := \|\Phi_{t,n}(x_i) - \Phi_{t,n}(x_j)\|_{\mathbb{R}^{n-1}}.$$
(B-7)

To study $d_{\text{DM},t,n}$, we take $a \in \mathcal{B}(M,g)$ and t > 0, and map $x \in M$ to the Hilbert space ℓ^2 by [11]

$$\Phi_t^a : x \mapsto \sqrt{\operatorname{vol}(\mathbf{M})} \left(e^{-\gamma_\ell t} \phi_\ell(x) \right)_{\ell=1}^\infty, \tag{B-8}$$

With the map Φ_t^a , we are able to define a new affinity between pairs of points:

$$d_{\mathrm{DM},t}(x,y) := \|\Phi_t^a(x) - \Phi_t^a(y)\|_{\ell^2}.$$
(B-9)

Due to Theorem B.3 and Theorem B.4, the DM (resp. DD) is a discretization of Φ_t^a (resp. $d_{\text{DM},t}$), so we may abuse the notation and call Φ_t^a DM and $d_{\text{DM},t}$ DD.

It has been shown that the DM satisfies the following "almost isometric" property [59]:

Theorem B.7. Take a Riemannian manifold (M,g). For all t > 0, Φ_t^a is diffeomorphic. Furthermore, suppose $x, y \in M$ so that $x = \exp_y v$, where $v \in T_y M$. When $||v||^2 \ll t \ll 1$ we have

$$d_{\mathrm{DM},t}^2(x,y) = \|v\|^2 + O(t\|v\|^2).$$
(B-10)

⁴In practice, we may consider the truncated diffusion maps (tDM) with diffusion time t and accuracy δ , which is defined as $\Phi_{t,n,m(\delta,t)}: x_i \mapsto \left(\nu_{n,l}^t u_{n,l}(i)\right)_{l=2}^{m(\delta,t)}$, where $m(\delta,t) \in \mathbb{N}$ such that $\lambda_{m(\delta,t)}^t > \delta \lambda_1^t$ and $\lambda_{m(\delta,t)+1}^t \leq \delta \lambda_1^t$ [21]



Figure 6: Left: the Trefoil knot; middle: the truncated DM with the first 2 non-trivial eigenvectors of the GL; right: the truncated DM with the first 3 non-trivial eigenvectors of the GL.

The above theorems, when combined with the above spectral convergence theorem, says that the VDD and DD provide an accurate estimation of the geodesic between two close points. While combined with the manifold sparsity property, we have the following practical fact – if we are allowed a positive small error when we estimate the geodesic distance, we do not need to recover the whole eigen-structure. Instead, the first few eigenvalues and eigenvectors are enough.

We have the following statement shown in [6, 52]. Fix $\epsilon > 0$ and (M, g) is a *d*-dim manifold satisfying $\operatorname{Ric}(g) \geq (d-1)kg$, $\operatorname{vol}(M) \leq V$, $\operatorname{inj}(M) \leq I$. Then there exists a $t_0 = t_0(d, k, I, \epsilon)$ such that for all $0 < t < t_0$, these is $N_E = N_E(d, k, I, V, \epsilon, t)$ so that if $N \geq N_E$, the truncated diffusion maps

$$\Phi^a_{t,N}: x \mapsto \sqrt{\operatorname{vol}(\mathbf{M})} \left(e^{-\gamma_\ell t} \phi_\ell(x) \right)_{\ell=1}^N$$
(B-11)

is an embedding of M into \mathbb{R}^N and

$$1 - \epsilon < \left| \frac{(2t)^{(n+2)/4} \sqrt{2} (4\pi)^{n/4}}{\sqrt{\text{vol}(\mathbf{M})}} \mathrm{d}\Phi_{N,t} |_x \right| < 1 + \epsilon.$$
 (B-12)

Before ending this section, we show an interesting example regarding the data visualization and embedding issue. Take the Trefoil knot M embedded in \mathbb{R}^3 by $\iota(t) = [\sin(t) + 2\sin(2t), \cos(t) - 2\cos(2t), -\sin(3t)]$, where $t \in [0, 2\pi)$. We refer to Figure 6 for an illustration. Note that the Trefoil knot is not homeomorphic to S^1 . We sample 1000 points uniformly from $[0, 2\pi)$ independently; that is, we sample 1000 points on M non-uniformly. If we want to visualize the dataset, we may apply the tDM to embed M to \mathbb{R}^3 (or \mathbb{R}^2). The result is shown in Figure 6. The results deserve some discussion. Note that the tDM maps the Trefoil knot is not preserved. Note that for the visualization purpose, we only choose the first 3 (or 2) eigenvectors, which leads to a map which deteriorate the topology. If we want to guarantee the preservation of the topology, we need the embedding theorem counting how many eigenvectors we need. This opens the following question, in particular when the dataset is noisy – how to balance between different data analysis results, for exampling, how to balance between preserving the topology information and data visualization?