*Information and Inference: A Journal of the IMA* (2022) **11**, 1173–1202 https://doi.org/10.1093/imaiai/iaac010 Advance Access publication on 6 August 2022

# Optimal recovery of precision matrix for Mahalanobis distance from high-dimensional noisy observations in manifold learning

MATAN GAVISH

School of Computer Science and Engineering, Hebrew University of Jerusalem, Jerusalem 9190401, Israel

> PEI-CHUN SU Department of Mathematics, Duke University, Durham, NC 27710, USA

> > RONEN TALMON

Viterbi Faculty of Electrical and Computer Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel

AND

Hau-Tieng Wu $^{\dagger}$ 

Department of Mathematics and Department of Statistical Science, Duke University, Durham, NC 27710, USA <sup>†</sup>Corresponding author. Email: hauwu@math.duke.edu

[Received on 28 December 2019; revised on 16 August 2021; accepted on 4 March 2022]

Motivated by establishing theoretical foundations for various manifold learning algorithms, we study the problem of Mahalanobis distance (MD) and the associated precision matrix estimation from highdimensional noisy data. By relying on recent transformative results in covariance matrix estimation, we demonstrate the sensitivity of MD and the associated precision matrix to measurement noise, determining the exact asymptotic signal-to-noise ratio at which MD fails, and quantifying its performance otherwise. In addition, for an appropriate loss function, we propose an asymptotically optimal shrinker, which is shown to be beneficial over the classical implementation of the MD, both analytically and in simulations. The result is extended to the manifold setup, where the nonlinear interaction between curvature and highdimensional noise is taken care of. The developed solution is applied to study a multi-scale reduction problem in the dynamical system analysis.

Keywords: Mahalanobis distance; large p large n; optimal shrinkage; precision matrix.

## 1. Introduction

High-dimensional datasets encountered in modern science often exhibit nonlinear low-dimensional structures. One prominent approach to deal with such point clouds is to model their nonlinear structures by manifolds. In the last two decades, this direction has led to the emergence of a multitude of *manifold learning* methods, including the classical ISOMAP [42], locally linear embedding (LLE) [29], Hessian LLE [7], eigenmap [4] and diffusion maps [5], as well as the more recent vector diffusion maps [32], multiview diffusion maps [21] and alternating diffusion maps [18, 39], to name but a few. Typically in manifold learning, point clouds in  $\mathbb{R}^p$  are assumed to be sampled from a *d*-dimensional smooth manifold *M* embedded in  $\mathbb{R}^p$ , usually with some additional contaminating noise. In this setting, the manifold represents the 'essence' or the 'signal' of the data. Consequently, the goal in manifold learning is to recover the geometric or topological structure of the manifold from the data points and,

1174

#### M. GAVISH ET AL.

in turn, to use the recovered structure to embed the high-dimensional data in a low-dimension space, facilitating a compact and informative representation of the data. This approach has been successfully applied to applications from a broad range of fields, e.g. dynamical systems modeling [41, 48], sleep stage assessment [22, 46], cryo electro microscope [30], image denoising [34], single channel blind source separation such as fetal electrocardiogram analysis [36] and stimulation artifact removal for the intracranial electrocephalogram [1], and long-term physiological signal visualization and analysis [20, 43].

Classical manifold learning methods heavily rely on meaningful measures of pairwise discrepancy between data points. In this so-called *metric design* problem, the data analyst aims to find a useful *metric* representing the relationship between data points embedded in a high-dimensional space. In this paper, we study the Mahalanobis Distance (MD)—a popular, and arguably the first method for metric design [23, 26]. MD was originally proposed in 1936 with the classical low-dimensional setting in mind, namely, for the case where the ambient dimension p is much larger than the dataset size n. Interestingly, due to its useful statistical and invariance properties, MD became the basis of several geometric data analysis techniques [44, 47, 49], aimed specifically at the high-dimensional regime  $p \simeq n$ .

In a recent line of work [31, 37], a variant of MD was proposed and used to reveal hidden intrinsic manifolds underlying high-dimensional, possibly multi-modal, observations. The main purpose of MD in this hidden manifold setup is handling possible deformations caused by the observation or sampling process. Broadly, this is carried out by estimating a quadratic form of the Jacobian of the (unknown) observation function, which is equivalent to estimating the precision matrix *locally* on the manifold. It was recently shown that MD is also implicitly used in the seminal LLE algorithm [29], when the barycenter step is properly expressed [24].

As the number of dimensions p in typical data analysis applications continues to grow, it becomes increasingly crucial to understand the behaviour of MD, as well as other metric design algorithms, in the high-dimensional regime  $p \approx n$ . At first glance, it might seem that this regime poses little more than a computational inconvenience for metric design using MD. Indeed, it is easy to show that in the absence of measurement noise, MD cares little about the increase in the ambient dimension p. This paper however calls attention to the following key observation. In the high-dimensional regime  $n \approx p$ , in the presence of ambient measurement noise, a new phenomenon emerges, which introduces various nontrivial effects on the performance of MD. Depending on the noise level, in the highdimensional regime, MD may be adversarially affected or even fail completely. Clearly, the assumption of measurement noise cannot be realistically excluded, and yet, to the best of our knowledge, this phenomenon has not been previously fully studied. A first step in this direction was taken in [6], with the calculation of the distribution of MD under specific assumptions.

Let us describe this key phenomenon informally at first. The computation of MD involves the estimation of the inverse covariance matrix, or precision matrix, corresponding to the data at hand. Classically, the estimation relies on the sample covariance, which is inverted using the Moore–Penrose pseudoinverse. It is well-known that, in the high-dimensional setup or in the regime  $n \approx p$ , the sample covariance matrix is a poor estimator of the underlying covariance matrix. Indeed, advances in random matrix theory from the last decade imply that the eigenvalues and eigenvectors of the sample covariance matrix are both biased, namely, do not converge to the corresponding eigenvalues and eigenvectors of the underlying population covariance matrix [14, 28]. Such biases in small eigenvalues, which lead to inaccurate covariance matrix estimation, become immense when applying the Moore–Penrose pseudo-inverse.

This challenge in the high-dimensional setup is amplified when we have a nonlinear manifold structure. Inverting small (inconsistent) eigenvalues is challenging when evaluating the precision matrix, and this issue becomes more challenging in the context of manifold learning, where the estimation

of MD is performed locally [31, 37]. Ideally, under the manifold assumption, in infinitely small neighborhoods and for a sufficiently large number of samples without noise, the rank of the local sample covariance equals the dimension of the manifold, which is typically smaller than the dimension of the ambient space, making the sample covariance low-rank with distinct strictly zero eigenvalues. However, in practice, due to the finite sample set, the considered neighborhoods cannot be sufficiently small. In such cases, depending on the manifold curvature, samples from the manifold depart from the tangent space to the manifold, and the rank of their sample covariance matrix undesirably grow, where the eigenvalues related to curvature are much smaller compared with those related to the tangent space. When noise exists, the situation is further complicated by the interaction between those curvature-related small eigenvalues and the relatively few points compared with the ambient space dimension.

In this paper, we study this problem and propose a remedy. By relying on formal existing results in covariance matrix estimation, we measure the sensitivity of MD to measurement noise. Under the assumption that, locally, the data on the manifold lie on a low-dimensional linear subspace embedded in the ambient space  $\mathbb{R}^p$  and that the measurement noise is Gaussian white, we are able to determine the exact asymptotic signal-to-noise ratio at which MD fails, and quantify its performance otherwise. Then, we propose a better MD estimator based on the idea of *shrinkage* of the associated precision matrix. It has been known since the 1970's [35] that by shrinking the sample covariance eigenvalues one can significantly mitigate the noise effects and improve the covariance estimation in high-dimensions. We formulate the classical MD as a particular choice of shrinkage estimator for the eigenvalues of the sample covariance. Building on recent results in high-dimensional covariance estimation, including the general theory in [8] and a special case with application in random tomography [32, Section 4.4), we find an *asymptotically optimal shrinker* for the precision matrix estimation, which is better than the classical implementation of the MD, whenever MD is computed from noisy high-dimensional data. We show that under a suitable choice of a loss function for the estimation of MD, our shrinker is the unique asymptotically optimal shrinker; the improvement in asymptotic loss it offers over the classical MD is calculated exactly. We then extend the above established results to handle the challenge of designing a better metric when applying diffusion map. This extension is nontrivial due to the nontrivial interaction of curvature and noise and the finite sampling size. Finally, we apply diffusion maps with the proposed estimate of MDto separate slow and fast dynamics when the observation is contaminated by highdimensional noise.

While the present paper focuses on MD, we posit that the same phenomenon holds much more broadly and in fact affects several widely used manifold learning, particular in metric learning algorithms. In this regard, the present paper seeks to highlight the fact that manifold learning and metric learning algorithms will *not* perform as predicted by the noiseless theory in high dimensions and may fail completely beyond a certain noise level.

### 2. Problem setup

#### 2.1 Manifold model

When a point cloud  $\mathscr{X} := \{x_i\}_{i=1}^n \subset \mathbb{R}^p$  has a nontrivial nonlinear structure, or even nontrivial topological structure, a common approach is to model this structure by a manifold. This is known as the so-called *manifold assumption*. Such manifold assumption holds for various practical data. Examples include cryo-electro microscopy [30], phase spaces of dynamical systems [41, 48] and various biomedical signals [1, 20, 36, 43]. The main feature of this manifold assumption is that the points are

distributed on a nonlinear set so that they are nonlinearly related, which generalizes the commonly used linear model.

To model  $\mathscr{X}$  by the manifold model, consider a *p*-dimensional random vector  $X : \Omega \to \mathbb{R}^p$ , which is a measurable function with respect to the probability space  $(\Omega, \mathscr{F}, \mathbb{P})$ , where  $\mathbb{P}$  is the probability measure defined on the sigma algebra  $\mathscr{F}$  in the event space  $\Omega$ . We assume that the range of X is supported on a *d*-dimensional compact, smooth Riemannian manifold (M, g) isometrically embedded in  $\mathbb{R}^p$  via  $\iota : M \hookrightarrow \mathbb{R}^p$ . In this work, we assume that M is boundary-free to simplify the exposition. We shall mention that the commonly considered linear model is a special manifold model, where the manifold is an affine linear subspace space of  $\mathbb{R}^p$ .

On the manifold, the associated statistical setup is as follows. Let  $\hat{\mathscr{B}}$  be the Borel sigma algebra of  $\iota(M)$ , and let  $\mathbb{P}_X$  denote the probability measure induced from X. Clearly,  $\mathbb{P}_X$  is defined on  $\tilde{\mathscr{B}}$ . Denote dV to be the Riemannian volume density of M associated with the metric g. For simplicity, we assume that  $\mathbb{P}_X$  is absolutely continuous with respect to the induced Riemannian measure on  $\iota(M)$ , denoted by  $\iota_* dV$ . By the Radon–Nikodym theorem, for any  $z \in \iota(M) \subset \mathbb{R}^p$ , there exists a non-negative measurable function P(z) defined on  $\iota(M)$  such that  $d\mathbb{P}_X(z) = P(z)\iota_* dV(z)$ . The probability density function (pdf) of X on M is defined to be P(z). We further assume that P(z) is bounded away from 0 and smooth. When P(z) is constant, we call X a *uniform* random sampling scheme; otherwise, it is *nonuniform*. Now, we introduce the key quantity of interest in this paper, the local covariance matrix.

DEFINITION 2.1. Fix  $x \in M$ . For an open simply connected neighborhood of x,  $\mathscr{O}(x) \subset \mathbb{R}^p$ , define

$$\Sigma_{\widehat{\mathcal{O}}(x)} := \mathbb{E}[(X - \mu_{\widehat{\mathcal{O}}(x)})(X - \mu_{\widehat{\mathcal{O}}(x)})^{\top}\chi_{\widehat{\mathcal{O}}(x)}(X)] \in \mathbb{R}^{p \times p}$$
(2.1)

as the *local covariance matrix* associated with  $\mathcal{O}(x)$  centered at  $\mu_{\mathcal{O}(x)}$ , where

$$\mu_{\mathscr{O}(x)} := \mathbb{E}[X\chi_{\mathscr{O}(x)}(X)] = \frac{1}{|\mathscr{O}(x) \cap \iota(M)|} \int_{\mathscr{O}(x) \cap \iota(M)} zP(z)dz,$$

 $|\mathscr{O}(x) \cap \iota(M)|$  is the volume of  $\mathscr{O}(x) \cap \iota(M)$  and  $\chi_{\mathscr{O}(x)}$  is the indicator function of the set  $\mathscr{O}(x)$ .

One main goal of considering this local covariance is capturing those directions with maximal variation of the dataset. From the knowledge of principal component analysis, those directions are the eigenvectors of the local covariance matrix with the largest eigenvalues. We have some remarks. First, by Nash's isometric embedding theorem,  $\Sigma_{\mathcal{O}(x)}$  is of rank *D*, where  $D \leq d(3d + 11)/2$  for any  $\mathcal{O}(x)$ . Second, in existing literature, there is another different definition of the local covariance matrix, in which the mean  $\mu_{\mathcal{O}(x)}$  is replaced with  $\iota(x)$ . In general, the two definitions are different, even when the set is perfectly symmetric, e.g.  $\mathcal{O}(x)$  is a ball and P(x) is uniform. Indeed, since the range of *X* is supported on  $\iota(M)$ , when  $\iota(M)$  is not flat,  $\iota(x)$  might deviate from the center of  $\mathcal{O}(x)$  due to the curvature of the manifold; that is,  $\frac{1}{|\mathcal{O}(x)\cap \iota(M)|} \int_{\mathcal{O}(x)\cap \iota M} (z - \iota(x)) dz \neq 0$ . While this seems to be a problem, it was shown in [24] that the difference between the center of  $\mathcal{O}(x)$  and  $\iota(x)$  is negligible (expressed as a higher order term in the error) when the diameter of  $\mathcal{O}(x)$  is sufficiently small. Broadly, since locally the manifold can be well approximated by an affine subspace, when the diameter of  $\mathcal{O}(x)$  is sufficiently small, the data located in  $\mathcal{O}(x)$  can be well approximated by the tangent space to the manifold at *x*, i.e.  $T_x M$ . This point will be further addressed in Section 2.2.

The above derivation leads to the following local statistical model, which enables to further study the local structure of a manifold. For  $x \in M$  and an open, simply connected neighborhood of x,  $\mathcal{O}(x) \subset \mathbb{R}^p$ ,

we define a new random vector

$$X_{x} := X \chi_{\widehat{\mathcal{O}}(x)}(X), \qquad (2.2)$$

with mean  $\mu_x := \mu_{\mathcal{O}(x)}$  and covariance matrix  $\Sigma_x := \Sigma_{\mathcal{O}(x)}$ . By definition,  $X_x$  is a bounded random vector, and hence, all its moments are finite. Also, as has been widely discussed in the literature (see [24] and reference therein), the first *d* dominant eigenvectors of  $\Sigma_x$  form an accurate estimate of the tangent space to the manifold at *x*. Specifically, if  $\Sigma_x u_l = \lambda_l u_l$ , for  $l = 1, \ldots, p$ , where  $\lambda_1 \ge \lambda_2 \ldots \ge \lambda_p$ , then the span of  $\{u_1, \ldots, u_d\}$  approximates  $\iota_* T_x M$ .

Often in applications, the manifold is not directly accessible due to additional noise, and we can only sample

$$Y = X + \sigma\xi, \tag{2.3}$$

where  $\xi \sim \mathcal{M}(0, I_p)$ . As a consequence, for  $x \in M$ , in the presence of noise, the corresponding local random vector can be recast as

$$Y_x := Y \chi_{\mathcal{O}(x)}(X), \tag{2.4}$$

with mean  $\mu_x := \mu_{\hat{\mathcal{O}}(x)}$  and covariance matrix  $\Sigma_x + \sigma^2 I_p$ . Although  $Y_x$  is not a bounded random vector, all its moments are finite.

We remark that the local set  $\mathcal{O}(x)$  can be defined in several plausible ways, depending on the problem at hand. A common choice is the ball  $B_x(\epsilon)$  with the center at  $\iota(x)$  and the radius  $\epsilon > 0$ . In other applications,  $\mathcal{O}(x)$  might be an ellipsoid [24, 31, 37] or a more general setup, depending on the metric of interest. We will revisit to this issue in the sequel.

#### 2.2 Linear spiked model

In manifold learning, local kernels are commonly used, e.g. kernels based on radial basis functions. The use of kernels implies that only points around the center of the kernel  $x \in M$  contribute to the algorithm outcome. Therefore, considering those points located inside the neighborhood of x,  $\mathcal{O}(x)$ , is sufficient for analyzing a manifold learning algorithm with a local kernel. Since the bandwidth of the kernel is typically small, the diameter of the local set  $\mathcal{O}(x)$  is small as well. Consequently, by the definition of a manifold, data in  $\mathcal{O}(x)$  can be well approximated by the tangent space to the manifold at x, which is a low-dimensional affine subspace. Note that in the special case that the manifold is a linear affine subspace, all points are located by a low-dimensional space.

Since *locally* the manifold can be viewed merely as a linear subspace, the local statistical model described in Section 2.1 can be well approximated by the classical linear spiked model (or spiked covariance model [14]), which is detailed next with slight modifications in the notation. We note that in this section, deviations of the samples from a linear space due to the manifold curvature will not be treated, and their affect will be considered together with the affect of the ambient noise. However, we will extend and test the ability of the proposed estimator to handle such phenomena in the simulation study in Section 6.

We now consider the spike model. In plain English, a spike model is a manifold model when the manifold is an affine subspace, where the dimension of the affine subspace is fixed. Consider a point cloud in  $\mathbb{R}^p$  supported on a *d*-dimensional linear subspace, where  $d \leq p$ . For simplicity, we assume that

the point cloud is sampled independently and identically (i.i.d.) from

$$X \sim \mathcal{N}(\mu, \Sigma_X) \in \mathbb{R}^p, \tag{2.5}$$

where  $\mu$  denotes the mean and  $\Sigma_X$  denotes the population covariance matrix, whose rank is equal to d. Note that we could consider a random vector with mean 0 and a finite fourth moment [8]; since this could only increase the notational burden without providing additional insights, we focus on this simplified model.

It is often convenient to note that a point cloud sampled from X can be understood as sampling i.i.d. from a p-dim random vector

$$X = \mu + \sum_{l=1}^{d} \sqrt{\lambda_l} \zeta_l u_l, \qquad (2.6)$$

where  $\lambda_l > 0$ ,  $\zeta_l \sim \mathcal{M}(0, 1)$ ,  $\mathbb{E}\zeta_l \zeta_k = \delta_{l,k}$ ,  $\Sigma_X u_l = \lambda_l u_l$  and  $||u_l||_{L^2} = 1$  for l = 1, ..., d. Thus, the *d*-dimensional linear affine subspace is the space spanned by  $u_1, ..., u_d$ , which could be understood as spikes and hence the nomination of the model, and shifted by  $\mu$ . We note that this global linear model is related to the local structure of a manifold in (2.2) in the following way:  $\mu$  here is the center point *x* on the manifold in (2.2), and the *d*-dimensional linear affine subspace spanned by  $u_1, ..., u_d$  is the tangent space at *x*, while we need extra components to capture the curvature of the manifold.

Similarly to the manifold model, suppose that the samples of *X*, which we refer to as the *signal*, are not directly observable. Instead, the observed data consist of samples from the random variable

$$Y = X + \sigma \xi , \qquad (2.7)$$

where  $0 \le \sigma < \infty$  and  $\xi \sim \mathcal{M}(0, I_p)$  is a Gaussian measurement noise independent of X, which we assume for simplicity to be white.

### 3. Precision matrix and MD estimation

The focal point of this paper is the estimation of the MD under the manifold model from a statistical standpoint, which, as described in Section 2, leads to the classical linear spiked model. The estimation of the MD involves the estimation of the local precision matrix. Therefore, we start this section with details on the estimation of the precision matrix in Section 3.1, followed by a detailed description of the estimation of the MD in Section 3.2.

### 3.1 Precision matrix estimation and its challenge

In Sections 2.1 and 2.2, we showed that the local and global statistical models are seemingly very similar. Indeed, at first glance, both models consist of a hidden 'signal' component ((2.2) and (2.5)) and noisy accessible observations ((2.4) and (2.7)). Furthermore, under both models, the local and global population covariance matrices of the 'signal' component ( $\Sigma_x$  and  $\Sigma_x$ ) are approximately low rank and precisely low rank, respectively.

However, one of the main claims of this work is that although the two models seem equivalent, subtle differences between them, particularly in the context of precision matrix estimation, become fundamental.

Let us consider first the estimation of the precision matrix in the simpler, classical linear spiked model described in (2.5) and (2.7). Without noise, the computation of the precision matrix of X could be simply implemented using the Moore–Penrose pseudo-inverse. Suppose the eigendecomposition of  $\Sigma_X$  is given by  $\Sigma_X = U \operatorname{diag}(\lambda_1, \ldots, \lambda_d, 0, \ldots, 0) U^{\top}$ , where  $\lambda_1 \ge \lambda_2 \ge \ldots$  and  $U \in O(p)$ . Then, the pseudo-inverse is given by  $\Sigma_X^{\dagger} = U \operatorname{diag}(1/\lambda_1, \ldots, 1/\lambda_d, 0, \ldots, 0) U^{\top}$ , namely, inverting the non-zero eigenvalues. The introduction of additive noise poses a significant challenge for such an estimation, since the small eigenvalues could be mixed with the noise. While the contribution of such small eigenvalues is limited in the composition of the covariance matrix, in the composition of the precision matrix their affect becomes significant.

In addition to this classical challenge, the estimation of the precision matrix under the manifold model poses another layer of complexity. To simplify the discussion, we assume that  $\mathcal{O}(x)$  is a simply connected ball, that is  $\mathcal{O}(x) = B_{\epsilon}^{\mathbb{R}^p}(\mu_x)$ , where  $B_{\epsilon}^{\mathbb{R}^p}(\mu_x)$  is a Euclidean ball in  $\mathbb{R}^p$  of radius  $\epsilon > 0$  centered at  $\mu_x$  with a sufficiently small  $\epsilon$ . In this case, the geometric picture of the local covariance matrix is well captured by [45, Proposition 3.1], which is summarized as follows. Fix  $x \in M$ . Assume that the manifold is translated and rotated properly, so that x = 0 and the tangent space in  $\mathbb{R}^p$ ,  $\iota_* T_x M$ , is spanned by  $\{e_1, \ldots, e_d\}$ , where  $e_j$  is a unit *p*-dim vector with 1 in the *j*th entry. We have the following asymptotical expansion of (2.2) when  $\epsilon > 0$  is sufficiently small:

$$\Sigma_{B_{\epsilon}^{\mathbb{R}^{p}}(\mu_{x})} = \frac{|S^{d-1}|P(x)|}{d(d+2)} \epsilon^{d+2} \left( \begin{bmatrix} I_{d \times d} & 0\\ 0 & 0 \end{bmatrix} + O(\epsilon^{2}) \right), \tag{3.1}$$

where  $S^{d-1}$  is a (d-1)-sphere,  $|S^{d-1}|$  is the volume of  $S^{d-1}$  and the implied constant in  $O(\epsilon^2)$  depends on the second fundamental form at x. From a statistical perspective, this expansion well captures the intuition that the variability of the data located in  $B_{\epsilon}^{\mathbb{R}^{p}}(\mu_{x})$  is restricted to only few directions aligned with  $T_{x}M$ . Indeed, by applying the perturbation theory to (3.1), the *d* eigenvectors of  $\Sigma_{B_{\epsilon}^{\mathbb{R}^{p}}(\mu_{x})}$ corresponding to the largest *d* eigenvalues provide an approximate basis for the embedded tangent plane  $\iota_{*}T_{x}M$ . The expression of the covariance matrix in (3.1) implies that there is a significant spectral gap between the top *d* eigenvalues and the remaining ones, which depend on the curvature of the manifold. Furthermore, these remaining eigenvalues could be very small with order higher than  $\epsilon^{d+4}$ . In other words, under the manifold setting, in contrast to the linear spiked model, even in the noiseless case, the covariance is not strictly low rank with only *d* nonzero eigenvalues. Such small eigenvalues due to the curvature of the manifold, if not properly taken care before taking inverse to estimate the precision matrix, the obtained precision matrix might be deformed. Such small eigenvalues can be added to the noise components and together they obscure the spectral gap. Thus, it is necessary for determining the 'effective rank' of the matrix, which is essential for the calculation of the precision matrix. For more details, we refer the readers to the detailed discussion in [24, 45].

In addition to the above challenge due to the curvature, the presence of noise particularly in the manifold setting imposes another challenge—that is, how to find the *true* neighbors? Specifically, note that the *neighboring points* of x denoted by  $\chi_{\mathcal{O}(x)}$  need to be identified from the noisy samples. When  $\sigma = 0$ , i.e. in the noiseless case, a neighbor can be easily identified if the diameter of  $\mathcal{O}(x)$  is sufficiently small. However, when  $\sigma > 0$ , i.e. in the presence of noise, it is not clear if a neighbor determined from the noisy point cloud is truly a neighbor. Concretely, let  $\mathscr{X} = \{x_i\}_{i=1}^n \subset \iota(M) \subset \mathbb{R}^p$  denote a set of identical and independent (i.i.d.) random samples from X, and  $\mathscr{Y} = \{y_i\}$ , where  $y_i = x_i + \sigma \xi_i$  is sampled from Y. Then, in general,  $\|y_i - y_j\|_{\mathbb{R}^p} \leq \epsilon$  does not imply  $\|x_i - x_j\|_{\mathbb{R}^p} \leq \epsilon$ . Thus, a naive approach to determine neighbors might fail. To the best of our knowledge, there are only few existing algorithms

for determining neighbors when the point cloud is noisy. For example, determining a neighbor by the diffusion distance [5], which has a solid theoretical support when noise exists [12, 13]. Since finding neighbors is still a challenging problem on its own, in this paper, we focus only on estimating the precision matrix, and subsequently, the MD, assuming that the *true* neighbors are known.

M. GAVISH ET AL.

### 3.2 Mahalanobis distance

We are now ready to define the MD: we start with the definition of the MD under the manifold model, and then for reference, we present the typical definition under the global linear spiked model.

Under the manifold model, since the covariance matrix may have no strictly zero eigenvalues even in the noiseless case due to the curvature, we consider the following definition of the MD.

Definition 3.1 MD under the manifold model. Suppose  $x \in M$ . The MD between  $z \in \iota(M)$  and  $X_x$  is defined by

$$d_{\Sigma_x}^2(z, X_x) := (z - \mu_x)^\top \mathscr{I}_d(\Sigma_x)(z - \mu_x), \tag{3.2}$$

where  $\mathscr{I}_d(\Sigma_x)$  is the truncated pseudo-inverse of degree d defined by

$$\mathscr{I}_d(\Sigma_x) := U \operatorname{diag}(1/\lambda_1, \dots, 1/\lambda_d, 0, \dots, 0) U^{\top},$$

 $\Sigma_x = U \operatorname{diag}(\lambda_1, \dots, \lambda_p) U^{\top}$  is the eigendecomposition of  $\Sigma_x$ , and  $\lambda_1 \ge \lambda_2 \ge \dots$ 

Note that the knowledge of the manifold dimension d is required for this definition; yet, it is in general not available and needs to be estimated.

For comparison, consider also the classical definition of the MD under the global linear spiked model.

Definition 3.2 MD under the linear spiked model. The MD between an arbitrary point  $z \in \mathbb{R}^p$  and the underlying signal distribution X (2.5) is defined by

$$d_{\Sigma_X}^2(z,X) = (z-\mu)^{\top} \Sigma_X^{\dagger}(z-\mu),$$
(3.3)

where † denotes the Moore–Penrose pseudo-inverse.

Let us take a closer look at the latter definition. Since  $\Sigma_X$  is semi-positive definite, by the Cholesky decomposition, we have  $\Sigma_X^{\dagger} = WW^{\top}$ , where  $W \in \mathbb{R}^{p \times d}$ . Hence,

$$d_{\Sigma_X}^2(z,X) = (z-\mu)^{\top} (WW^{\top})(z-\mu) = \|W^{\top}(z-\mu)\|_{\mathbb{R}^p}^2,$$

which indicates that geometrically, MD evaluates the relationship between z and (the mean of) X by a proper linear transform. In Section 3.3, we relate  $W^{\top}$  to the inverse of the Jacobian of arbitrary unknown observation functions and show that it gives rise to an important invariance property of the MD in the context of manifold learning. Here, we only demonstrate a primary merit of MD, which stems from its invariance to rotation and rescaling. Importantly, this invariance property holds for both definitions: under the linear spiked model as well as under the manifold model. Consider the linear spiked model and a random variable  $\tilde{X} = cAX$ , where  $c \in \mathbb{R}$  models rescaling and  $A \in O(p)$  models rotation. Here,

O(p) denotes the group of p-by-p orthogonal matrices. The population covariance matrix of  $\widetilde{X}$  is

$$\Sigma_{\widetilde{X}} = c^2 A \Sigma_X A^\top \tag{3.4}$$

and its population mean is rotated and rescaled to  $\tilde{\mu} = cA\mu$ . To demonstrating the invariance, suppose  $\tilde{z} = cAz$  and observe the MD between  $\tilde{z}$  and  $\tilde{X}$ 

$$d_{\Sigma_{\widetilde{X}}}^2(\widetilde{z},\widetilde{X}) = (cA(z-\mu))^\top \Sigma_{\widetilde{X}}^\dagger(cA(z-\mu)) = (z-\mu)^\top \Sigma_X^\dagger(z-\mu) = d_{\Sigma_X}^2(z,X).$$

The same argument holds for the MD under the manifold model; for brevity, we omit the details.

Now, recall that the goal in this work is to estimate the MD between a point  $z \in \mathbb{R}^p$  and X, when we only have access to noisy data sampled from Y in (2.7). Concretely, assume that  $y_1, \ldots, y_n \stackrel{\text{iid}}{\sim} Y$  is a sample of n data points. Since  $\Sigma_X$  is unknown, the quantity  $d_{\Sigma_X}(z, X)$ , or  $\Sigma_X^{\dagger}$  in (3.3), must be estimated from the (noisy) data. For simplicity, we assume below that  $\mu$  and  $\sigma$  are known; these assumptions can easily be removed in applications to real data. As discussed in Section 3.1, the notable challenge in estimating the MD from local samples of Y is the estimation of the precision matrix via the pseudoinverse, stemming from the interaction of the small eigenvalues of  $\Sigma_X$  and the noise. In addition, as discussed in Section 3.1, the main difference between (2.5) and (2.2) is that in (2.5), the covariance matrix  $\Sigma_X$  is global and strictly of low rank, whereas in (2.2), the covariance matrix  $\Sigma_x$  depends on x, and its rank is in general higher than d when the manifold is of dimension d, yet only its first d eigenvalues are dominant. In the context of MD, this rank difference might lead to unwanted consequences if we consider Definition 3.2. Below, we demonstrate it with an example.

Assume for simplicity that the pdf is uniform,  $\mathcal{O}(x) = B_{\epsilon}^{\mathbb{R}^p}(\iota(x))$  and  $\epsilon$  is sufficiently small. From (3.1), we observe a clear spectral gap between the first *d* eigenvalues and the remaining ones. Therefore, we would expect that the behavior of  $X_x$  in (2.2) is similar to that of *X* in (2.5). However, we need to be careful with the associated precision matrices, and hence, with the two definitions of MD. Suppose that the main interest in estimating the MD is recovering the geodesic distance when the manifold cannot be directly accessed. It has been shown in [24, Theorem 8 (3)] that even if the manifold can be directly accessed, defining the MD between a point  $z \in \iota(M)$  and  $X_x$  by  $\tilde{d}_{\Sigma_x}^2(z, X_x) = (z - x)^{\top} \Sigma_x^{\dagger}(z - x)$  as in (3.3) might lead to a biased estimate of the geodesic distance. To be more precise, when the rank of  $\Sigma_x$  is greater than *d* and  $||x - z||_{\mathbb{R}^p} = t$  is sufficiently small, asymptotically we have

$$\frac{d(d+2)}{|S^{d-1}|\epsilon^{d+2}}\tilde{d}_{\Sigma_x^{\dagger}}(z,X_x) = t + O(t), \qquad (3.5)$$

which means that the MD cannot even recover the basic geodesic distance since the order of the error is of the same order as the geodesic distance. Broadly, this bias is the result of the interaction of the small eigenvalues related to the curvature, particularly the second fundamental form, and the fact that the vector z - x is not intrinsic to the manifold and might contain component that is normal to the tangent space at x.

### 3.3 Motivating example

A motivating example for the centrality of the MD under the manifold model in data analysis was presented in [31] and later elaborated in the empirical intrinsic geometry (EIG) framework presented in [37, 38] with various applications, e.g., [9–11, 17, 22, 40, 46].

#### M. GAVISH ET AL.

Suppose the manifold M is merely an image of a different, intrinsic, and inaccessible manifold of interest N, that is  $M = \phi(N)$ , where  $\phi$  is a diffeomorphic map. We may call N the latent space. In applications,  $\phi$  could represent the distortion of the data introduced by some measurement equipment. This 'distortion model' is often of key important since it critically affects the intrinsic information N we have interest in, even if noise does not exist. Under this model, it is shown in [31, Section 3] that when M and N are both Euclidean spaces, for  $x, \bar{x} \in M$ , the following holds:

$$\|\theta - \bar{\theta}\|_{2}^{2} = \frac{1}{2} (x - \bar{x})^{\top} \left[ C^{\dagger} + \bar{C}^{\dagger} \right] (x - \bar{x}) + O(\|x - \bar{x}\|_{2}^{4}), \qquad (3.6)$$

where  $\theta = \phi^{-1}(x)$ ,  $\bar{\theta} = \phi^{-1}(\bar{x})$  and  $C = \nabla \phi|_{\theta} \nabla \phi|_{\theta}^{\top}$ ,  $\bar{C} = \nabla \phi|_{\bar{\theta}} \nabla \phi|_{\bar{\theta}}^{\top}$ . A similar statement for the case when *M* and *N* are both manifolds and  $\phi$  is a diffeomorphism is given in [24]. It was further shown in [31] and [24] that if we have i.i.d. sampled from *M* adhering to the statistical manifold models discussed above, then

$$\Sigma_{x} = \nabla \phi|_{\theta} \nabla \phi|_{\theta}^{\top}$$
(3.7)

for a sufficiently small neighborhood  $\mathcal{O}(x)$ . Remarkably, by combining (3.6) and (3.7), a small variant of the MD can recover the distance between  $\theta$  and  $\overline{\theta}$  from the hidden intrinsic manifold  $M_{\theta}$  based on samples from M without explicit information on the map  $\phi$ , thereby solving a completely blind inverse problem.

The above distortion model was studied in [37] in the context of a nonlinear dynamical system. Denote the dataset or the point cloud as  $\mathscr{U} := \{u_j\}_{j=1}^n \subset \mathbb{R}^q$ , where  $u_j$  is sampled at the *j*th time stamp. The key assumption is that  $\mathscr{U}$  comes from observing an inaccessible intrinsic dynamics  $\theta(t) \in \mathbb{R}^p$  that satisfies the stochastic differential equation (SDE)

$$d\theta(t) = a(\theta(t))dt + d\omega(t), \qquad (3.8)$$

where *a* is an unknown drift function and  $\omega$  is the standard *d*-dimensional Brownian motion. We call the subset of  $\mathbb{R}^p$  that hosts  $\theta(t)$  the *phase space*; for example, an open subset of  $\mathbb{R}^p$ , or a smooth manifold embedded in  $\mathbb{R}^p$ .

The observation is modeled by a diffeomorphic function  $\Phi : \mathbb{R}^p \to \mathbb{R}^q$  so that  $u_j = \Phi(\theta_j)$ , where  $\theta_j$  is sampled from the intrinsic dynamics  $\theta(t)$  at the *j*-th time stamp. We call  $\Phi$  the *observation transform*. Based on (3.8), and the fact that

$$du_t = \left(\frac{1}{2}\Delta\Phi|_{\theta_t} + \nabla\Phi|_{\theta_t}a(\theta_t)\right)dt + \nabla\Phi|_{\theta_t}d\omega_t$$
(3.9)

by the Ito's formula, we obtain that

$$\operatorname{Cov}\left(\mathrm{d}u_{t}\right) = \nabla \Phi|_{\theta_{t}} \nabla \Phi|_{\theta_{t}}^{\top},\tag{3.10}$$

since  $(\frac{1}{2}\Delta\Phi|_{\theta_t} + \nabla\Phi|_{\theta_t}a(\theta_t))dt$  is the drift. By combining the above facts, it is shown in [31, Section 3] that when  $u_i$  and  $u_i$  are sufficiently close and the phase space is flat, we could recover the intrinsic

distance between  $\theta_i$  and  $\theta_j$  by

$$\|\theta_i - \theta_j\|_{\mathbb{R}^d}^2 = \frac{1}{2} (u_i - u_j)^\top \left[ C_i^\dagger + C_j^\dagger \right] (u_i - u_j) + O(\|u_i - u_j\|^4), \qquad (3.11)$$

where  $C_i = \nabla \Phi(\theta_i) [\nabla \Phi(\theta_i)]^\top$  is the covariance matrix associated with the observation process (i.e. the deformed Brownian motion). Furthermore, it is shown in [31, 37] that  $C_i$  can be estimated by the covariance matrix of  $\{u_k\}_{k=i-L}^{i+L}$ , where  $L \in \mathbb{N}$  is chosen by the user. The key relevant fact here is that since the covariance matrix of  $du_t$  is  $\nabla \Phi|_{\theta_t} \nabla \Phi|_{\theta_t}^\top$ , the related open set  $O(u_t)$  is an ellipsoid with principle semi-axes described by the non-degenerate eigenvalues and eigenvectors of  $\nabla \Phi|_{\theta_t} \nabla \Phi|_{\theta_t}^\top$ .

To conclude, the intrinsic signal model in (3.8) and the observation model (3.9) are generic and can describe a broad range of applications. Therefore, the ability to recover the distances between samples of the intrinsic process from observations in a non-parametric and unsupervised manner using an estimate of the MD is very powerful. Indeed, this model was used for system identification [40], molecular dynamics [10], sleep analysis [22, 46], model reduction [11], speech processing [9] and gene expression data [17].

#### 3.4 Shrinkage estimators

For any *p*-by-*p* matrix  $M_n$  estimated from  $y_1, \ldots, y_n$ , consider the estimator for MD

$$d_{M_n}^2(z,X) = (z-\mu)^{\top} M_n(z-\mu)$$
(3.12)

using Definition 3.2. The extension to Definition 3.1 is straight-forward. In order to quantitatively measure the performance of any MD estimator  $d_{M_n}(z, X)$ , it is useful to introduce a loss function. For any estimator of the form (3.12), the absolute value of the estimation error with respect to the true value (3.3) is

$$\left| d_{\Sigma_X}^2(z,X) - d_{M_n}^2(z,X) \right| = \left| (z-\mu)^\top [\Sigma_X^\dagger - M_n](z-\mu) \right|.$$

As the test vector z is arbitrary, it is natural to consider the *worst* case and define the loss of  $M_n$  at the (unknown) underlying low-dimensional covariance  $\Sigma_X$ :

DEFINITION 3.3. The worst case loss function of an estimator of the form (3.12) for (3.3) is defined as

$$L_n(M_n, \Sigma_X^{\dagger}) := \sup_{z: \, \|z-\mu\|_{\mathbb{R}^p} = 1} \left| (z-\mu)^{\top} [\Sigma_X^{\dagger} - M_n] (z-\mu) \right| = \|\Sigma_X^{\dagger} - M_n\|_{\text{op}},$$
(3.13)

where  $\|\cdot\|_{op}$  is the matrix operator norm.

It is also reasonable to consider the root mean squared estimation error of all possible test vectors. The discussion below follows the same line. To keep the notation light, the dependence of  $L_n$  on  $\mu$  and  $\sigma$  as well as the dependence of  $M_n$  on the sample  $y_1, \ldots, y_n$  are implicit.

Consider matrices of the form  $M_n^{\eta} := \eta(S_n)$ , where  $\eta : [0, \infty) \to [0, \infty)$  and  $S_n$  is the sample covariance. We call  $M_n^{\eta}$  the *shrinkage estimator* of  $\Sigma_X^{\dagger}$  with  $\eta$ . A typical example is the *classical MD* 

*estimator*, which is a shrinkage estimator with  $\eta = \eta_{\sigma}^{\text{classical}}$ , where

$$\eta_{\sigma}^{\text{classical}}(\alpha) = \begin{cases} 1/(\alpha - \sigma^2) & \alpha > \sigma^2 \\ 0 & \alpha \le \sigma^2 \end{cases} .$$
(3.14)

From [23], in the traditional setup when the dimension *p* is fixed and  $n \to \infty$ , the classical MD estimator obtains zero loss asymptotically.

THEOREM 3.1. Let p be fixed independently of n. Then,

$$\lim_{n \to \infty} L_n(\eta_{\sigma}^{\text{classical}}(S_n), \Sigma_X^{\dagger}) = 0.$$

*Proof.* Since it is well known that  $(S_n - \sigma^2 I_p) \to \Sigma_X$  as  $n \to \infty$ , substituting  $M_n$  with  $\eta_{\sigma}^{\text{classical}}(S_n)$  in (3.13) and taking limit with  $n \to \infty$  complete the proof.

When p grows with n, such that  $p = p_n \to \infty$  with  $p_n/n \to \beta > 0$ , the situation is quite different. It is known that in this situation, the sample covariance matrix is an inconsistent estimate of the population covariance matrix [15], and Theorem 3.1 might not hold; that is, the classical MD estimator might not be optimal. The following questions naturally arise when  $\beta > 0$ :

- 1. Is there an optimal shrinkage (OS) estimator with respect to the loss  $L_n$ ?
- 2. How does the loss of the optimal shrinkage estimator compare with the loss  $L_n(\eta_{\sigma}^{\text{classical}}(S_n), \Sigma_x^{\dagger})$ ?

In the sequel, we attempt to answer these questions.

#### 4. Optimal recovery of precision matrix for MD under the spiked model

We start the derivation of the OS for MD under the linear spiked model, which involves the OS for the precision matrix. Its extension to the manifold model requires only an additional mild condition, which will be discussed in Section 5. Without loss of generality, we set the noise level  $\sigma = 1$  and will discuss the general case subsequently.

Assumption 4.1 Asymptotic( $\beta$ ). The number of variables  $p = p_n$  grows with the number of observations *n*, such that  $p/n \rightarrow \beta$  as  $n \rightarrow \infty$ , for  $0 < \beta \le 1$ .

Assumption 4.2 Spiked model. Suppose  $\Sigma_X = \Sigma_Y - \sigma^2 I_p$  with the eigendecomposition

$$\Sigma_X = U \begin{bmatrix} \Sigma_d & 0\\ 0 & 0_{p-d} \end{bmatrix} U^\top \in \mathbb{R}^{p \times p}, \qquad (4.1)$$

where  $d \ge 0$ ,  $\Sigma_d = \text{diag}(\ell_1, \dots, \ell_d)$  is a  $d \times d$  matrix whose diagonal consists of d spikes  $\ell_1 > \dots > \ell_d > 0$ , which are fixed and independent of p and n, and the off-diagonal elements are set to zero. For completeness, denote  $\ell_{d+1} = \dots = \ell_p = 0$ . Note that we assume that all spikes are simple. When d = 0, it is the null case.

1184

Denote the eigendecomposition of  $S_n$  as

$$S_n = V_n \operatorname{diag}(\lambda_{1,n}, \dots, \lambda_{p,n}) V_n^\top \in \mathbb{R}^{p \times p},$$
(4.2)

where  $\lambda_{1,n} \ge ... \lambda_{p,n} \ge 0$  are the empirical eigenvalues and  $V_n \in O(p)$  is the matrix, whose columns are the empirical eigenvectors  $v_{i,n} \in \mathbb{R}^p$ , i = 1, ..., p. Under Assumptions 4.1 and 4.2, results collected from [2, 3, 25, 28] imply three important facts about the sample covariance matrix  $S_n$ .

 Eigenvalue spread. Suppose Assumption 4.1 holds and considers the null case where Σ<sub>d</sub> = 0. As n → ∞, the spread of the empirical eigenvalues λ<sub>i,n</sub> converges to a continuous distribution called the 'Marcenko–Pastur' law [25]

$$\frac{\sqrt{(\lambda_{+} - x)(x - \lambda_{-})}}{2\pi\beta x} \mathbf{1}_{[\lambda_{-}, \lambda_{+}]} dx, \qquad (4.3)$$

where  $\lambda_{+} = (1 + \sqrt{\beta})^2$  and  $\lambda_{-} = (1 - \sqrt{\beta})^2$  are the *limiting bulk edges*.

2. Top eigenvalue bias. Suppose Assumptions 4.1 and 4.2 hold. For  $1 \le i \le d$ , the empirical eigenvalues

$$\lambda_{i,n} \xrightarrow{a.s.} \lambda(\ell_i) =: \lambda_i$$

as  $n \to \infty$ , where

$$\lambda(\alpha) = \begin{cases} 1 + \alpha + \beta + \frac{\beta}{\alpha} & \alpha > \ell_+ \\ (1 + \sqrt{\beta})^2 & 0 \le \alpha \le \ell_+ \end{cases}$$
(4.4)

is defined on  $\alpha \in [0, \infty)$  and  $\ell_+ := \sqrt{\beta}$ . For  $d + 1 \le i \le p$ , since  $\ell_i = 0$ , the empirical eigenvalues  $\lambda_{i,n}$  follow the Marcenko–Pastur law (4.3).

3. Top eigenvector inconsistency. Suppose Assumptions 4.1 and 4.2 hold. Let  $c_{i,n}$  and  $s_{i,n}$  be the cosine and sine values of the angle between the *i*th population eigenvector and the *i*th empirical eigenvector after properly adjusting the sign of each empirical eigenvector, respectively. Note that there exists a sequence of  $R_n \in O(p)$  so that  $R_n V_n$  converges almost surely (a.s.) to  $V \in O(p)$ . In the following, we assume that the empirical eigenvectors have been properly rotated. It is known that when  $n \to \infty$ ,  $c_{i,n} \xrightarrow{a.s.} c(\ell_i)$  and  $s_{i,n} \xrightarrow{a.s.} s(\ell_i)$ , where

$$c(\alpha) := \begin{cases} \sqrt{\frac{\alpha^2 - \beta}{\alpha^2 + \beta \alpha}} & \alpha > \ell_+ \\ 0 & 0 \le \alpha \le \ell_+ \end{cases}, \tag{4.5}$$

and

$$s(\alpha) := \sqrt{1 - c^2(\alpha)} \tag{4.6}$$

are defined on  $\alpha \in [0, \infty)$ .

#### M. GAVISH ET AL.

The above three properties imply that the classical estimator  $\eta_{\sigma}^{\text{classical}}(S_n)$  may not be the best estimator in general, and for the purpose of estimating MD in particular. Inspired by [8], we may 'correct' the bias of the eigenvalues to improve the estimation.

DEFINITION 4.1. The asymptotic loss function is defined as

$$L_{\infty}(\eta|\ell_1,\ldots,\ell_d) := \lim_{n \to \infty} L_n(M_n^{\eta}, \Sigma_X^{\dagger}), \qquad (4.7)$$

assuming the limit exists.

To find a shrinkage estimator  $\eta$  that minimizes  $L_{\infty}(\eta|\ell_1, \ldots, \ell_d)$ , it is natural to construct the estimator by recovering the spikes  $\ell_i$  using the biased eigenvalues  $\lambda_i$ . From the inversion of (4.4), recalling that  $\ell_+ = \sqrt{\beta}$ , we can define

$$\ell(\alpha) := \frac{\alpha + 1 - \beta + \sqrt{(\alpha + 1 - \beta)^2 - 4\alpha}}{2} - 1$$
(4.8)

when  $\alpha > \lambda_+$  and consider the shrinkage function

$$\eta^{\text{inv}}(\alpha) = \begin{cases} 1/\ell(\alpha) & \alpha > \lambda_+ \\ 0 & \text{otherwise.} \end{cases}$$
(4.9)

Note that since taking inverse of a matrix is nonlinear, it is not clear if the above naive idea will lead to the OS estimator of the precision matrix for the MD estimate. Nevertheless, it is reasonable to expect the existence of an optimal shrinkage function  $\eta^*$  satisfying

$$L_{\infty}(\eta^*|\ell_1,\ldots,\ell_d) \le L_{\infty}(\eta^{\mathrm{inv}}|\ell_1,\ldots,\ell_d)$$

for any spikes  $\ell_1, \ldots, \ell_d$ . Below we show that this naive idea,  $\eta^{inv}$ , is in fact the OS estimator for our purpose.

#### 4.1 Derivation of the optimal shrinker when $\sigma = 1$

DEFINITION 4.2. A function  $\eta : [0, \infty) \to [0, \infty)$  is called a *shrinker* if it is continuous when  $\lambda > \lambda_+$ , and  $\eta(\lambda) = 0$  when  $0 \le \lambda \le \lambda_+$ .

Note that this shrinker is a *bulk shrinker* considered in [8, Definition 3]. Based on the assumption of a shrinker  $\eta$ , the associated shrinkage estimator converges almost surely, that is

$$M_n^{\eta} \xrightarrow{a.s.} M^{\eta} := V \operatorname{diag}(\eta(\lambda_1), \dots, \eta(\lambda_p)) V^{\top}, \qquad (4.10)$$

where the right-hand side is the eigendecomposition of  $M^{\eta}$ . Thus, the sequence of loss functions also almost surely converges as

$$L_n(M_n^{\eta}, \Sigma_X^{\dagger}) = \|\Sigma_X^{\dagger} - M_n^{\eta}\|_{op} \xrightarrow{a.s.} \|\Sigma_X^{\dagger} - M^{\eta}\|_{op}.$$

$$(4.11)$$

As a result, the limit in (4.7) exists when  $\eta$  is a shrinker, and we have the following theorem which in turn gives rise to the optimal shrinker. Note that while the biased eigenvalues could be recovered by the quadratic relationship between the biased eigenvalues and the population eigenvalues (4.4), true (population) eigenvectors and empirical eigenvectors are not collinear [8] and so far we do not have a way to recover the biased eigenvectors. The main idea beyond the proof is respecting this fact, and when we find the optimal way to correct the eigenvalues, the biased eigenvectors should be taken into account. With the control of this eigenvector bias, the OS will be derived.

Theorem 4.1 Characterization of the asymptotic loss. Suppose  $\sigma = 1$ . Consider the spiked covariance model satisfying Assumptions 4.1 and 4.2 and a shrinkage function  $\eta : [0, \infty) \rightarrow [0, \infty)$ . We have a.s.

$$L_{\infty}(\eta|\ell_1,\ldots,\ell_d) = \max_{i=1,\ldots,p} \{ \Delta(\ell_i,\eta(\lambda_i)) \}, \qquad (4.12)$$

where  $\Delta : [0, \infty) \times [0, \infty) \rightarrow [0, \infty)$  is given by

$$\Delta(\alpha,\zeta) = \begin{cases} u_{+}(\alpha,\zeta) & \alpha > \ell_{+} \text{ and} \zeta \le \frac{1}{\alpha} \\ -u_{-}(\alpha,\zeta) & \alpha > \ell_{+} \text{ and} \zeta > \frac{1}{\alpha} \\ 1/\alpha & 0 < \alpha \le \ell_{+} \\ 0 & \alpha = 0 , \end{cases}$$
(4.13)

where

$$u_{+}(\alpha,\zeta) = \frac{1}{2} \left( \frac{1}{\alpha} - \zeta + \sqrt{\left(\frac{1}{\alpha} - \zeta\right)^{2} + 4\frac{\zeta s(\alpha)^{2}}{\alpha}} \right), \qquad (4.14)$$

$$u_{-}(\alpha,\zeta) = \frac{1}{2} \left( \frac{1}{\alpha} - \zeta - \sqrt{\left(\frac{1}{\alpha} - \zeta\right)^2 + 4\frac{\zeta s(\alpha)^2}{\alpha}} \right).$$
(4.15)

*Proof.* Based on the property of 'simultaneous block-diagonalization' for  $\Sigma_X^{\dagger}$  and  $M_n^{\eta}$  in [8, Section 2], the properties of 'orthogonal invariance' and 'max-decomposability' for the operator norm in [8, Section 3], and the convergence of  $c_{i,n}$  and  $s_{i,n}$  in (4.5) and (4.11), we have

$$L_n(M_n^{\eta}, \Sigma_X^{\dagger}) = \max_i \|A_i - B_{i,n}\|_{\text{op}},$$

where

$$A_i = \begin{bmatrix} 1/\ell_i & 0\\ 0 & 0 \end{bmatrix}$$



FIG. 1. The asymptotic loss for several  $\beta = p/n$  values as a function of the spike strength in a single spike model, i.e.  $\ell_1 \in [0.1, 3]$ , and  $\ell_2 = \ell_3 = \ldots = \ell_d = 0$ . The noise level is fixed and set to  $\sigma = 1$ .

when  $\ell_i \neq 0$  and  $A_i = 0_{2 \times 2}$  otherwise, and

$$B_{i,n} = \eta(\lambda_{i,n}) \begin{bmatrix} c_{i,n}^2 & c_{i,n}s_{i,n} \\ c_{i,n}s_{i,n} & s_{i,n}^2 \end{bmatrix}.$$

When  $n \to \infty$ , the loss converges a.s. to  $\max_i ||A_i - B_i||_{op}$ , where

$$B_i = \eta(\lambda_i) \begin{bmatrix} c(\ell_i)^2 & c(\ell_i)s(\ell_i) \\ c(\ell_i)s(\ell_i) & s(\ell_i)^2 \end{bmatrix}.$$

Now, we evaluate  $||A_i - B_i||_{op}$  for different  $\ell_i$ .

When  $\ell_i > \ell_+$ , denote the eigenvalues of  $A_i - B_i$  as  $u_+(\ell_i, \eta(\lambda_i))$  and  $u_-(\ell_i, \eta(\lambda_i))$ . If  $\eta(\lambda_i) > 1/\ell_i$ , we have  $0 \le u_+(\ell_i, \eta(\lambda_i)) \le -u_-(\ell_i, \eta(\lambda_i))$ , and hence,  $||A_i - B_i||_{op} = -u_-(\ell_i, \eta(\lambda_i))$ ; otherwise, we have  $u_+(\ell_i, \eta(\lambda_i)) \ge -u_-(\ell_i, \eta(\lambda_i)) \ge 0$ , and hence,  $||A_i - B_i||_{op} = u_+(\ell_i, \eta(\lambda_i))$ . For  $0 < \ell_i \le \ell_+$ , since  $c(\ell_i) = 0$ , we have

$$B_i = \begin{bmatrix} 0 & 0 \\ 0 & \eta(\lambda_i) \end{bmatrix},$$

which equals  $0_{2\times 2}$  since  $\eta(\lambda_i) = 0$  by the definition of shrinkage function. Thus,  $||A_i - B_i||_{op} = 1/\ell_i$ . Finally, for  $\ell_i = 0$ ,  $A_i$  is a 2 × 2 zero matrix, and thus,  $||A_i - B_i||_{op} = \eta(\lambda_i) = 0$ . This concludes the proof.

Figure 1 illustrates the obtained asymptotic loss for several  $\beta = p/n$  values as a function of the spike strength in a single spike model. It is clear that for each  $\beta$ , there is a transition at  $\ell_+ = \sqrt{\beta}$ . An immediate consequence of Theorem 4.1 is that  $\eta^{\text{inv}}$  is an optimal shrinker.

COROLLARY 4.1. Suppose  $\sigma = 1$  and Assumptions 4.1 and 4.2 hold. Define the asymptotically optimal shrinkage function as

$$\eta^* := \arg\min_{\eta} \{ L_{\infty}(\eta | \ell_1, \dots, \ell_d) \}, \qquad (4.16)$$

where argmin is evaluated on the set of all possible shrinkage functions. Then,  $\eta^*$  is unique and equals  $\eta^{\text{inv}}$  given in (4.9). Moreover, its associated loss is

$$\max_{i} \left\{ \Delta(\ell(\lambda_i)) \right\}, \tag{4.17}$$

where

$$\Delta(\alpha) = \begin{cases} s(\alpha)/\alpha = \frac{\sqrt{\beta}}{\alpha^{3/2}} \sqrt{\frac{1+\alpha}{\beta+\alpha}} & \alpha > \ell_+ \\ 1/\alpha & 0 < \alpha \le \ell_+ \\ 0 & \alpha = 0 \,. \end{cases}$$
(4.18)

Note that this result coincides with the findings reported in [8]. Precisely, it is shown in [8, (1.12)] that for the operator norm,  $\ell(\alpha)$  (4.8) is the optimal shrinkage for the covariance matrix and precision matrix. In this corollary, we show that for the MD, which is related to the precision matrix, the optimal estimator is also achieved by the optimal shrinkage, taking  $\ell(\alpha)$  into account.

*Proof.* Based on Theorem 4.1, the optimal shrinker  $\eta^*$  leads to  $\min_{\eta \ge 0} \max_{i=1,...,d} \{\Delta(\ell_i, \eta(\lambda_i))\}$ . Note that for  $j = \arg\max_{i=1,...,d} \{\Delta(\ell_i, \eta(\lambda_i))\}$ , the optimal shrinker achieves  $\min_{\eta \ge 0} \{\Delta(\ell_j, \eta(\lambda_j))\}$ . Thus, by the same argument in [8], if we could solve  $\arg\min_{\eta \ge 0} \{\Delta(\alpha, \eta(\lambda(\alpha)))\}$  for any  $\alpha > 0$ , we find the optimal shrinker. To simplify the notation, we abbreviate  $\eta(\lambda(\alpha))$  by  $\eta$ .

For  $\alpha > \ell_+$  and  $\eta > \frac{1}{\alpha}$ , we have  $\Delta(\alpha, \eta) = -u_-(\alpha, \eta)$ . By a direct calculation, we get

$$\partial_{\eta} \Delta(\alpha, \eta) = \frac{1}{2} + \frac{-(\frac{1}{\alpha} - \eta) + \frac{2s(\alpha)}{\alpha}}{\sqrt{(\frac{1}{\alpha} - \eta)^2 + 4\frac{\eta s(\alpha)^2}{\alpha}}} > 0.$$

For  $\alpha > \ell_+$  and  $0 \le \eta \le \frac{1}{\alpha}$ , we have  $\Delta(\alpha, \eta) = u_+(\alpha, \eta)$ , and similarly by taking the derivative of (4.14), we have

$$\partial_{\eta} \Delta(\alpha, \eta) = -\frac{1}{2} + \frac{-(\frac{1}{\alpha} - \eta) + \frac{2s(\alpha)}{\alpha}}{2\sqrt{(\frac{1}{\alpha} - \eta)^2 + 4\frac{\eta s(\alpha)^2}{\alpha}}} \ge 0.$$

As a result, the partial derivative of the loss function is decreasing when  $0 \le \eta \le 1/\alpha$  and increasing when  $\eta > 1/\alpha$  with a discontinuity at  $\eta = 1/\alpha$ , while the loss function is continuous. Thus, the loss function reaches the minimum when  $\eta = 1/\alpha$ . These facts imply that  $\eta^*(\lambda_i) = 1/\ell(\lambda_i)$  when  $\lambda_i > \ell_+$ . Furthermore, by substituting  $\eta$  with  $\eta^*$  in (4.14) or (4.15), we get  $\Delta(\alpha) = s(\alpha)/\alpha$ . By definition,  $\eta^* = 0$  when  $0 \le \alpha \le \ell_+$ . Thus, for  $0 < \alpha \le \ell_+$ ,  $\Delta(\alpha) = 1/\alpha$ , and for  $\ell = 0$ ,  $\Delta(\ell) = 0$ . Finally, it is clear



FIG. 2. The obtained optimal shrinker with the classical shrinker overlay, for  $\beta = p/n = 1$  and  $\sigma = 1$ . To enhance the visualization, the y-axis of the figure is truncated at 1.5 on the right-hand side.

that  $\eta^*$  is continuous when  $\alpha > \lambda_+$ , and  $\eta(\alpha) = 0$  when  $0 \le \alpha \le \lambda_+$ . We thus conclude that  $\eta^*$  is the optimal shrinker.

REMARK 4.1. We compare our result with another naive approach; that is, obtaining the covariance by the optimal shrinkage with respect to the operator norm, and then taking the Moore–Penrose pseudo-inverse. Let  $\eta^{cov}$  denote the optimal shrinker for the covariance matrix recovery obtained from [8] with respect to the operator norm loss function. With the same notation, that is  $\Sigma_Y = \Sigma_X + \sigma^2 I_p$  and  $\sigma = 1$  as in (4.2), the optimal shrinker satisfies

$$\eta^{cov}(\alpha) = \begin{cases} \ell(\alpha) + 1 & \alpha > \lambda_+ \\ 1 & \text{otherwise} , \end{cases}$$
(4.19)

where  $\ell$  is defined in (4.8). Note that in [8], the authors aimed to recover  $\Sigma_Y$ , while we recover  $\Sigma_X^{\dagger}$ . In other words, the authors in [8, (1.12)] showed that when the operator norm is considered, the OS estimator is the same as correcting the eigenvalues according to the relationship (4.4). Thus, our result  $\eta^*$  coincides with the inverse of  $\eta^{cov} - 1$  when  $\alpha > \lambda_+$ .

We note the following interesting phenomenon stemming from Theorem 4.1 and Corollary 4.1. If there exists a non-trivial spike  $\ell_i > 0$  that is weak enough so that  $\ell_i$  is sufficiently small compared with  $\ell_+$ , then  $L_{\infty}(\eta | \ell_1, \dots, \ell_d)$  is dominated by  $1/\ell_i$ . Consequently, in this large *p* large *n* regime, we cannot 'rescue' this spike, and the associated signal is lost in the noise, as can be seen in Corollary 4.1.

Figure 2 illustrates the obtained optimal shrinker with the classical shrinker overlay, for  $\beta = p/n = 1$  and  $\sigma = 1$ . Clearly, compared with the classical shrinker, the obtained optimal shrinker truncates the eigenvalues more aggressively.

#### 4.2 Derivation of the optimal shrinker when $\sigma \neq 1$

To handle the general case when  $\sigma \neq 1$ , we first re-scale the data and model by setting  $\ell'_i := \ell_i / \sigma^2$  and  $\lambda'_{i,n} := \lambda_{i,n} / \sigma^2$  and consider the following shrinker defined on  $[0, \infty)$ :

$$\eta_{\sigma}(\alpha) := \frac{\eta(\alpha/\sigma^2)}{\sigma^2} \,. \tag{4.20}$$

Note that since  $\eta$  plays the role of estimating the precision matrix, we re-normalize it by dividing  $\eta(\alpha/\sigma^2)$  by  $\sigma^2$ . The shrinkage estimator for  $\Sigma_X^{\dagger}$  becomes  $M_n^{\eta_\sigma} := \eta_{\sigma}(S_n)$ , the general optimal shrinker becomes

$$\tilde{\eta}^*(\alpha) = \begin{cases} \frac{1}{\sigma^2 \ell(\alpha/\sigma^2)} & \alpha > \sigma^2 \ell_+ \\ 0 & 0 \le \alpha \le \sigma^2 \ell_+ \end{cases}$$
(4.21)

and the associated loss is

$$\max_{i} \left\{ \frac{\Delta(\ell(\frac{\lambda_i}{\sigma^2}))}{\sigma^2} \right\} \, .$$

#### 5. Extension to the manifold model

We now come back to the manifold learning problem with the manifold setup described in Section 2.1. We argue that despite the challenge mentioned in Section 3.1, the developed theorem in Section 4.1 can be extended to study the manifold model in the large p large n setup with proper modifications. Note that For each n, there exists an orthonormal basis  $\{u_{n,l}\}_{l=1}^{p}$  of  $\mathbb{R}^{p}$  so that the d-dimensional compact smooth manifold M is isometrically embedded in the subspace spanned by  $\{u_{n,l}\}_{l=1}^{K}$  for a fixed  $K \in \mathbb{N}$ . In other words, while the rank of  $\Sigma_{x}$  associated with  $\mathcal{O}(x)$  depends on x, it is bounded uniformly from above by K. Note that K in general can be much larger than d, yet it is fixed due to the well-known Nash's isometric embedding theory [27], which guarantees the existence of K that is independent of pand  $K \leq d(3d + 11)/2$ . We put the following assumption.

Assumption 5.1. Assume  $\mathcal{O}(x) = \phi(B_{\epsilon}(0))$ , where  $B_{\epsilon}(0)$  is a Euclidean ball centered at 0 with the radius  $\epsilon > 0$ , and  $\phi : \mathbb{R}^d \to M$  is diffeomorphic on  $B_{\epsilon}(0)$ . We call  $\mathcal{O}(x)$  an ellipsoid.

Note that since  $\phi$  is diffeomorphic, in general,  $\mathcal{O}(x)$  is not really an ellipsoid unless  $\phi$  is linear. But to simplify the terminology, we abuse the notation and still call it an ellipsoid. Also note that the elliptic radii of  $\mathcal{O}(x)$  is of order  $\epsilon > 0$  with the implied constant depending on the Jacobian of  $\phi$  when  $\epsilon$  is sufficiently small.

Based on the developed theorem in Sections 4.1 and 4.2, we state the following theorem that secures the recovery of MD under the manifold model in Definition 3.1. The basic idea beyond this theorem is twofold. First, it relies on the result from [24, Lemmas 3 and 6] stating that when  $\epsilon > 0$  is sufficiently small, there is a sufficiently large gap between the first *d* eigenvalues of  $\Sigma_x$  and the remaining small eigenvalues. Second, as discussed in Section 4.1, any non-trivial eigenvalue  $\ell_i$  that satisfies  $\ell_i \ll \ell_+$  is ignored by the optimal shrinkage.

#### M. GAVISH ET AL.

THEOREM 5.1. Assume Assumptions 4.1–5.1 hold. Fix  $x \in M$ . Suppose  $\sigma = \sigma(\epsilon)$  so that  $\sigma^2 \sqrt{\beta} \epsilon^{-d-2} \to 0$  and  $\sigma^2 \sqrt{\beta} \epsilon^{-d-4} \to \infty$  when  $\epsilon \to 0$ . Assume the maximal elliptic radius of  $\mathcal{O}(x)$  is  $m\epsilon$ , where  $m = m(\epsilon) \approx 1$ , and the ratio of the maximal and minimal elliptic radii is fixed for all  $\epsilon > 0$ . When  $\epsilon$  is sufficiently small, all non-trivial eigenvalues of  $\Sigma_x$  except the top *d* eigenvalues are set to zero by the optimal shrinkage  $\tilde{\eta}^*$ .

*Proof.* By [24, Lemma 6], the local covariance matrix has the following asymptotical expansion when  $\epsilon$  is sufficiently small:

$$\Sigma_{x} = \frac{|S^{d-1}|P(x)\epsilon^{d+2}}{d(d+2)} [\iota_{*}|_{x} \nabla \phi|_{\phi^{-1}(x)}] [\iota_{*}|_{x} \nabla \phi|_{\phi^{-1}(x)}]^{\top} + O(\epsilon^{d+4}),$$
(5.1)

where  $|S^{d-1}|$  is the volume of  $S^{d-1}$ . A derivation, similar to the derivation in [24, Lemma 3], yields that when  $\epsilon$  is sufficiently small, the top *d* eigenvalues of  $\Sigma_x$  are of order  $\epsilon^{d+2}$  and the remaining eigenvalues are of order equal to or higher than  $\epsilon^{d+4}$ . In other words, the 'signal strength' is of order  $\epsilon^{d+2}$ , while the noise strength is  $\sigma > 0$ . Note that there are at most *K* non-zero eigenvalues.

By (4.21), all eigenvalues smaller than  $\sigma^2 \ell_+ = \sigma^2 \sqrt{\beta}$  are eliminated by the optimal shrinkage. Combining the above, since  $\beta > 0$ , P(x) and m are all fixed, when  $\epsilon$  is sufficiently small so that  $\sigma^2 \sqrt{\beta} \epsilon^{-d-2}$  is sufficiently small and  $\sigma^2 \sqrt{\beta} \epsilon^{-d-4}$  is sufficiently large, all non-trivial eigenvalues of  $\Sigma_x$  except the top d eigenvalues are set to zero by the optimal shrinkage  $\tilde{\eta}^*$ . Indeed, by the assumption that  $\sigma^2 \sqrt{\beta} \epsilon^{-d-2} \to 0$  and  $\sigma^2 \sqrt{\beta} \epsilon^{-d-4} \to \infty$ , we know that  $c_1 \epsilon^{d+2} \ge \sigma^2 \sqrt{\beta}$  and  $c_2 \epsilon^{d+4} \le \sigma^2 \sqrt{\beta}$  when  $\epsilon$  is sufficiently small for some universal constants  $c_1$  and  $c_2$ . Therefore,  $\lambda_l \ge \sigma^2 \sqrt{\beta}$  for  $l = 1, \ldots, d$  and  $\lambda_l \le \sigma^2 \sqrt{\beta}$  for  $l = d + 1, \ldots, K$ , and hence, we conclude the proof.

We conclude this section with several remarks on this theorem. First, the condition  $\sigma^2 \sqrt{\beta} \epsilon^{-d-2} \rightarrow 0$ seems to be limited since the noise level goes to zero when  $\epsilon$  goes to zero. This condition is needed if we want to properly estimate the precision matrix locally within an ellipsoid over a manifold with elliptic radii of order  $\epsilon$ . In practice,  $\epsilon$  plays the role of 'bandwidth' which reflects the 'resolution of how accurate we could estimate the quantity we have interest. For example, in the motivating EIG example discussed in Section 3.3, we need an accurate precision matrix estimation so that the geodesic distance in the latent space can be accurately estimated. Such precision matrix via (3.10) is less affected by the curvature if  $\epsilon$ is small. However, if the noise level  $\sigma$  is fixed, then  $\epsilon$  is bounded from below, so that we have a limited accuracy when we recover the geodesic distance on the latent space, and hence the latent space itself by DM.

Second, this theorem also suggests that we *always* need the noise to attain a reasonable estimate of MD; that is  $\sigma^2 \sqrt{\beta} \epsilon^{-d-4} \rightarrow \infty$ . This statement is certainly counterintuitive, since in the linear spiked model, noise absence is desired and beneficial. To clarify this point, note that in general, the dimension of the manifold is unknown, but it is needed so that we can define a sensible MD on the manifold in Definition 3.1. See the discussion in Section 3.2. Thus, a dimension estimation is needed. This theorem mainly says that when  $\sigma > 0$  is sufficiently large so that the condition in Theorem 5.1 holds, we could get the MD in Definition 3.1 even if we do not know the dimension. If the noise is not sufficiently big and we do not know the dimension, then we cannot obtain the desired MD. If the dimension is known, or can be robustly estimated when noise exists, then this lower bound condition can be removed. The above facts stem from the complicated nonlinear interaction of nonlinear manifold structure and high dimensional noise. Finally, we mention that this result extends the EIG study in [24] under the manifold model when noise exists.



FIG. 3. (top) The performance of the optimal shrinker and the classical shrinker when d = 1. (bottom) The performance of the optimal shrinker and the classical shrinker when d = 4. The black curve represents the median of the difference between the loss of the classical shrinker and that of the theoretical optimal loss presented in log scale. The blue curve represents the median of the difference between the loss of the optimal shrinker and the theoretical optimal loss presented in log scale. The blue curve represents the median of the difference between the loss of the optimal shrinker and the theoretical optimal loss presented in log scale. The error bars depict the interquartile range of each shrinker (in log scale). The vertical blue line is  $\sigma = 1/\sqrt{\ell_+}$ , which indicates the tolerable noise level for the given  $\beta$  and signal strength.

### 6. Simulation study

To numerically compare the optimal shrinker and the classical shrinker, we set  $\beta = 0.2, 0.4, \ldots, 1$  and consider the number of samples n = 300 so that  $p = \beta n$ . For simplicity, we set  $\ell_i = i$  for  $i = 1, \ldots, d$ . We consider  $d = \{1, 4\}$ . Suppose  $x_i, i = 1, \ldots, n$  are sampled i.i.d. from the random vector  $\sum_{\ell=1}^{d} \zeta_{\ell} e_{\ell}$ , where  $e_{\ell} \in \mathbb{R}^p$  is the unit vector with  $\ell$ th entry 1,  $\zeta_{\ell} \sim N(0, 1)$  for  $\ell = 1, \ldots, d$ , and  $\zeta_{\ell}$  is independent of  $\zeta_k$  when  $\ell \neq k$ . The noisy data are simulated by  $y_i = Ax_i + \sigma^2 \xi$ , where  $\xi \sim \mathcal{N}(0, I_p)$  is the noise matrix,  $\xi$  is independent of  $\zeta_{\ell}$  and  $A \in O(p)$  is randomly sampled from O(p). In the simulation, we take  $\sigma = 0.225, 0.45, \ldots, 1.8$ . For each  $\sigma$ , we repeat the experiment 200 times and report the mean and variance of the loss  $L_p$ .

Figure 3 shows the loss of the optimal and classical shrinkers when d = 1 and d = 4 in a logarithmic scale. We observe that the loss using the classical shrinker is significantly larger. This stems from the fact that in the large p and large n regime, there are eigenvalues greater than  $\sigma^2$  that are not associated with the signal. When applying the classical shrinker (3.14) (the Moore–Penrose pseudo-inverse), these irrelevant eigenvalues contribute significantly, leading to high loss. Conversely, the optimal shrinker is much more 'selective' (as illustrated in Fig. 2), associating larger eigenvalues with the noise, thereby increasing the robustness of the estimator.

Recall that our main motivation for considering MD in the high-dimensional regime  $p \approx n$  comes from manifold learning. Next, we test the performance of the proposed OS algorithm on high-dimensional data lying on a lower dimensional manifold. Consider the following model:

$$Y = X + \sigma \xi , \qquad (6.1)$$

		<b>Error</b> $(M_n, y_1)$		<b>Error</b> $(M_n, y_2)$	
		$\eta_{\sigma}^{\mathbf{classical}}(S_n)$	$\tilde{\eta}^*(S_n)$	$\eta_{\sigma}^{\mathbf{classical}}(S_n)$	$\tilde{\eta}^*(S_n)$
	$\sigma = 1$	$18.78 \pm 1.16$	$0.78 \pm 0.54$	$55.98 \pm 2.09$	$1.32 \pm 0.93$
$\beta = 0.1$	$\sigma = 1.5$	$23.72 \pm 1.19$	$1.41\pm0.87$	$59.42 \pm 2.10$	$2.59 \pm 1.74$
	$\sigma = 2$	$33.54 \pm 1.53$	$2.18 \pm 1.31$	$64.04 \pm 1.69$	$5.19 \pm 2.99$
	$\sigma = 1$	$26.86 \pm 2.70$	$2.41 \pm 1.55$	$60.88 \pm 4.22$	$4.06 \pm 2.54$
$\beta = 0.5$	$\sigma = 1.5$	$42.66 \pm 3.43$	$4.78 \pm 2.65$	$69.06 \pm 3.43$	$10.65 \pm 5.38$
	$\sigma = 2$	$58.59 \pm 3.24$	$9.84 \pm 4.63$	$77.24 \pm 2.81$	$31.52 \pm 17.56$
	$\sigma = 1$	$34.70 \pm 4.89$	$4.05 \pm 2.35$	$64.11 \pm 5.57$	$8.39 \pm 3.91$
$\beta = 1$	$\sigma = 1.5$	$54.72 \pm 4.46$	$10.62 \pm 4.69$	$75.54 \pm 3.63$	$23.97 \pm 12.49$
	$\sigma = 2$	$69.65 \pm 3.97$	$21.35 \pm 7.94$	$83.30 \pm 2.78$	$62.99 \pm 19.34$

TABLE 1 The normalized loss of the optimal shrinker estimator  $M_n = \tilde{\eta}^*(S_n)$  and the classical estimator  $M_n = \eta_{\sigma}^{classical}(S_n)$  in the manifold setup. The mean and the standard deviation over 500 realizations are reported.

where X is sampled from a curved manifold with one chart  $\mathcal{M}$  embedded in  $\mathbb{R}^p$  that is parametrized by

$$\left[s, t, 4\left(\frac{s}{3}\right)^2 + 5\left(\frac{t}{3}\right)^2, 0, \dots, 0\right]^{\top} \in \mathbb{R}^p,$$
(6.2)

 $s,t \in [-5,5], \xi \sim \mathcal{N}(0,I_p) \in \mathbb{R}^p$  and  $\sigma > 0$  is the noise level. Note that since the precision matrix estimation and MD estimation are both local, we consider this one-chart manifold without loss of generality. Suppose the ambient dimension is fixed at p = 100. For various values of  $\beta > 0$ , we sample  $n = \lceil p/\beta \rceil$  pairs (s,t) uniformly from  $[-5,5] \times [-5,5]$  and generate *n* nonuniform points on  $\mathcal{M}$  by (6.2). Then, these points are corrupted by additive noise with various levels  $\sigma$  according to (6.1). The normalized loss of MD is computed by

$$\operatorname{Error}(M_n, y) := \frac{\left| d_{M_n}(y, X) - d_{\Sigma_X^{\dagger}}(y, X) \right|}{d_{\Sigma_X^{\dagger}}(y, X)},$$

where  $y \in \mathcal{M} \subset \mathbb{R}^p$  is an arbitrary point on ehe manifold and  $M_n$  is the estimated covariance matrix. We examine two reference points:  $y_1 = [0, ..., 0] \in \mathcal{M}$  and  $y_2 = [2, 2, 4, 0, ..., 0] \in \mathcal{M}$ . For each case, the simulation was repeated for 500 times, and the mean and standard deviation of errors are reported. In Table 1, we compare the performance of the OS estimator,  $\tilde{\eta}^*(S_n)$ , where  $S_n$  is the sample covariance, to the performance of the classical estimator  $\eta_{\sigma}^{\text{classical}}(S_n)$ . We observe that the OS outperforms the classical estimator in this well-controlled manifold setup, and we could see that the larger the noise, the worse the performance is. Moreover, the higher the dimension is, that is, the larger the  $\beta$  is, the worse the performance is.

### 7. Application to dynamical system analysis

Next, we apply our OS estimator on the multi-scale reduction problem studied in [11]. We consider the following two-dimensional SDE:

$$dx_{1}(t) = 3dt + dW_{1}(t)$$
  

$$dx_{2}(t) = -\frac{x_{2}(t)}{\epsilon}dt + \frac{1}{\sqrt{\epsilon}}dW_{2}(t),$$
(7.1)

where  $W_1$  and  $W_2$  are independent standard Bronwian motion and  $\epsilon > 0$  is a small constant quantifying the scale of  $x_2$ . This SDE defines a dynamical system with two time scales, where  $x_1$  is a slow variable and  $x_2$  is a fast variable. Suppose the state of the system  $(x_1(t), x_2(t))$  is hidden and we have access to it through an embedding in a high-dimension space via the map

$$\mathbf{y}(t) = \begin{bmatrix} y_1(t) \\ y_2(t) \\ y_3(t) \\ \vdots \\ y_p(t) \end{bmatrix} = \begin{bmatrix} f_1(x_1(t), x_2(t)) \\ f_2(x_1(t), x_2(t)) \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$
 (7.2)

In addition, suppose the high-dimensional observation  $\mathbf{y}(t)$  is contaminated by noise

$$\mathbf{z}(t) = \begin{bmatrix} z_1(t) \\ z_2(t) \\ z_3(t) \\ \vdots \\ z_p(t) \end{bmatrix} = \mathbf{y}(t) + \begin{bmatrix} w_1(t) \\ w_2(t) \\ w_3(t) \\ \vdots \\ w_p(t) \end{bmatrix},$$
(7.3)

where

$$dw_i(t) = \sigma * d\Omega_i(t),$$

 $\Omega_i(t)$  are independent standard Brownian motion, and  $\sigma$  is the noise level.

In [11], diffusion maps [5] with a Gaussian kernel based on the MD [31] were applied to effectively reduce the dimensionality of the system by attenuating the fast variable and recovering the slow variable. Specifically, the following kernel matrix  $\mathbf{W} \in \mathbb{R}^{N \times N}$  was used:

$$W_{ij} = \exp\left(-\frac{\|\mathbf{y}(t_i) - \mathbf{y}(t_j)\|_M^2}{\sigma_{kernel}^2}\right),\tag{7.4}$$

where  $\|\cdot\|_M$  is defined by

$$\|\mathbf{y}(t_i) - \mathbf{y}(t_j)\|_M^2 := \frac{1}{2} (\mathbf{y}(t_i) - \mathbf{y}(t_j))^T \left( \mathbf{C}^{\dagger}(\mathbf{y}(t_i)) + \mathbf{C}^{\dagger}(\mathbf{y}(t_j)) \right) (\mathbf{y}(t_i) - \mathbf{y}(t_j)),$$
(7.5)

where  $C(\mathbf{y}(t))$  is the local population covariance of the observed stochastic process *at the point*  $\mathbf{y}(t)$ , and  $\sigma_{kernel}$  is the selected kernel scale.

We generate N samples of the noisy observation  $\mathbf{z}(t_i)$ , where  $t_i := idt$  and i = 1, ..., N. At each sampling point  $t_i$ , we generate a trajectory consisting of q samples of noisy observations with sampling time interval  $\delta t$  starting at  $\mathbf{z}(t_i)$ , such that we have q samples:  $\mathbf{z}(t'_1), \mathbf{z}(t'_2), ..., \mathbf{z}(t'_q)$  drawn approximately from the local distribution at  $\mathbf{z}(t_i)$ . Here, we set  $\delta t = o(\epsilon)$ , where  $\epsilon > 0$  is from (7.1). Let **N** denote the local sample mean of  $\mathbf{z}(t'_1), ..., \mathbf{z}(t'_q)$ . The local sample covariance is given by

$$\hat{\mathbf{C}}(\mathbf{y}(t_i)) = \frac{1}{\delta t(q-1)} \sum_{j=1}^{q} (\mathbf{z}(t'_j) - \bar{\mathbf{z}}) (\mathbf{z}(t'_j) - \bar{\mathbf{z}})^{\top}.$$
(7.6)

For the estimation of the local precision at  $\mathbf{y}(t_i)$ , that is  $\mathbf{C}^{\dagger}(\mathbf{y}(t_i))$  from the noisy samples  $\mathbf{z}(t)$ , we test two methods. The first method is based simply on the pseudo-inverse of  $\hat{\mathbf{C}}(\mathbf{y}(t_i))$ . The second method is based on the OS estimator  $\eta_{\sigma}$  that is applied to (7.6) to estimate the local precision estimation. Once the local precision matrices are estimated, we construct the following distance matrix **D**, whose elements are given by

$$\mathbf{D}_{i,j} = \frac{1}{2} (\mathbf{z}(t_i) - \mathbf{z}(t_j))^T \left( \hat{\mathbf{C}}^{\dagger}(\mathbf{y}(t_i)) + \hat{\mathbf{C}}^{\dagger}(\mathbf{y}(t_j)) \right) (\mathbf{z}(t_i) - \mathbf{z}(t_j)).$$
(7.7)

We examine three cases of observation functions  $f_1$  and  $f_2$  **Case I.** Let

$$f_1(x_1(t), x_2(t)) = x_1(t), \quad f_2(x_1(t), x_2(t)) = x_2(t).$$
 (7.8)

and set  $\epsilon = 10^{-3}$ ,  $dt = 10^{-4}$ ,  $\delta t = 10^{-7}$ , N = 3000, p = 50, q = 50,  $\sigma = 0.1$  and  $\sigma_{kernel}$  to the 20 percent quantile of the distance matrix **D** acquired from (7.7). These parameters are taken from [11].

Case II. Let

$$f_1(x_1(t), x_2(t)) = x_1(t) + 2x_2(t), \quad f_2(x_1(t), x_2(t)) = x_2(t), \tag{7.9}$$

and set the parameters as in Case I.

Case III. Let

$$f_1(x_1(t), x_2(t)) = x_1(t) + x_2^2(t), \quad f_2(x_1(t), x_2(t)) = x_2(t).$$
 (7.10)

We set the parameters as in Case I and II, except for  $\delta t$ , which is set to a finer value  $10^{-10}$ , and  $\sigma_{kernel}$ , which is set to 5% quantile of the distance matrix **D**. These observations functions create a 'half-moon' shape, and where considered in [11].

In Fig. 4, we compare the results obtained for Case I by diffusion maps based on three distance matrices. The first distance matrix is a variant of (7.7), where the noisy signal  $\mathbf{z}(t)$  is replaced by the clean signal  $\mathbf{y}(t)$ . These results are plotted in the leftmost column. The second distance matrix is (7.7), where the estimates of the precision matrix are based on the pseudo-inverse. These results are plotted in the middle column. The third distance matrix is (7.7), where the estimates of the precision matrix are plotted in the rightmost column. In the first row, we



FIG. 4. The results obtained for Case I by diffusion maps with different distances. Left column: the distance matrix is a variant of (7.7), where the noisy signal  $\mathbf{z}(t)$  is replaced by the clean signal  $\mathbf{y}(t)$ . Middle column: the distance matrix is defined in (7.7), where the estimates of the precision matrix are based on the pseudo-inverse. Right: the third distance matrix is defined in (7.7), where the estimates of the precision matrix are based on the proposed OS. First row: the scatter plot of the slow variable  $x_1$  and the most correlated eigenvector of the transition matrix associated with diffusion maps. Second row: the scatter plot of the fast variable  $x_2$  and its most correlated eigenvector of the transition matrix associated with diffusion maps. The correlation between the respective variables and eigenvectors is shown in red above each subplot. Third row: the eigenvalues of diffusion maps, where  $d_k$  is the *k*th largest eigenvalue. The eigenvalues corresponding to the eigenvectors presented in the first and second rows (corresponding to the slow and fast variables) are marked by red and blue circles, respectively.

show a scatter plot of the slow variable  $x_1$  against its most correlated eigenvector of the transition matrix associated with the diffusion map. In the second row, we show the scatter plot of the fast variable  $x_2$ against its most correlated eigenvector of the transition matrix associated with the diffusion map. The correlation between the respective variables and eigenvectors is shown in red above each subplot. In the third row, we present the eigenvalues of the transition matrix associated with the diffusion map, where  $d_k$  is the *k*th largest eigenvalue. The eigenvalues corresponding to the eigenvectors presented in the first and second rows (corresponding to the slow and fast variables) are marked by red and blue circles, respectively. Figures 5 and 6 are the same as Fig. 4, but for Case II and Case III.

We observe that all three figures present consistent results and trends. Focusing first on the leftmost column, we see that the use of MD in diffusion maps based on the clean signal recovers accurately the



FIG. 5. The results obtained for Case II by diffusion maps with different distances. The meaning of the plot is the same as that detailed in the caption of Fig. 4.

slow variable  $x_1$  and attenuates the fast variable  $x_2$ , conforming to the results presented in [11] for only a two-dimensional observations. In the middle column, we see that the addition of noise hinders both the recovery of the slow variable as well as the attenuation of the fast variable, when the pseudo-inverse is used for the estimation of the MD. In contrast, in the rightmost column, we see that when the proposed OS is used, the obtained results based on the noisy signal are comparable to the results obtained based on the clean signal in the leftmost column.

These results demonstrate, in the context of manifold learning, that the estimation of the MD based on the pseudo-inverse fails in the high-dimensional regime with the presence of noise. In addition, they show that using the proposed optimal shrinker indeed offers a remedy and gives rise to accurate manifold learning.

### 8. Conclusions

We proposed a new estimator for MD based on precision matrix shrinkage. For an appropriate loss function, we show that the proposed estimator is asymptotically optimal and outperforms the classical implementation of MD using the Moore–Penrose pseudo-inverse of the sample covariance. Importantly,



FIG. 6. The results obtained for Case III by diffusion maps with different distances. The meaning of the plot is the same as that detailed in the caption of Fig. 4.

the proposed estimator is particularly beneficial when the data are noisy and in high-dimension, a case in which the classical MD estimator might completely fail. Consequently, we believe that the new estimator may be useful in modern data analysis applications, involving for example, local principal component analysis, metric design and manifold learning.

In this work, we focused on the case in which the intrinsic dimensionality of the data (the rank of the covariance matrix) d is unknown, and therefore, it was not explicitly used in the estimation. Yet, in many scenarios, this dimension is known. In this case, it could be beneficial to consider a direct truncation and use only the top d eigen-pairs for the estimation of the precision matrix. While the benefit from such a truncation has been shown empirically in several applications [22, 41, 46, 48], it still requires a systematic investigation. For example, identifying the rank of the signal, or estimating the dimension of a manifold, are by themselves highly challenging tasks [16, 19]. Note that in the particular manifold setup, knowing the manifold dimension is essentially different from the *rank-aware shrinker* discussed in [8]; as we showed here, under the manifold setup, the rank of the covariance matrix associated with points residing inside a small neighborhood of any point on the manifold could be much larger than d. Since the focus of this paper is MD recovery, the associated loss function for the OS of the precision matrix is the operator norm. As is discussed in [8], there are other loss functions that

## M. GAVISH ET AL.

we can choose, such as the Frobenius norm, the nuclear norm, etc. It would be interesting to explore if the OS for the precision matrix under those norms when they are needed. We shall also mention that the manifold model discussed in this paper is a simplified model for more complicated datasets with more nonlinear structure. While it is possible to apply the principle component analysis approach to denoise the data when the smooth compact manifold assumption holds thanks to the fixed *K* mentioned in Section 5, in general, this approach might not be optimal, particularly when the geometric structure is more complicated. The current work explores this situation with a simplified manifold model and paves a way toward more complicated setups, and these cases will be explored in our future work.

## Data availability statement

No new data were generated or analyzed in support of this review.

# Acknowledgements

We thank the anonymous reviewers for their constructive comments and critiques that highly improve the original manuscript.

# Funding

H-CSRC Security Research Center and Israeli Science Foundation (1523/16 to M.G.); Tel-Aviv University ICRC Research Center (to M.G. and R.T.); Technion Hiroshi Fujiwara cyber security research center (to R.T.); Pazy Foundation (to R.T.).

# **Conflict of Interest**

No conflict of interest.

## References

- 1. ALAGAPAN, S., SHIN, H. W., FRÖHLICH, F. & WU, H.-T. (2019) Diffusion geometry approach to efficiently remove electrical stimulation artifacts in intracranial electroencephalography. *J. Neural Eng.*, **16**, 036010.
- BAIK, J. & SILVERSTEIN, J. W. (2006) Eigenvalues of large sample covariance matrices of spiked population models. J. Multivar. Anal., 97, 1382–1408.
- 3. BAIK, J., BEN AROUS, G. & PECHE, S. (2005) Phase transition of the largest eigenvalue for non-null complex sample covariance matrices. *Ann. Probab.*, **33**, 1643–1697.
- 4. BELKIN, M. & NIYOGI, P. (2003) Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.*, **15**, 1373–1396.
- 5. COIFMAN, R. R. & LAFON, S. (2006) Diffusion maps. Appl. Comput. Harmon. Anal., 21, 5-30.
- 6. DAI, D. & HOLGERSSON, T. (2018) High-dimensional CLTs for individual mahalanobis distances. *Trends and Perspectives in Linear Statistical Inference*. (M. TEZ & D. von ROSEN eds)., pp. 57–68. Springer, Cham.
- DONOHO, D. L. & GRIMES, C. (2003) Hessian eigenmaps: locally linear embedding techniques for highdimensional data. *Proc. Natl. Acad. Sci. U. S. A.*, 100, 5591–5596.
- DONOHO, D. L., GAVISH, M. & JOHNSTONE, I. M. (2018) Optimal Shrinkage of Eigenvalues in the Spiked Covariance Model. Ann. Stat., 46, 1742–1778.
- 9. DOV, D., TALMON, R., COHEN, I., DOV, D., TALMON, R. & COHEN, I. (2016) Kernel method for voice activity detection in the presence of transients. *IEEE/ACM Trans. Audio, Speech, Language Process.*, 24, 2313–2326.
- 10. DSILVA, C. J., TALMON, R., RABIN, N., COIFMAN, R. R. & KEVREKIDIS, I. G. (2013) Nonlinear intrinsic variables and state reconstruction in multiscale simulations. *J. Chem. Phys.*, **139**, 11B608\_1.

- 11. DSILVA, C. J., TALMON, R., GEAR, C. W., COIFMAN, R. R. & KEVREKIDIS, I. G. (2016) Data-driven reduction for a class of multiscale fast-slow stochastic dynamical systems. *SIAM J. Appl. Dyn. Syst.*, **15**, 1327–1351.
- 12. EL KAROUI, N. (2010) On information plus noise kernel random matrices. Ann. Stat., 38, 3191–3216.
- 13. EL KAROUI, N. & WU, H.-T. (2016) Connection graph Laplacian methods can be made robust to noise. *Ann. Stat.*, **44**, 346–372.
- 14. JOHNSTONE, I. M. (2001) On the distribution of the largest eigenvalue in principal components analysis. *Ann. Stat.*, **29**, 295–327. https://doi.org/10.2307/2674106.
- 15. JOHNSTONE, I. M. (2006) High Dimensional Statistical Inference and Random Matrices. *Proceedings of International Congress of Mathematicians*, EMS, pp. 307–333.
- KRITCHMAN, S. & NADLER, B. (Oct. 2009) Non-parametric detection of the number of signals: Hypothesis testing and random matrix theory. *Trans. Sig. Proc.*, 57, 3930–3941 ISSN 1053-587X.
- LAHAV, A., TALMON, R. & KLUGER, Y. (2018) Mahalanobis distance informed by clustering. *Inf. Inference*, 8, 377–406.
- LEDERMAN, R. R. & TALMON, R. (2018) Learning the geometry of common latent variables using alternatingdiffusion. *Appl. Comput. Harmon. Anal.*, 44, 509–536.
- 19. LEVINA, E. & BICKEL, P. J. (2005) Maximum likelihood estimation of intrinsic dimension. *Advances in neural information processing systems*, NIPS, pp. 777–784.
- 20. LIN, Y.-T., MALIK, J. & WU, H.-T. (2021) Wave-shape oscillatory model for nonstationary periodic time series analysis. *Foundations of Data Science*, **3**, 99–131.
- 21. LINDENBAUM, O., YEREDOR, A., SALHOV, M. & AVERBUCH, A. (2020) Multi-view diffusion maps. *Inf. Fusion*, **55**, 127–149.
- 22. LIU, G.-R., LO, Y.-L., MALIK, J., SHEU, Y.-C. & WU, H.-T. (2020) Diffuse to fuse eeg spectra–intrinsic geometry of sleep dynamics for classification. *Biomed. Signal Process. Control*, **55**, 101576.
- MAHALANOBIS, P. C. (1936) On the generalized distance in statistics. *Proc. Natl. Acad. Sci. U. S. A. (Calcutta)*, 2, 49–55.
- 24. MALIK, J., SHEN, C., WU, H.-T. & WU, N. (2019) Connecting dots from local covariance to empirical intrinsic geometry and locally linear embedding. *Pure Appl. Anal.*, **1**, 515–542.
- 25. MARCENKO, V. A. & PASTUR, L. A. (1967) Distribution of eigenvalues for some sets of random matrices. *Math. USSR-Sb.*, **1**, 457.
- 26. MCLACHLAN, G. J. (1999) Mahalanobis distance. Resonance, 4, 20–26.
- 27. NASH, J. (1956) The imbedding problem for riemannian manifolds. Ann. Math., (2), 20-63.
- PAUL, D. (2007) Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Stat. Sin.*, **17**, 1617–1642.
- ROWEIS, S. T. & SAUL, L. K. (2000) Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290, 2323–2326. https://doi.org/10.1126/science.290.5500.2323.
- SINGER, A. (2018) Mathematics for cryo-electron microscopy. Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018. World Scientific, pp. 3995–4014.
- SINGER, A. & COIFMAN, R. R. (2008) Non-linear independent component analysis with diffusion maps. *Appl. Comput. Harmon. Anal.*, 25, 226–239.
- 32. SINGER, A. & WU, H.-T. (2012) Two-dimensional tomography from noisy projections taken at unknown random directions. *SIAM J. Imaging Sci.*, **6**, 136–175.
- SINGER, A. & WU, H.-T. (2012) Vector diffusion maps and the connection Laplacian. *Commun. Pure Appl. Math.*, 65, 1067–1144.
- 34. SINGER, A., SHKOLNISKY, Y. & NADLER, B. (2009) Diffusion interpretation of nonlocal neighborhood filters for signal denoising. *SIAM J. Imaging Sci.*, **2**, 118–139.
- 35. STEIN, C. M. (1986) Lectures on the theory of estimation of many parameters. J. Sov. Math., 34, 1373–1403.
- 36. SU, P.-C., MILLER, S., IDRISS, S., BARKER, P. & WU, H.-T. (2019) Recovery of the fetal electrocardiogram for morphological analysis from two trans-abdominal channels via optimal shrinkage. *Physiol. Meas.*, **40**, 115005.

#### M. GAVISH ET AL.

- 37. TALMON, R. & COIFMAN, R. R. (2013) Empirical intrinsic geometry for nonlinear modeling and time series filtering. *Proc. Natl. Acad. Sci.*, **110**, 12535–12540.
- TALMON, R. & COIFMAN, R. R. (2015) Intrinsic modeling of stochastic dynamical systems using empirical geometry. *Appl. Comput. Harmon. Anal.*, 39, 138–160.
- 39. TALMON, R. & WU, H.-T. (2019) Latent common manifold learning with alternating diffusion: analysis and applications. *Appl. Comput. Harmon. Anal.*, **47**, 848–892.
- 40. TALMON, R., KUSHNIR, D., COIFMAN, R. R., COHEN, I. & GANNOT, S. (2011) Parametrization of linear systems using diffusion kernels. *IEEE Trans. Signal Process.*, **60**, 1159–1173.
- 41. TALMON, R., MALLAT, S., ZAVERI, H. & COIFMAN, R. R. (2015) Manifold learning for latent variable inference in dynamical systems. *IEEE Trans. Signal Process.*, **63**, 3843–3856.
- TENENBAUM, J. B., DE SILVA, V. & LANGFORD, J. C. (2000) A global geometric framework for nonlinear dimensionality reduction. *Science*, **290**, 2319–2323.
- 43. WANG, S.-C., WU, H.-T., HUANG, P.-H., CHANG, C.-H., TING, C.-K. & LIN, Y.-T. (2020) Novel imaging revealing inner dynamics for cardiovascular waveform analysis via unsupervised manifold learning. *Anesth. Analg.*, **130**, 1244–1254.
- 44. WEINBERGER, K. Q. & SAUL, L. K. (2009) Distance metric learning for large margin nearest neighbor classification. J. Mach. Learn. Res., 10, 207–244.
- 45. WU, H.-T. & WU, N. (2018) Think globally, fit locally under the Manifold Setup: Asymptotic Analysis of Locally Linear Embedding. *Ann. Stat.*, **46**, 3805–3837.
- 46. WU, H.-T., TALMON, R. & LO, Y.-L. (2015) Assess sleep stage by modern signal processing techniques. *IEEE Trans. Biomed. Engineering*, **62**, 1159–1168.
- 47. XIANG, S., NIE, F. & ZHANG, C. (2008) Learning a Mahalanobis distance metric for data clustering and classification. *Pattern Recogn.*, **41**, 3600–3612.
- 48. YAIR, O., TALMON, R., COIFMAN, R. R. & KEVREKIDIS, I. G. (2017) Reconstruction of normal forms by learning informed observation geometries from data. *Proc. Natl. Acad. Sci.*, **114**, E7865–E7874.
- 49. YANG, L. & JIN, R. (2006) Distance metric learning: a comprehensive survey. Michigan State University, 2, 4.