
Log-Euclidean Signatures for Intrinsic Distances Between Unaligned Datasets

Tal Shnitzer¹ Mikhail Yurochkin^{2,3} Kristjan Greenewald^{2,3} Justin Solomon¹

Abstract

The need for efficiently comparing and representing datasets with unknown alignment spans various fields, from model analysis and comparison in machine learning to trend discovery in collections of medical datasets. We use manifold learning to compare the intrinsic geometric structures of different datasets by comparing their diffusion operators, symmetric positive-definite (SPD) matrices that relate to approximations of the continuous Laplace-Beltrami operator from discrete samples. Existing methods typically assume known data alignment and compare such operators in a pointwise manner. Instead, we exploit the Riemannian geometry of SPD matrices to compare these operators and define a new theoretically-motivated distance based on a lower bound of the log-Euclidean metric. Our framework facilitates comparison of data manifolds expressed in datasets with different sizes, numbers of features, and measurement modalities. Our log-Euclidean signature (LES) distance recovers meaningful structural differences, outperforming competing methods in various application domains.

1. Introduction

With the increased availability of data-driven methods, the ability to compare data collections has become crucial. In many cases, such comparisons are impossible without computationally-expensive alignment and manipulation, especially when the datasets are of different dimensions or are embedded in different domains. This problem is prominent in learning, where the need for comparing datasets arises, e.g., in the evaluation and comparison of neural network (NN) architectures (Kornblith et al., 2019), data manifolds of generative models (Heusel et al., 2017), and tasks in transfer and meta learning (Achille et al., 2019).

¹MIT CSAIL ²IBM Research ³MIT-IBM Watson AI Lab. Correspondence to: Tal Shnitzer <talsd@mit.edu>.

A common approach to dataset comparison assumes that complex high-dimensional data lie on some manifold of lower dimension, i.e., the manifold hypothesis (Fefferman et al., 2016). Manifold learning techniques exploit this assumption to provide geometrically-motivated intrinsic data representations, commonly by constructing discrete approximations of the underlying manifold’s Laplace-Beltrami operator or heat kernel (Belkin & Niyogi, 2003; Coifman & Lafon, 2006). Comparing such representations usually requires point-to-point correspondences, yielding an NP-hard quadratic assignment problem (Cela, 2013).

One approach to avoiding the need for data alignment is via spectral properties of the operators. Tsitsulin et al. (2020) compare unaligned datasets by computing distances between approximations of the heat kernel traces at different temperature scales, lower bounding an optimal transport-based metric computing a pointwise distance between heat kernels (Mémoli, 2011). This pointwise comparison, providing the theoretical motivation for their distance as in many related methods, overlooks the symmetric positive-definite (SPD) structure of the heat kernel approximation used in practice. Prior studies show that using appropriate metrics on the space of SPD matrices, forming a differentiable Riemannian manifold, significantly outperforms pointwise Euclidean comparisons (Pennec et al., 2006; Arsigny et al., 2006; Barachant et al., 2013).

In this work, we take into account the SPD structure of heat kernel approximations and use the log-Euclidean (LE) metric (Arsigny et al., 2006) to compare them. We propose the *log-Euclidean signature (LES) distance* for unaligned datasets based on a lower bound of this metric. Our distance function can be formulated in terms of efficient descriptors, which capture the underlying intrinsic geometry of the data manifold, facilitating the comparison of heterogeneous datasets. We demonstrate properties of the LES distance on various data types, including 3-4D point clouds of geometric shapes, single cell RNA-sequencing data, and NN embeddings, where we show that the LES distance provides useful insights on architectures and embedding structures.

2. Related Work

Many frameworks for comparing datasets have been proposed. We outline the most related approaches proposing

geometrically-motivated measures.

Geometrically-motivated dataset descriptors. An early work proposing a geometric global fingerprint of a dataset is ‘Shape-DNA’ (Reuter et al., 2006), constructed from approximate Laplace-Beltrami eigenvalues. This signature was formulated for $2D$ and $3D$ manifolds, and its extension to high-dimensional data is nontrivial. More recent works include the Geometry Score (GS) (Khrulkov & Oseledets, 2018), the Network Laplacian Spectral Descriptor (NetLSD), and the Intrinsic Multi-scale Distance (IMD) (Tsitsulin et al., 2018; 2020). The GS characterizes datasets by their topological properties using persistent homology, which captures multi-scale structural differences of the data manifolds but comes with high computational costs. NetLSD and IMD propose dataset signatures based on the heat kernel trace at different temperature scales, for graphs (NetLSD) and for general datasets (IMD). IMD is the most related work to ours, discussed in Subsection 4.3 and Section 5.

Distances between datasets. Methods for computing distances between datasets often estimate some alignment between them. For example, Peyré et al. (2016) introduce a technique for computing the Gromov-Wasserstein (GW) distance between unaligned kernels. Such methods typically lead to non-convex optimization problems with extremely high computational costs. Distance measures that do not rely on data alignment include the Maximum Mean Discrepancy (MMD) (Gretton et al., 2012), Manifold Topology Divergence (MTD) (Barannikov et al., 2021) and the Cross Local Intrinsic Dimensionality (CrossLID) (Barua et al., 2019). MTD measures topological differences of datasets by constructing simplicial complexes between them and evaluating the obtained topology. CrossLID compares data distributions in small neighborhoods of the two datasets. The MMD, MTD and CrossLID require that datasets are embedded in the same space with the same dimensionality.

Finally, the Fréchet Inception Score (FID) (Heusel et al., 2017) computes the Wasserstein-2 distance for Gaussians between outputs of an Inception convolutional NN layer (Szegedy et al., 2015). Due to the Inception layer, the FID score and its many variants are not suited for the evaluation of general datasets, which are the focus of this paper.

For the pairwise measures above, every coupling must be evaluated separately, yielding high computational costs when comparing many datasets. In contrast, our framework, while derived from a pairwise distance, first embeds the datasets, allowing for simple Euclidean comparisons.

3. Preliminaries

We briefly describe relevant background for the proposed method, first describing a method for approximating the Laplace-Beltrami operator and heat kernel from high-

dimensional data, with minimal prior assumptions on the data type. We then outline the LE metric, which defines a Riemannian manifold on the space of SPD matrices.

3.1. Diffusion Operator Approximations from Data

Manifold learning techniques typically construct positive-semidefinite kernel matrices from sampled data. We focus here on diffusion maps, a manifold learning technique that defines the diffusion operator, which was shown to have desirable convergence properties (Coifman & Lafon, 2006).

Given a set of points $x_i \in \mathbb{R}^d$, $i = 1, \dots, N$, assumed to lie on some low-dimensional manifold \mathcal{M} embedded in the ambient space \mathbb{R}^d , an affinity matrix is constructed by:

$$K_{i,j} = \exp\left(-\frac{d^2(x_i, x_j)}{\sigma^2}\right), \quad (1)$$

where $d(\cdot, \cdot)$ denotes some notion of distance suitable to the data (e.g., Euclidean distance) and σ is the kernel scale, typically the median distance between data points multiplied by some constant (Ding & Wu, 2020). The affinity matrix is then normalized, producing the diffusion operator \mathbf{W}_{DM} :

$$\tilde{\mathbf{W}} = \tilde{\mathbf{D}}^{-1} \mathbf{K} \tilde{\mathbf{D}}^{-1}, \quad \tilde{D}_{i,i} = \sum_{j=1}^N K_{i,j}, \quad (2)$$

$$\mathbf{W}_{DM} = \mathbf{D}^{-1} \tilde{\mathbf{W}}, \quad D_{i,i} = \sum_{j=1}^N \tilde{W}_{i,j}, \quad (3)$$

where the normalization in (2) mitigates the effect of non-uniform data sampling (Coifman & Lafon, 2006).

The diffusion operator satisfies several convergence properties. Specifically, $\mathbf{W}_{DM}^{t/\sigma^2}$ converges pointwise to the Neumann heat kernel of the underlying manifold as $N \rightarrow \infty$ and $\sigma^2 \rightarrow 0$, and $\mathbf{L} = (\mathbf{I} - \mathbf{W}_{DM})/\sigma^2$ converges to the Laplace-Beltrami operator (Coifman & Lafon, 2006). Both continuous operators capture geometric properties of the manifold. In addition, eigenvalues and eigenvectors of \mathbf{L} converge to Laplace-Beltrami eigenvalues and eigenfunctions (Belkin & Niyogi, 2007; Dunson et al., 2021).

The diffusion operator has proven to be extremely useful in numerous tasks and diverse applications, providing data-driven low-dimensional representations that preserve the structure of the underlying manifold, see e.g. (Talmon et al., 2013; Mishne & Cohen, 2012; Talmon & Coifman, 2013).

Following previous work (Coifman & Hirn, 2014; Katz et al., 2020; Shnitzer et al., 2022), we define an SPD matrix that is similar to \mathbf{W}_{DM} given by:

$$\mathbf{W} = \mathbf{D}^{1/2} \mathbf{W}_{DM} \mathbf{D}^{-1/2}. \quad (4)$$

Since \mathbf{W} and \mathbf{W}_{DM} are similar, they share eigenvalues, and their eigenvectors are related by $\mathbf{D}^{1/2}$. Our proposed distance function relies on the eigenvalues of \mathbf{W} . The convergence properties of \mathbf{W}_{DM} and its similarity to \mathbf{W} thus imply that our distance compares (a transformation of) discrete approximations of the heat kernel spectrum.

3.2. Riemannian Manifold of SPD Matrices

The space of $N \times N$ SPD matrices forms a Riemannian manifold when endowed with an appropriate metric. The most common metrics are the affine-invariant (AI) (Pennec et al., 2006) and LE metrics (Arsigny et al., 2006), both providing closed-form geodesic distances. We focus on the LE metric, since it shares useful properties with the AI metric—i.e., invariance to inversion, translation, scaling and orthogonal transformations—while requiring lower computational cost.

Let Sym_N^+ denote the set of $N \times N$ real SPD matrices, and let Sym_N denote the set of $N \times N$ real symmetric matrices. Define the logarithmic product for $\mathbf{W}_1, \mathbf{W}_2 \in \text{Sym}_N^+$ as $\mathbf{W}_1 \odot \mathbf{W}_2 = \exp(\log \mathbf{W}_1 + \log \mathbf{W}_2)$, where \exp and \log denote the matrix exponent and logarithm functions. This defines a Lie group structure on Sym_N^+ . By commutativity of the logarithmic product, the Lie group can be equipped with a bi-invariant Riemannian metric induced by any inner product on $\text{Sym}_N = \mathcal{T}_1 \text{Sym}_N^+$, the tangent space at the identity (Lie algebra) (Arsigny et al., 2007); the associated geodesic distance is then $d(\mathbf{W}_1, \mathbf{W}_2) = \|\log \mathbf{W}_1 - \log \mathbf{W}_2\|$, where the norm is induced by the inner product. The L^2 inner product (and corresponding Frobenius norm) yields the following LE distance:

$$d_{\text{LE}}(\mathbf{W}_1, \mathbf{W}_2)^2 = \|\log \mathbf{W}_1 - \log \mathbf{W}_2\|_F^2. \quad (5)$$

This metric is invariant to similarity transformations (scaling and orthogonal transformations).

From the SPD manifold perspective, the matrix log in the LE distance maps SPD matrices onto the tangent space at the identity, i.e., the space of symmetric matrices, which is closed under Euclidean operations. Therefore, the distances between SPD matrices are computed as the Euclidean distance between their mappings at the identity, reducing the space of SPD matrices to a flat space (Arsigny et al., 2007).

The LE metric improves performance in various applications involving SPD matrices, such as diffusion tensor imaging (Arsigny et al., 2006), image set classification (Huang et al., 2015), and land cover classification (Guan et al., 2019). Computing LE distances, however, requires pointwise correspondence of the two SPD matrices; such data alignment may be expensive or impossible to obtain. One of the key contributions of our work is to alleviate the need for data alignment via a lower bound on the LE distance. Therefore, the choice of the LE metric is intentional: it is simpler relatively to other SPD metrics, while performing well in practice. Overcoming the need for alignment in other metrics may be more challenging, e.g. due to matrix multiplication/inversion in the AI metric (Pennec et al., 2006). In Appendix B.3 we empirically compare the LE metric, the AI metric, and Euclidean distance to motivate the use of the SPD Riemannian geometry and specifically the LE metric.

4. Distance Between Intrinsic Representations of Unaligned Data

Based on the tools in Section 3, we propose a distance measure for comparing the underlying manifold geometry of unaligned datasets. Our method conceptually consists of two steps. First, we represent each dataset using an SPD diffusion operator capturing underlying geometry of the data, constructed in Subsection 3.1. Second, we account for the Riemannian structure of the space of SPD matrices by defining our new theoretically motivated distance measure for unaligned datasets as a lower bound of the LE metric. The resulting distance is based on the ordered eigenvalues of the matrices representing the data. We refer to this distance as the *LES (log-Euclidean signature) distance*. This procedure embeds whole datasets into a new space based on their underlying geometry, allowing us to compare datasets of *different sizes*, with *different numbers of features*, and from *various measurement modalities*.

Below, we lay out the details and derivations leading to the LES distance and describe implementation considerations for large datasets. Algorithm 2 summarizes our approach.

4.1. Bounding the Log-Euclidean Distance

We begin with a setting involving datasets of the same size, but relax this assumption in Subsection 4.2.2. Let $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{N \times N}$ denote SPD diffusion operators, constructed according to (4) from two datasets. We aim to compare these matrices in a manner that acknowledges their SPD structure using the LE distance (5). Such a comparison, however, requires point-to-point correspondences.

Existing methods for aligning and comparing datasets typically come with extremely high computational costs. Therefore, we propose a different approach for comparing the matrix representations of the datasets through a lower bound of the LE distance. The lower bound is described in the following proposition, along with an upper bound.

Proposition 1. *Given two SPD matrices, $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{N \times N}$, the following holds:*

$$\begin{aligned} \|\log \mathbf{W}_1 - \log \mathbf{W}_2\|_F^2 &\geq \sum_{i=1}^N \left(\log \lambda_i^{(1)} - \log \lambda_i^{(2)} \right)^2 \quad (6) \\ \|\log \mathbf{W}_1 - \log \mathbf{W}_2\|_F^2 &\leq \sum_{i=1}^N \left(\log \lambda_i^{(1)} - \log \lambda_{N-i+1}^{(2)} \right)^2 \end{aligned}$$

where $\lambda_i^{(\ell)}$ denotes the i 'th eigenvalue of \mathbf{W}_ℓ and the eigenvalues of both matrices are organized in decreasing order, i.e., $\lambda_1^{(\ell)} \geq \lambda_2^{(\ell)} \geq \dots \geq \lambda_N^{(\ell)}$.

Proof. The proof follows directly from Theorem III.4.4 in Bhatia (2013), which states that for any pair of Hermitian

matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{N \times N}$ and any symmetric gauge function Φ on \mathbb{R}^N , we have $\Phi(\lambda^{(\mathbf{A})\downarrow} - \lambda^{(\mathbf{B})\uparrow}) \geq \Phi(\lambda^{(\mathbf{A}-\mathbf{B})}) \geq \Phi(\lambda^{(\mathbf{A})\downarrow} - \lambda^{(\mathbf{B})\downarrow})$, where $\lambda^{(\mathbf{A})\downarrow}$ denotes the eigenvalues of \mathbf{A} organized in decreasing order and $\lambda^{(\mathbf{B})\uparrow}$ denotes the eigenvalues of \mathbf{B} organized in increasing order. This theorem holds due to the symmetry of $\log \mathbf{W}_\ell$ and since the l^2 norm is a symmetric gauge function. \square

Equality for the lower bound is obtained when the matrices are mutually diagonalizable and have an *identical* order of basis elements w.r.t. the ordered eigenvalues. Equality for the upper bound is obtained when the matrices are mutually diagonalizable and have an *opposite* order of basis elements.

Both bounds in Proposition 1 are easily computable from the matrices \mathbf{W}_ℓ . However, due to the reversed order of the eigenvalues, the upper bound is loose and is not suitable in practice due to numerical limitations on the estimation of small eigenvalues. We construct LES from the *lower* bound because the eigenvalues of the matrices are ordered consistently in the sum. In particular, the right-hand side of (6) is a squared Euclidean distance between vectors whose elements are $\log \lambda_i^{(\ell)}$. This choice is further motivated by recent distances based on lower bounds of other common metrics (Tsitsulin et al., 2020; Barannikov et al., 2021).

4.2. Accommodating Large Datasets of Different Sizes

We use the lower bound from (6) to define our proposed distance, presented later in Subsection 4.3, Eq. (11). However, a few computational considerations are required in its design to compare large datasets with different numbers of samples. First, we add kernel regularization to \mathbf{W} and provide theoretical reasoning in Subsection 4.2.1. Second, we address datasets of different sizes by truncating the full spectrum to a fixed length K . Third, we describe methods for estimating the spectrum and the matrix \mathbf{W} in (4) to facilitate the use of the LES distance in large datasets. The latter two points are discussed in Subsection 4.2.2.

4.2.1. OPERATOR REGULARIZATION

As the number of data points increases, the diffusion operator in (3) converges to the heat kernel. The matrix logarithm, however, is bounded only for finite matrices, since as $N \rightarrow \infty$, the eigenvalues decay and $\lim_{k \rightarrow \infty} \log \lambda_k^{(\ell)} = -\infty$ (Ha Quang et al., 2014). This is especially prominent due to numerical instabilities and errors, causing zero or even negative eigenvalues in large matrices that are expected to be SPD. Common practice adds kernel regularization of the form $\mathbf{W}_\ell + \gamma \mathbf{I}$, where $0 < \gamma \in \mathbb{R}$ and \mathbf{I} denotes the identity.

Regularizing the operators serves two purposes. First, it allows us to compute the matrix logarithm for matrices whose eigenvalues decrease below the minimal numerical error

limit. Second, it reduces the effect of noise, since noise commonly governs the smaller eigenvalues (for relatively high signal-to-noise ratios); in particular, this regularization reduces the effect of eigenvalues smaller than γ in $\log(\lambda^{(\ell)} + \gamma)$. Hence, this regularization can be thought of as truncating the log-spectrum at a specific precision.

This regularization leads to theoretical complications, however, since the infinite-dimensional identity operator is unbounded in the Hilbert-Schmidt norm, which extends the Frobenius norm to infinite dimensions. Ha Quang et al. (2014; 2016) tackle this issue by generalizing the LE metric to an infinite-dimensional setting via positive definite unitized Hilbert-Schmidt operators. They define the Log-Hilbert-Schmidt metric for regularized positive operators, which is bounded for both finite and infinite operators. In infinite dimensions, they modify the LE metric as follows:

$$d_{\log\text{HS}}^2(\mathbf{W}_1 + \gamma \mathbf{I}, \mathbf{W}_2 + \mu \mathbf{I}) = (\log \gamma - \log \mu)^2 + \|\log(\mathbf{W}_1/\gamma + \mathbf{I}) - \log(\mathbf{W}_2/\mu + \mathbf{I})\|_{\text{HS}}^2 \quad (7)$$

where $\|\cdot\|_{\text{HS}}^2$ denotes the Hilbert-Schmidt norm, $0 < \gamma, \mu \in \mathbb{R}$, and $\mathbf{W}_1 + \gamma \mathbf{I}$ and $\mathbf{W}_2 + \mu \mathbf{I}$ are regularized positive definite operators. When $\gamma = \mu$ and in the finite case, this formulation coincides with the standard LE metric applied to regularized operators. In this work, we follow Ha Quang et al. (2016) and set $\gamma = \mu$. This derivation motivates the use of the regularized diffusion operators in the LE metric, guaranteeing that the defined distance will not diverge as the number of samples, N , increases.

Remark. Another way to address the numerical limitations, i.e. eigenvalues that decrease below the minimal numerical error limit, is with a fixed-rank framework for symmetric positive semidefinite matrices. In fact, such a framework can give additional justification for our distance. Our *exact* distance can also be obtained by lower-bounding a fixed-rank framework in the spirit of Bonnabel & Sepulchre (2010); Collard et al. (2014), by zeroing the term related to the rotation group (see Eq (49) and Alg. 3 in Collard et al. (2014)), which requires the alignment.

Remark. The choice of the regularization parameter γ can be motivated by expected noise levels in the data. For very noisy data, the spectrum components that represent noise are expected to correspond to eigenvalues with larger magnitudes, since the noise accounts for more variability (Ding & Wu, 2020). In such scenarios, it is advised to take larger γ values, which can be motivated by the spectral decay. For example, as implied by Ding & Wu (2020), the ratio between consecutive eigenvalues, i.e., λ_i/λ_{i+1} , can distinguish between parts of the spectrum that relate to the data and parts that relate to the noise. In our experiments, we set $\gamma \in [10^{-8}, 10^{-5}]$, which empirically led to good results.

4.2.2. SPECTRUM TRUNCATION AND ESTIMATION

The formulation in (6) employs the entire spectrum of the operator. However, since we aim to compare datasets of different sizes, we truncate the spectrum and use only the leading (largest) eigenvalues. This choice is motivated by the regularization above, which reduces the effects of smaller eigenvalues, and by the fact that smaller eigenvalues are commonly governed by noise. Therefore, truncating the spectrum typically will not affect the ability to capture geometric structures in the data and is common in similar scenarios (Reuter et al., 2006; Cosmo et al., 2019).

Remark. Truncating the spectrum for comparing datasets of different sizes is significantly more advantageous than modifying the size of the dataset, e.g. by removing or repeating samples. With direct data manipulations, geometric structures might be distorted or ignored, whereas with spectrum truncation, the dominant geometric structures are preserved, since they are typically captured by the largest eigenvalues.

Remark. Spectrum truncation preserves the lower bound of the LE metric in (6) since it removes nonnegative terms.

Spectrum estimation. Computing even the truncated spectrum for large matrices is computationally demanding. We use a recent spectral approximation technique based on the Nyström method for low-rank positive semidefinite (PSD) matrix approximation (Tropp et al., 2017). Their approach implements a more stable truncation of the Nyström approximation to rank $K < N$ based on Li et al. (2017).

We outline the relevant parts of their method in Algorithm 1. This algorithm approximates the leading K eigenvalues of a PSD matrix \mathbf{W} , with theoretical guarantees. For the eigenvalue estimation in our setting, we obtain the following theoretical bounds based on a result by Tropp et al. (2017):

Proposition 2. *The rank- K approximation of a real symmetric PSD matrix, \mathbf{W} , computed according to Algorithm 1, and denoted by $\hat{\mathbf{W}}_K$, satisfies the following error bounds.*

1. *The eigenvalue approximation error is bounded by:*

$$\mathbb{E} \left[\sum_{i=1}^K \left| \lambda_i^{(\mathbf{W})} - \lambda_i^{(\hat{\mathbf{W}}_K)} \right| \right] \leq \frac{K}{M-K-1} \sum_{i=K+1}^N \lambda_i^{(\mathbf{W})} \quad (8)$$

2. *The γ -regularized log-eigenvalue error is bounded by:*

$$\begin{aligned} & \mathbb{E} \left[\sum_{i=1}^K \left| \log(\lambda_i^{(\mathbf{W})} + \gamma) - \log(\lambda_i^{(\hat{\mathbf{W}}_K)} + \gamma) \right| \right] \\ & \leq \frac{1.5K}{(M-K-1)(\lambda_K^{(\mathbf{W})} + \gamma)} \sum_{i=K+1}^N \lambda_i^{(\mathbf{W})} \quad (9) \end{aligned}$$

for $\frac{|\lambda_i^{(\mathbf{W})} - \lambda_i^{(\hat{\mathbf{W}}_K)}|}{\lambda_i^{(\mathbf{W})} + \gamma} \leq 0.5828$, where the expectations are over the normal distribution (due to the random Ω in Algorithm

1), $\lambda_i^{(\mathbf{W})}$ and $\lambda_i^{(\hat{\mathbf{W}}_K)}$ denote the eigenvalues of \mathbf{W} and $\hat{\mathbf{W}}_K$, respectively, organized in decreasing order, and M denotes the number of random vectors (columns) in Ω .

These bounds provide additional guidelines for setting γ and M . Note that due to the spectral decay, as N increases, if we appropriately increase M , the error tends to 0.

Proof. Applying Theorem 4.1 in Tropp et al. (2017) to symmetric PSD matrices, we have

$$\mathbb{E} \left[\left\| \mathbf{W} - \hat{\mathbf{W}}_K \right\|_1 \right] \leq \left(1 + \frac{K}{M-K-1} \right) \sum_{i=K+1}^N \lambda_i^{(\mathbf{W})}, \quad (10)$$

where $\|\cdot\|_1$ denotes the Schatten 1-norm. Using Theorem III.4.4 in Bhatia (2013) (see proof of Proposition 1),

$$\begin{aligned} \left\| \mathbf{W} - \hat{\mathbf{W}}_K \right\|_1 &= \sum_{i=1}^N \left| \lambda_i^{(\mathbf{W} - \hat{\mathbf{W}}_K)} \right| \geq \sum_{i=1}^N \left| \lambda_i^{(\mathbf{W})} - \lambda_i^{(\hat{\mathbf{W}}_K)} \right| \\ &= \sum_{i=1}^K \left| \lambda_i^{(\mathbf{W})} - \lambda_i^{(\hat{\mathbf{W}}_K)} \right| + \sum_{i=K+1}^N \lambda_i^{(\mathbf{W})}, \end{aligned}$$

since \mathbf{W} and $\hat{\mathbf{W}}_K$ are symmetric PSD and the ℓ^1 norm is a symmetric gauge function. The last equality is due to the construction of $\hat{\mathbf{W}}_K$ with rank K . Substituting this result into (10) and rearranging terms yields (8). The proof of the second part of the proposition appears in Appendix A. \square

The proposition holds in our case, since the matrices for which we estimate the eigenvalue decomposition, i.e., \mathbf{W} in (4), are symmetric positive-definite (SPD).

Proposition 2 demonstrates that the eigenvalue estimation error depends on the spectrum tail, the spectral decay, and the number M of random vectors. The spectral decay of our \mathbf{W} can be controlled to some extent by the kernel scale, σ^2 in (1), where larger σ^2 yields faster spectral decay. In our experiments, we set $M = 2K$ to obtain a reasonable coefficient multiplying the spectrum tail in the bound.

The computational complexity of this approximation algorithm is $\mathcal{O}(MN^2 + M^2N)$ (Tropp et al., 2017). For very large datasets, we have $M \ll N$ due to spectrum truncation.

Note that the dense matrix \mathbf{W} in (4) requires storing N^2 floating point numbers, which may be limiting in large datasets. To reduce this cost, the matrix \mathbf{W} can be computed in small batches of rows (in two iterations due to its normalization), since Algorithm 1 relies on \mathbf{W} only through matrix multiplications with Ω .

4.3. The Log-Euclidean Signature Distance

We can now define the log-Euclidean signature distance between unaligned datasets. Given two datasets $\{x_i^{(1)} \in$

Algorithm 1 Fixed-Rank approx. (Tropp et al., 2017)

Input: PSD matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, rank K , size M .
Output: Diagonal matrix $\mathbf{\Lambda} \in \mathbb{R}^{K \times K}$, such that $\mathbf{W} \approx \hat{\mathbf{W}}_K = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, for some $\mathbf{U} \in \mathbb{R}^{N \times K}$ with orthonormal columns.

Function APPROXEIGENVALUES(\mathbf{W} , K , M)
 $\mathbf{\Omega} \leftarrow \text{randn}(N, M)$
 $\mathbf{\Omega} \leftarrow \text{orth}(\mathbf{\Omega})$ {Orthonormal basis for range($\mathbf{\Omega}$)}
 $\mathbf{Y} \leftarrow \mathbf{W}\mathbf{\Omega}$
 $\nu \leftarrow \mu_{\text{norm}}(\mathbf{Y})$ { $\mu = 2.2 \cdot 10^{-16}$ }
 $\mathbf{Y} \leftarrow \mathbf{Y} + \nu\mathbf{\Omega}$
 $\mathbf{B} \leftarrow \mathbf{\Omega}^T \mathbf{Y}$
 $\mathbf{C} \leftarrow \text{chol}((\mathbf{B} + \mathbf{B}^T)/2)$ {Cholesky decomp.}
 $(\mathbf{U}, \mathbf{\Sigma}, \sim) \leftarrow \text{svd}(\mathbf{Y}/\mathbf{C})$
 $\mathbf{\Lambda} \leftarrow \max\{0, \mathbf{\Sigma}_{1:K, 1:K}^2 - \nu\mathbf{I}\}$
return sorted $\mathbf{\Lambda}$

\mathbb{R}^{d_1} and $\{x_i^{(2)}\}_{i=1}^{N_2} \in \mathbb{R}^{d_2}$, we construct diffusion operators $\mathbf{W}_1 \in \mathbb{R}^{N_1 \times N_1}$ and $\mathbf{W}_2 \in \mathbb{R}^{N_2 \times N_2}$ according to (4). We approximate the K leading eigenvalues $\{\hat{\lambda}_i^{(\ell)}\}_{i=1}^K$ for each matrix using Algorithm 1. The LES distance is then:

$$d_{\text{LES}}^2(\mathbf{W}_1, \mathbf{W}_2) := \sum_{i=1}^K \left(\log(\hat{\lambda}_i^{(1)} + \gamma) - \log(\hat{\lambda}_i^{(2)} + \gamma) \right)^2 \quad (11)$$

This procedure is summarized in Algorithm 2.

Remark. The LES distance is not a metric, since it does not satisfy the identity of indiscernibles, i.e. $d_{\text{LES}}(\mathbf{W}_1, \mathbf{W}_2) \neq 0$ is not guaranteed for $\mathbf{W}_1 \neq \mathbf{W}_2$. Specifically, datasets with underlying non-isometric but isospectral manifolds (or isospectral up to the K th eigenvalue) are indistinguishable. Isospectral manifolds exist (Kac, 1966) but typically require careful construction and are rarely encountered in practice, especially in low dimensions (Reuter et al., 2006). In addition, for $\mathbf{W}_1 \equiv \mathbf{W}_2$, $d_{\text{LES}}(\mathbf{W}_1, \mathbf{W}_2) = 0$ is guaranteed only with exact eigenvalue computation. However, our experimental results depict that the distance is close to 0 also with the eigenvalue approximation, and even for *different* point clouds sampled from the same manifold.

Remark. The proposed framework is not limited to diffusion operators \mathbf{W} . While our choice of \mathbf{W} is motivated by its convergence to the heat kernel, the LES distance can be computed using any SPD matrix, e.g., covariance matrices.

Relation to IMD. The LES distance is directly related to IMD (Tsitsulin et al., 2020), written using our notation as

$$d_{\text{IMD}}(\mathbf{W}_1, \mathbf{W}_2) = \sup_{t>0} c(t) \left| \sum_i (\lambda_i^{(1)})^t - \sum_i (\lambda_i^{(2)})^t \right|$$

where $c(t) = e^{-2(t+t^{-1})}$, $\lambda_i^{(\ell)} = \exp(-\mu_i^{(\ell)})$ and $\mu_i^{(\ell)}$ are (approx.) Laplace-Beltrami eigenvalues. This formulation demonstrates two key differences between LES and IMD. First, LES compares aligned spectral locations, whereas in

IMD the overall decay of the spectrum is compared in different temperature scales, controlled by the variable t . Second, log is applied to the eigenvalues by LES but not IMD, indicating that by taking a proper metric into account when comparing discrete approximations of heat kernels, LES actually compares approximate Laplace-Beltrami eigenvalues (related to eigenvalues of the heat kernel via a log), while IMD compares heat kernel eigenvalues, which decay quickly. We demonstrate empirically in Section 5 that the LES distance outperforms IMD in all considered applications, particularly when the datasets contain complex structures with differences in multiple scales.

Algorithm 2 Log-Euclidean Signature Distance

Input: r datasets $X^{(\ell)} = \{x_i^{(\ell)} \in \mathbb{R}^{d_\ell}\}_{i=1}^{N_\ell}$, $\ell = 1, \dots, r$.

Output: LES descriptors, $F_{\text{LES}}^{(\ell)} \in \mathbb{R}^K$, and LES distances, $d_{\text{LES}}(\ell, q)$, for $\ell, q \in [1, r]$.

for $\ell = 1$ **to** r **do**
 $\mathbf{W}_\ell \leftarrow \text{DM-SPD}(X^{(\ell)})$ {According to (4)}
 $\{\hat{\lambda}_i^{(\ell)}\}_{i=1}^K \leftarrow \text{APPROXEIGENVALUES}(\mathbf{W}_\ell)$
 $F_{\text{LES}}^{(\ell)} \leftarrow [\log(\hat{\lambda}_1^{(\ell)} + \gamma), \dots, \log(\hat{\lambda}_K^{(\ell)} + \gamma)]$
end for
for $\ell, q \in [1, r]$ **do**
 $d_{\text{LES}}^2(\ell, q) \leftarrow \|F_{\text{LES}}^{(\ell)} - F_{\text{LES}}^{(q)}\|_2^2$
end for

Spectral descriptors of data. Spectral descriptors appear in several applications, e.g., in shape analysis (Reuter et al., 2006; Cosmo et al., 2019) and molecular representations (Schrier, 2020). We similarly define a dataset descriptor based on the bound in (6) and the proposed distance (11), as

$$F_{\text{LES}}^{(\ell)} = [\log(\hat{\lambda}_1^{(\ell)} + \gamma), \dots, \log(\hat{\lambda}_K^{(\ell)} + \gamma)].$$

$F_{\text{LES}}^{(\ell)}$ embeds datasets into a K -dimensional space based on intrinsic geometry, facilitating dimensionality reduction and visualization and allowing for intuitive analysis of clusters.

5. Experimental Results

Throughout this section we compute LES according to Algorithm 2 with $\sigma^2 = 2 \text{ median}(d^2(x_i, x_j))$, $K \in \{200, 500\}$, $M = 2K$ and $\gamma \in [10^{-8}, 10^{-5}]$. The exact parameters used in each application are described in the appendices.

An additional experiment using LES to match hidden representations of NN is presented in Appendix E.

The code is available in <https://github.com/shnitzer/les-distance>.

5.1. Toy Example: Comparison of Geometric Shapes

We demonstrate the main advantages of the LES distance over competing methods on a toy example consisting of 3D and 4D point clouds. We sample points from two 2D tori, $T_2 = \mathcal{S}^1 \times \mathcal{S}^1$, and two 3D tori, $T_3 = \mathcal{S}^1 \times \mathcal{S}^1 \times \mathcal{S}^1$, which have the same major radius R_1 . The two T_2 tori differ in their minor radius, denoted by R_2 and $R_2^{Sc} = c \cdot R_2$, and the two T_3 tori share R_2 and differ in their minor radius, denoted by R_3 and $R_3^{Sc} = c \cdot R_3$. The torus equations are formulated such that $\lim_{R_3 \rightarrow 0} T_3 = T_2$, and $\lim_{R_2 \rightarrow 0} T_2 = T_1$, i.e., converges to a 1D circle. We denote the two tori with the scaled radius by T_2^{Sc} and T_3^{Sc} . We sample the 4 tori independently such that no pointwise correspondence exists. The exact equations of the tori embedding are given in Appendix B, eq. (13), along with a detailed description of the chosen parameters in this subsection.

We compute distances between the four tori for different c 's using LES, IMD, Geometry Score (GS) (Khrukov & Oseledets, 2018), and Gromov-Wasserstein (GW) (Peyré et al., 2016). To emphasize the conceptual differences between LES and IMD, we approximate IMD according to $d_{\text{IMD}}(\mathbf{W}_1, \mathbf{W}_2)$ (Section 4.3), using our estimated eigenvalues of \mathbf{W} . We denote this method by *IMD (our approach)*. Due to the different dimensions of the tori datasets, methods requiring that all datasets are embedded in the same space, such as MTD and CrossLID, cannot be applied. We compare our results with other standard topological persistence diagram metrics in Appendix B.2.

Figure 1 presents the mean and standard deviation of the distances between the different tori, averaged over 10 trials. Solid lines denote distances between tori with the same dimensionality and dashed lines denote distances between tori with different dimensions. As c decreases, it is expected that distances $d(T_2, T_2^{Sc})$, $d(T_3, T_3^{Sc})$ and $d(T_3, T_2^{Sc})$ will increase, since they converge to a distance between shapes of different dimensions, and that $d(T_2, T_3^{Sc})$ will decrease, since $\lim_{c \rightarrow 0} T_3^{Sc} = T_2$. Based on Figure 1, the only methods that exhibit the expected trends as c decreases are LES, GW, and IMD computed using our approach, which indicates that our approach for estimating leading heat kernel eigenvalues yields a more stable and accurate distance. The IMD distance in Figure 1(b) captures mainly the major dimensionality differences and is not sensitive to small differences in the radius scale. Moreover, the LES distance distinguishes between dimensionality differences and scale differences for a wider range of c values compared with GW and IMD (our approach), as depicted by the intersection between the distance lines of $d(T_2, T_2^{Sc})$ and $d(T_2, T_3^{Sc})$.

The distance trends captured by LES follow the trends of the exact log-Euclidean distances between the SPD matrices representing the tori data with ground truth pointwise correspondence. We present an example in Appendix B.3.

Running time and stability. Using the tori simulation, we compare computational complexity and stability of LES with IMD (with the exact k -nn distance), GW, and GS, for different dataset sizes, N . Figure 2(a) demonstrates that LES and IMD require significantly less computational time than GW and GS. While LES requires more computational time than IMD, its advantages in stability and accuracy are significant, as depicted in the examples in this section.

To evaluate the effect of different sample sizes on LES, we compute the distance $d(T_2, T_3)$ for different N values. Figure 2(b) presents this distance normalized by the distance at $N = 10^4$. This plot demonstrates that LES is the most robust to different sample sizes, even for relatively small N .

5.2. Single Cell Gene Expression Data Analysis

We further illustrate the power of LES in recovering fine structural details on single-cell RNA sequencing (scRNA-seq) data. ScRNA-seq is powerful for analyzing cell populations and developmental trajectories (Tanay & Regev, 2017). Recent works propose visualization techniques of such data for the analysis/comparison of cell sub-populations (Haghverdi et al., 2015; Moon et al., 2019). However, comparing datasets of scRNA-seq in different gene spaces, e.g., from different species or measurement modalities, remains a challenge (Shafer, 2019). Since in such tasks a clear ground truth is typically unavailable, we demonstrate the applicability of LES to identifying intrinsic structural similarities between gene-expression datasets via an illustrative task of recovering the time axis in cell differentiation data. In such experiments, all the cells start from a stem cell-like state and differentiate into various cell types. The gene-expressions of (random, non-repeating) subsets of these cells are measured at different time points throughout the developmental trajectory. Therefore, the time axis relates to the developmental stage of the cells and can indicate the expected similarity of the different subsets. This application inherently lacks pointwise alignment between samples, since scRNA-seq destroys the cell during measurement.

We apply LES to data of reprogrammed mouse embryonic cells (Schiebinger et al., 2019), used for learning developmental trajectories. It contains expressions of $\sim 2.5 \cdot 10^5$ cells and 1500 genes (after filtration), collected across 18 days of differentiation in half day intervals, resulting in 39 time points; a visualization is in Appendix C. Each time point consists of gene expressions of 2000–12000 cells from the same developmental stage. Our goal is to measure the gene-space similarities of these different cell groups. Each cell group (from every time point) is treated as a separate dataset, and we measure the distances between every pair of datasets using LES, IMD, Wasserstein distance (OT), GS, and MTD (Barannikov et al., 2021); we omit GW since it requires unreasonable processing times for all pairs of 39 time points.

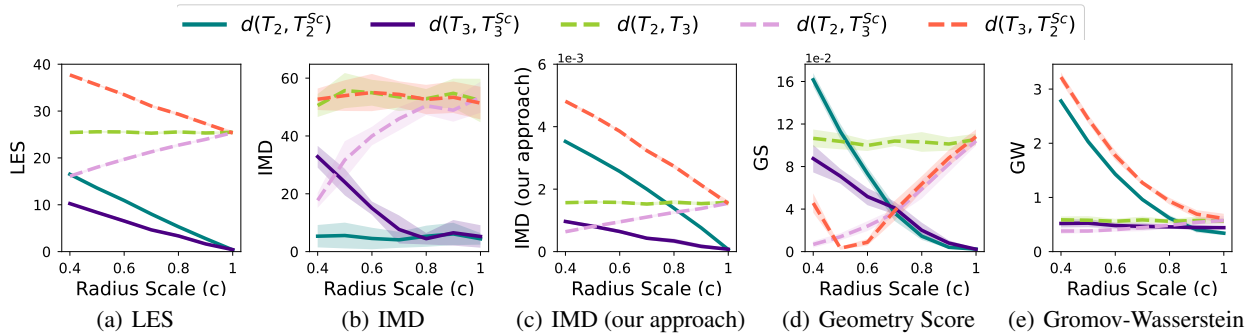


Figure 1. Comparison of distance measures on 2D and 3D tori with radii related by scaling c .




Figure 2. Run time and stability for different dataset sizes.

To visualize the resulting distance matrices, we compute their embeddings using diffusion maps (Coifman & Lafon, 2006). Figure 3 presents the leading embedded coordinates, denoted by f_1 and f_2 , where each point corresponds to a cell group from one time. The points are colored according to their measurement day, denoting their location along the developmental trajectory. The correlations between the time axis and the distances are below each plot. LES captures the developmental trajectory significantly better than IMD, GS, and MTD, as emphasized by their correlation with time and plot coloring. LES obtains a slightly higher correlation with time than OT, even though OT uses geometric proximity of the datasets while LES is intrinsic.

To slightly complicate the problem, we repeat this experiment and sample a different subset of genes at each time point ($\sim U[1200, 1500]$), so that cells from different days are represented in gene spaces of different dimensions. LES outperforms IMD in this setting as well, with slightly reduced correlations of 0.69 (LES) and 0.27 (IMD). Computation of OT and MTD in this setting is not possible, since cell groups from different time points are in different domains (different gene subsets and dimensionality). These results suggest that LES could be highly useful in comparing cell population structures from different gene spaces.

5.3. Identifying Difficult Tasks in Few-Shot Learning

Meta-learning algorithms, e.g., Model Agnostic Meta Learning (Finn et al., 2017) and MetaOptNet (Lee et al., 2019),

enable training models that can adapt to new tasks. In few-shot learning, they obtain classifiers for previously unseen classes using as little as a single labeled example per class. However, meta-learning can fail on tasks that are not sufficiently similar to the training tasks (Fallah et al., 2021). To anticipate such failures a practitioner may consider measuring dissimilarity between the tasks used during training and the incoming test task. A good dissimilarity score correlates negatively with the performance on tasks to provide guidance on whether it is safe to deploy a meta-learning model on a new task without collecting labels. Namely, in such a setting we aim to predict whether a specific embedder (NN) will result in high few-shot classification accuracy in an unsupervised manner. Our study complements recent literature on meta-learning safety (Goldblum et al., 2020; Slack et al., 2020; Agarwal et al., 2021).

Suppose we have several training tasks (datasets) $\{\tilde{X}^{(\ell)} \in \mathbb{R}^{N_\ell \times d}\}_{\ell=1}^n$ and a new task $\tilde{X}' \in \mathbb{R}^{N' \times d}$. We train MetaOptNet-SVM (Lee et al., 2019) for benchmark datasets CIFAR-FS (Bertinetto et al., 2018) and FC100 (Oreshkin et al., 2018) in 1-, 5-, and 10-shot settings. Its architecture has a feature extractor shared across tasks, generating task embeddings denoted by $X^{(\ell)}$ and X' for train and new tasks, respectively, and a classification head for each new task, learned during adaptation for evaluating accuracy. We compute dissimilarity between embeddings of the new and train tasks as $\mathcal{A}(\{d(X^{(\ell)}, X')\}_{\ell=1}^n)$, where \mathcal{A} is some aggregator, e.g. the mean, and $d(\cdot, \cdot)$ is a dataset distance.

For each benchmark and shot setting we sample several new tasks from train/test task distributions and their mixtures. In Table 1 we report the correlation between the task accuracy and the dissimilarity score (with \mathcal{A} as the mean) based on various dataset distances. We compare LES against IMD, GW, GS and Wasserstein distance (OT) (here the dimensionality of all datasets is the same, so OT can be applied). LES outperforms IMD, GS and OT in all settings by a large margin. Performance of GW is comparable, but it is orders of magnitude slower. The accuracy on the test tasks is also reported in Table 1, which clearly distinguishes the two datasets (accuracy on train tasks for both is above 90% in all

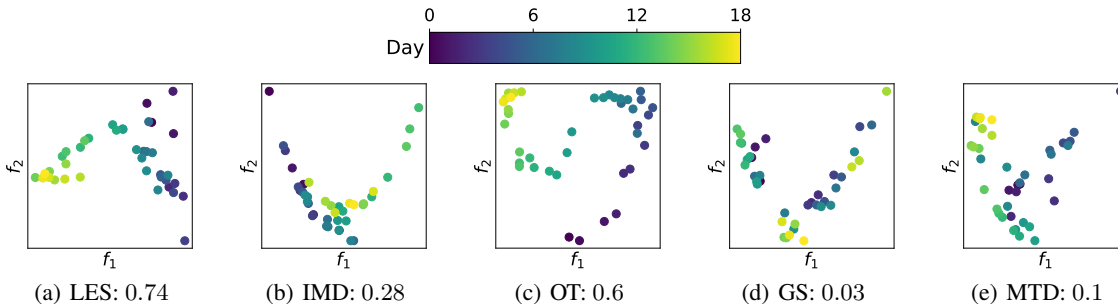


Figure 3. Gene-expression datasets embedding based on the different distances. Correlation with time is reported by the method labels.

Table 1. Few-shot learning: MetaOptNet test accuracies and correlations (20 trials) with task dissimilarities.

	FC100 1-shot	FC100 5-shot	FC100 10-shot	CIFAR-FS 1-shot	CIFAR-FS 5-shot	CIFAR-FS 10-shot
Acc.	38.19±0.48%	54.45±0.49%	60.52±0.48%	70.79±0.69%	83.98±0.44%	87.11±0.40%
LES	-0.934 ±0.034	-0.945 ±0.032	-0.935 ±0.032	-0.671±0.128	-0.698 ±0.155	-0.657 ±0.175
IMD	-0.184±0.250	-0.303±0.225	-0.210±0.235	-0.151±0.225	-0.015±0.220	-0.032±0.281
OT	0.739±0.102	0.605±0.126	0.579±0.133	0.453±0.170	0.269±0.214	0.180±0.231
GW	-0.919±0.034	-0.921±0.046	-0.914±0.045	-0.677 ±0.126	-0.672±0.167	-0.582±0.171
GS	-0.735±0.101	-0.570±0.154	-0.581±0.218	-0.384±0.201	-0.034±0.258	-0.012±0.281

settings). FC100 is a more challenging variant of CIFAR-FS that was created to reduce similarity between train and test tasks (Oreshkin et al., 2018). As demonstrated by the correlations for LES, which are close to -1 (ideal), LES captures the train/test task dissimilarity in FC100 and can be used to anticipate poor performance of meta-learners. On the other hand, CIFAR-FS train and test tasks are more similar, as indicated by the high accuracy, and thus are harder to distinguish; the correlation is worse for all dataset distances. We report experiment details, run-time comparison, and additional results in Appendix D.

We note some related works. Achille et al. (2019) and Nguyen et al. (2020) proposed methods for embedding meta-learning tasks using gradients of an auxiliary neural network and a specialized topic model correspondingly. Both methods require labeled examples to produce the embeddings, while LES is unsupervised. Venkitaraman et al. (2020) and Zhou et al. (2021) incorporated task similarity into training of meta-learners. Their approaches learn task descriptors during training and require specialized architectures and loss functions. Using LES for similarity-aware training of meta-learners can be a promising direction to explore.

6. Conclusion

We proposed the LES distance, a geometric, data-driven intrinsic distance for unaligned datasets of different dimensions that lower-bounds the log-Euclidean metric between SPD matrices. LES provides a stable and efficient distance

measure with capabilities of recovering meaningful structural differences on gene expression data and NN embeddings. Due to LES’s generality, it can be applied to a wide range of other tasks and datasets. In particular, our results on task similarity in few-shot learning suggest applicability to quantifying distribution shifts. Realistic distribution shifts can cause failures of many existing machine learning algorithms (Koh et al., 2021; Santurkar et al., 2021). Using LES to anticipate and diagnose poor out-of-distribution generalization is an interesting future work direction.

Acknowledgements

We thank the anonymous reviewers for their valuable feedback and helpful recommendations. We also wish to thank Madeleine Udell for useful suggestions. The MIT Geometric Data Processing group acknowledges the generous support of Army Research Office grants W911NF2010168 and W911NF2110293, of Air Force Office of Scientific Research award FA9550-19-1-031, of National Science Foundation grants IIS-1838071 and CHS-1955697, from the CSAIL Systems that Learn program, from the MIT-IBM Watson AI Laboratory, from the Toyota-CSAIL Joint Research Center, and from a gift from Adobe Systems. TS acknowledges the generous support of the Schmidt Futures Israeli Women’s Postdoctoral Award and the Viterbi Fellowship, Technion.

References

- Achille, A., Lam, M., Tewari, R., Ravichandran, A., Maji, S., Fowlkes, C. C., Soatto, S., and Perona, P. Task2vec: Task embedding for meta-learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6430–6439, 2019.
- Agarwal, M., Yurochkin, M., and Sun, Y. On sensitivity of meta-learning to support data. *Advances in Neural Information Processing Systems*, 34, 2021.
- Arsigny, V., Fillard, P., Pennec, X., and Ayache, N. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 56(2):411–421, 2006.
- Arsigny, V., Fillard, P., Pennec, X., and Ayache, N. Geometric means in a novel vector space structure on symmetric positive-definite matrices. *SIAM journal on matrix analysis and applications*, 29(1):328–347, 2007.
- Barachant, A., Bonnet, S., Congedo, M., and Jutten, C. Classification of covariance matrices using a Riemannian-based kernel for BCI applications. *Neurocomputing*, 112: 172–178, 2013.
- Barannikov, S., Trofimov, I., Sotnikov, G., Trimbach, E., Korotin, A., Filippov, A., and Burnaev, E. Manifold topology divergence: a framework for comparing data manifolds. *Advances in Neural Information Processing Systems*, 34, 2021.
- Barua, S., Ma, X., Erfani, S. M., Houle, M. E., and Bailey, J. Quality evaluation of GANs using cross local intrinsic dimensionality. *arXiv preprint arXiv:1905.00643*, 2019.
- Belkin, M. and Niyogi, P. Laplacian Eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- Belkin, M. and Niyogi, P. Convergence of Laplacian Eigenmaps. *Advances in Neural Information Processing Systems*, 19:129, 2007.
- Bertinetto, L., Henriques, J. F., Torr, P., and Vedaldi, A. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2018.
- Bhatia, R. *Matrix analysis*, volume 169. Springer Science & Business Media, 2013.
- Bonnabel, S. and Sepulchre, R. Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1055–1070, 2010.
- Cela, E. *The quadratic assignment problem: theory and algorithms*, volume 1. Springer Science & Business Media, 2013.
- Coifman, R. R. and Hirn, M. J. Diffusion maps for changing data. *Applied and computational harmonic analysis*, 36(1):79–107, 2014.
- Coifman, R. R. and Lafon, S. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.
- Collard, A., Bonnabel, S., Phillips, C., and Sepulchre, R. Anisotropy preserving dti processing. *International journal of computer vision*, 107(1):58–74, 2014.
- Cosmo, L., Panine, M., Rampini, A., Ovsjanikov, M., Bronstein, M. M., and Rodola, E. Isospectralization, or how to hear shape, style, and correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7529–7538, 2019.
- Ding, X. and Wu, H.-T. Impact of signal-to-noise ratio and bandwidth on graph laplacian spectrum from high-dimensional noisy point cloud. *arXiv preprint arXiv:2011.10725*, 2020.
- Dunson, D. B., Wu, H.-T., and Wu, N. Spectral convergence of graph Laplacian and heat kernel reconstruction in L^∞ from random samples. *Applied and Computational Harmonic Analysis*, 2021.
- Edelsbrunner, H., Harer, J., et al. Persistent homology—a survey. *Contemporary mathematics*, 453:257–282, 2008.
- Fallah, A., Mokhtari, A., and Ozdaglar, A. Generalization of model-agnostic meta-learning algorithms: Recurring and unseen tasks. In *Advances in Neural Information Processing Systems*, 2021.
- Fefferman, C., Mitter, S., and Narayanan, H. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pp. 1126–1135. PMLR, 2017.
- Goldblum, M., Fowl, L., and Goldstein, T. Adversarially robust few-shot learning: A meta-learning approach. *Advances in Neural Information Processing Systems*, 33, 2020.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.

- Guan, D., Xiang, D., Tang, X., Wang, L., and Kuang, G. Covariance of textural features: a new feature descriptor for SAR image classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(10):3932–3942, 2019.
- Ha Quang, M., San Biagio, M., and Murino, V. Log-Hilbert-Schmidt metric between positive definite operators on Hilbert spaces. In *Advances in neural information processing systems*, pp. 388–396, 2014.
- Ha Quang, M., Biagio, M. S., Bazzani, L., and Murino, V. Approximate log-Hilbert-Schmidt distances between covariance operators for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5195–5203, 2016.
- Haghverdi, L., Buettner, F., and Theis, F. J. Diffusion maps for high-dimensional single-cell analysis of differentiation data. *Bioinformatics*, 31(18):2989–2998, 2015.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Huang, Z., Wang, R., Shan, S., Li, X., and Chen, X. Log-Euclidean metric learning on symmetric positive definite manifold with application to image set classification. In *International conference on machine learning*, pp. 720–729. PMLR, 2015.
- Jacomy, M., Venturini, T., Heymann, S., and Bastian, M. ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PLoS one*, 9(6):e98679, 2014.
- Kac, M. Can one hear the shape of a drum? *The american mathematical monthly*, 73(4P2):1–23, 1966.
- Katz, O., Lederman, R. R., and Talmon, R. Spectral flow on the manifold of SPD matrices for multimodal data processing. *arXiv preprint arXiv:2009.08062*, 2020.
- Khrulkov, V. and Oseledets, I. Geometry score: A method for comparing generative adversarial networks. In *International Conference on Machine Learning*, pp. 2621–2629. PMLR, 2018.
- Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Gao, I., et al. WILDS: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021.
- Kornblith, S., Norouzi, M., Lee, H., and Hinton, G. Similarity of neural network representations revisited. In *International Conference on Machine Learning*, pp. 3519–3529. PMLR, 2019.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Lee, K., Maji, S., Ravichandran, A., and Soatto, S. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10657–10665, 2019.
- Li, H., Linderman, G. C., Szlam, A., Stanton, K. P., Kluger, Y., and Tygert, M. Algorithm 971: An implementation of a randomized algorithm for principal component analysis. *ACM Transactions on Mathematical Software (TOMS)*, 43(3):1–14, 2017.
- Mead, A. Review of the development of multidimensional scaling methods. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 41(1):27–39, 1992.
- Mémoli, F. A spectral notion of Gromov–Wasserstein distance and related methods. *Applied and Computational Harmonic Analysis*, 30(3):363–401, 2011.
- Mishne, G. and Cohen, I. Multiscale anomaly detection using diffusion maps. *IEEE Journal of selected topics in signal processing*, 7(1):111–123, 2012.
- Moon, K. R., van Dijk, D., Wang, Z., Gigante, S., Burkhardt, D. B., Chen, W. S., Yim, K., van den Elzen, A., Hirn, M. J., Coifman, R. R., et al. Visualizing structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12):1482–1492, 2019.
- Nguyen, C., Do, T.-T., and Carneiro, G. Similarity of classification tasks. In *NeurIPS Meta Learning Workshop*, 2020.
- Oreshkin, B. N., Rodríguez, P., and Lacoste, A. TADAM: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Pennec, X., Fillard, P., and Ayache, N. A Riemannian framework for tensor computing. *International Journal of computer vision*, 66(1):41–66, 2006.
- Peyré, G., Cuturi, M., and Solomon, J. Gromov-Wasserstein averaging of kernel and distance matrices. In *International Conference on Machine Learning*, pp. 2664–2672. PMLR, 2016.
- Reuter, M., Wolter, F.-E., and Peinecke, N. Laplace–Beltrami spectra as ‘Shape-DNA’ of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.
- Santurkar, S., Tsipras, D., and Madry, A. BREEDS: Benchmarks for subpopulation shift. In *International Conference on Learning Representations*, 2021.

- Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube, P., et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- Schrier, J. Can one hear the shape of a molecule (from its Coulomb matrix eigenvalues)? *Journal of Chemical Information and Modeling*, 60(8):3804–3811, 2020.
- Shafer, M. E. Cross-species analysis of single-cell transcriptomic data. *Frontiers in cell and developmental biology*, 7:175, 2019.
- Shnitzer, T., Wu, H.-T., and Talmon, R. Spatiotemporal analysis using Riemannian composition of diffusion operators. *arXiv preprint arXiv:2201.08530*, 2022.
- Slack, D., Friedler, S. A., and Givental, E. Fairness warnings and Fair-MAML: learning fairly with minimal data. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 200–209, 2020.
- Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. Striving for simplicity: The all convolutional net. In *International Conference on Learning Representations Workshop*, 2015.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.
- Talmon, R. and Coifman, R. R. Empirical intrinsic geometry for nonlinear modeling and time series filtering. *Proceedings of the National Academy of Sciences*, 110(31):12535–12540, 2013.
- Talmon, R., Cohen, I., Gannot, S., and Coifman, R. R. Diffusion maps for signal processing: A deeper look at manifold-learning techniques based on kernels and graphs. *IEEE signal processing magazine*, 30(4):75–86, 2013.
- Tanay, A. and Regev, A. Scaling single-cell genomics from phenomenology to mechanism. *Nature*, 541(7637):331–338, 2017.
- Tralie, C., Saul, N., and Bar-On, R. Ripsr.py: A lean persistent homology library for python. *The Journal of Open Source Software*, 3(29):925, Sep 2018. doi: 10.21105/joss.00925. URL <https://doi.org/10.21105/joss.00925>.
- Tropp, J. A., Yurtsever, A., Udell, M., and Cevher, V. Fixed-rank approximation of a positive-semidefinite matrix from streaming data. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Tsitsulin, A., Mottin, D., Karras, P., Bronstein, A., and Müller, E. NetLSD: hearing the shape of a graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2347–2356, 2018.
- Tsitsulin, A., Munkhoeva, M., Mottin, D., Karras, P., Bronstein, A., Oseledets, I., and Müller, E. The shape of data: Intrinsic distance for data distributions. In *International Conference on Learning Representations*, 2020.
- Venkitaraman, A., Hansson, A., and Wahlberg, B. Task-similarity aware meta-learning through nonparametric kernel regression. *arXiv preprint arXiv:2006.07212*, 2020.
- Zhou, P., Zou, Y., Yuan, X.-T., Feng, J., Xiong, C., and Hoi, S. Task similarity aware meta learning: Theory-inspired improvement on MAML. In *Uncertainty in Artificial Intelligence*, pp. 23–33. PMLR, 2021.

A. Proof of the Second Bound in Proposition 2

We reiterate the claim of the second part of Proposition 2 for convenience.

Proposition 2. *The rank- K approximation of a real symmetric PSD matrix, \mathbf{W} , computed according to Algorithm 1, and denoted by $\hat{\mathbf{W}}_K$, satisfies the following error bounds.*

2. *The γ -regularized log-eigenvalue error is bounded by:*

$$\mathbb{E} \left[\sum_{i=1}^K \left| \log(\lambda_i^{(\mathbf{W})} + \gamma) - \log(\lambda_i^{(\hat{\mathbf{W}}_K)} + \gamma) \right| \right] \leq \frac{1.5K}{(M - K - 1)(\lambda_K^{(\mathbf{W})} + \gamma)} \sum_{i=K+1}^N \lambda_i^{(\mathbf{W})}$$

for $\frac{|\lambda_i^{(\mathbf{W})} - \lambda_i^{(\hat{\mathbf{W}}_K)}|}{\lambda_i^{(\mathbf{W})} + \gamma} \leq 0.5828$, where the expectations are over the normal distribution (due to the random Ω in Algorithm 1), $\lambda_i^{(\mathbf{W})}$ and $\lambda_i^{(\hat{\mathbf{W}}_K)}$ denote the eigenvalues of \mathbf{W} and $\hat{\mathbf{W}}_K$, respectively, organized in decreasing order, and M denotes the number of random vectors (columns) in Ω .

Proof. To bound the log-eigenvalues, we denote the i -th approximation error by $\lambda_i^{(\hat{\mathbf{W}}_K)} = \lambda_i^{(\mathbf{W})} \pm \epsilon_i$:

$$\begin{aligned} \sum_{i=1}^K \left| \log(\lambda_i^{(\hat{\mathbf{W}}_K)} + \gamma) - \log(\lambda_i^{(\mathbf{W})} + \gamma) \right| &= \sum_{i=1}^K \left| \log(\lambda_i^{(\mathbf{W})} \pm \epsilon_i + \gamma) - \log(\lambda_i^{(\mathbf{W})} + \gamma) \right| = \sum_{i=1}^K \left| \log \left(1 \pm \frac{\epsilon_i}{\lambda_i^{(\mathbf{W})} + \gamma} \right) \right| \\ &\stackrel{(*)}{\leq} 1.5 \sum_{i=1}^K \left| \frac{\epsilon_i}{\lambda_i^{(\mathbf{W})} + \gamma} \right| \stackrel{(**)}{\leq} 1.5 \sum_{i=1}^K \left| \frac{\epsilon_i}{\lambda_K^{(\mathbf{W})} + \gamma} \right| = \frac{1.5}{\lambda_K^{(\mathbf{W})} + \gamma} \sum_{i=1}^K |\lambda_i^{(\mathbf{W})} - \lambda_i^{(\hat{\mathbf{W}}_K)}| \end{aligned} \quad (12)$$

where transition $(*)$ is due to $\log(1+x) \leq x$ for $x \geq 0$ and $|\log(1-x)| \leq 1.5x$ for $0 \leq x \leq 0.5828$, and $(**)$ is since the eigenvalues $\lambda_i^{(\mathbf{W})}$ are positive and decreasing. Therefore, this derivation holds when $\frac{|\lambda_i^{(\mathbf{W})} - \lambda_i^{(\hat{\mathbf{W}}_K)}|}{\lambda_i^{(\mathbf{W})} + \gamma} \leq 0.5828$. Taking the expectation of (12) and bounding it with (8) yields the error bound in the proposition. \square

B. Tori Simulation: Additional Details and Results

B.1. Tori Equations and Parameter Choices

The tori embedding equations used to generate the data for the toy example in Subsection 5.1 are given by:

$$T_2 : \left\{ \begin{array}{l} x[i] = (R_1 + R_2 \cos(\theta_2[i])) \cos(\theta_1[i]) \\ y[i] = (R_1 + R_2 \cos(\theta_2[i])) \sin(\theta_1[i]) \\ z[i] = R_2 \sin(\theta_2[i]) \end{array} \right\}, T_3 : \left\{ \begin{array}{l} x[i] = (R_1 + (R_2 + R_3 \cos(\theta_3[i])) \cos(\theta_2[i])) \cos(\theta_1[i]) \\ y[i] = (R_1 + (R_2 + R_3 \cos(\theta_3[i])) \cos(\theta_2[i])) \sin(\theta_1[i]) \\ z[i] = (R_2 + R_3 \cos(\theta_3[i])) \sin(\theta_2[i]) \\ w[i] = R_3 \sin(\theta_3[i]) \end{array} \right\} \quad (13)$$

where $\theta_1[i], \theta_2[i], \theta_3[i] \sim \text{U}[0, 2\pi]$, $i = 1, \dots, N$ and $R_1 = 10$, $R_2 = 3$, $R_3 = 1$. T_2^{Sc} was defined similarly to T_2 by replacing R_2 with $R_2^{Sc} = c \cdot R_2$ and T_3^{Sc} was defined by replacing R_3 in T_3 with $R_3^{Sc} = c \cdot R_3$. Figure 4 presents examples of data sampled from the 2D torus and the scaled 2D torus.

For Figure 1, we sampled $N = 3000$ points of θ_1 , θ_2 and θ_3 in each torus independently, such that no known pointwise correspondence exists. In our comparisons we used the default parameters for GW, IMD with the exact k -nearest neighbors distance matrix and GS with $n = 2000$. Using non-default parameters did not lead to significant differences in the results.

LES hyperparameters. In the computation of LES, we set $K = 200$ and $\gamma = 10^{-8}$ since the simulated data is clean and the eigenvalues decay quickly. In all experiments, the choice of γ was motivated by the magnitude of the smallest eigenvalues.

B.2. Comparison to Persistence Diagram Metrics

As additional examples of the differences between our approach and topological data analysis methods, we compare LES to the bottleneck distance between H_0 , H_1 , and H_2 persistence diagrams (Edelsbrunner et al., 2008) on the tori datasets.




Figure 4. Visualization of (a) the 2D torus and (b) the scaled 2D torus with $c = 0.5$.

We compute the H_0 , H_1 , and H_2 persistence diagrams (Tralie et al., 2018) for samples from each of the 4 tori and then compute the pairwise bottleneck distances between them. We sample only $N = 1000$ points from each torus in this example due to the long run times of the persistence diagram computations. Figure 5 presents the average and standard deviation of the distances (avg. over 10 trials) as a function of the scale parameter c , similarly to Figure 1, along with LES for comparison. While the H_0 and H_1 bottleneck distances do capture meaningful differences between the tori, as expected, LES distinguishes between dimensionality and scale differences for a wider range of c values.

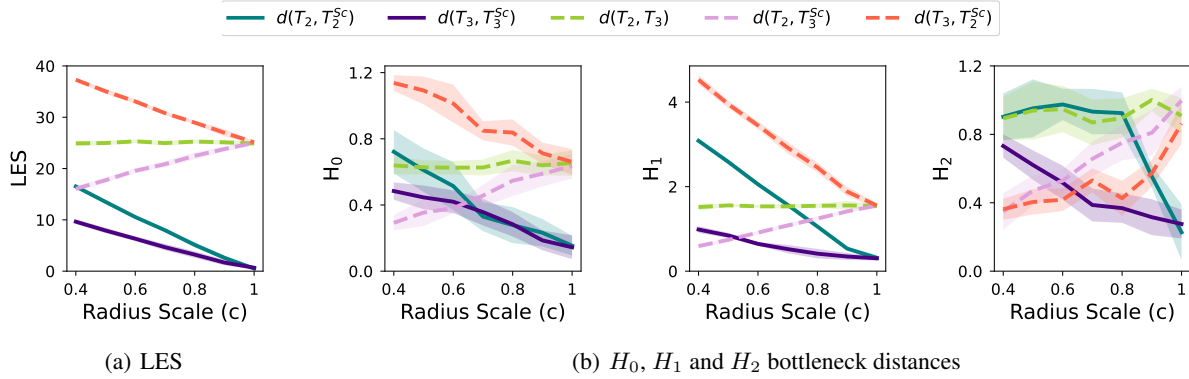


Figure 5. LES vs. H_0 , H_1 and H_2 bottleneck distances on the 2D and 3D tori example.

B.3. Comparisons to the Log-Euclidean Distance Between Aligned Datasets

To evaluate the use of Riemannian SPD geometry in our setting and the relation between the log-Euclidean (LE) distance and LES, we generated the tori datasets *with* pointwise correspondence by using the same samples of θ_1 , θ_2 and θ_3 for all tori. We constructed the SPD matrix \mathbf{W} (4) for each torus and computed the LE distance, the affine-invariant (AI) distance, $d_{AI}(\mathbf{W}_1, \mathbf{W}_2) = \|\log(\mathbf{W}_1^{-1}\mathbf{W}_2)\|_F$, the Euclidean distance, $d_{EU}(\mathbf{W}_1, \mathbf{W}_2) = \|\mathbf{W}_1 - \mathbf{W}_2\|_F$, and the Spectral Gromov-Wasserstein (SpecGW) distance (Mémoli, 2011), $d_{SpecGW}(\mathbf{W}_1, \mathbf{W}_2) = \sup_{t>0} e^{-(t+t^{-1})} \|\mathbf{W}_1^t - \mathbf{W}_2^t\|$ (for known correspondence). All distances were computed using the ground truth correspondence. To avoid violations of the SPD structure of \mathbf{W} due to numerical issues, we limit the number of samples on the tori to $N = 100$ in this comparison. Recall that LES is a lower bound of the LE distance and that IMD is a lower bound of SpecGW. Figure 6 presents the average and standard deviation of the computed distances (avg. over 10 trials) between the tori for different values of c , similarly to Figure 1. This figure demonstrates that both the LE and AI distances follow a similar trend and are able to distinguish between scale and dimensionality differences for a wider range of c values, compared with the Euclidean and SpecGW distances. Note that both LES and IMD (our approach), Figures 1(a) and 1(c), approximately follow the same trends as LE and SpecGW. This justifies our approach for approximating and comparing the eigenvalues of the heat kernel and indicates that our lower bound of the LE distance indeed captures similar structures as the LE distance in this example. Note that the values of SpecGW are smaller than IMD (our approach) in Figure 1(c). This may be due to the truncation of

the spectrum in IMD (our approach).

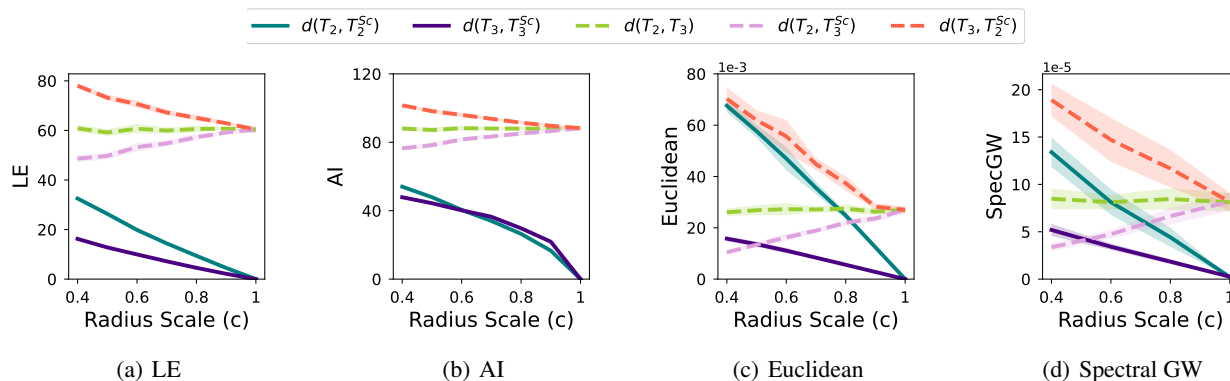


Figure 6. LE, AI, Euclidean and SpecGW distances on the 2D and 3D tori example using true pointwise correspondence (coupled angles).

C. Gene Expression Data: Additional Details

Figure 7 presents the gene space of the scRNA-seq data used in Subsection 5.2, colored according to the time of collection. At each time point, gene expressions from 2000–12000 cells were collected. This figure was created using code provided by Schiebinger et al. (2019), which applies force-directed layout embedding (Jacomy et al., 2014) for the visualization.

In this experiment, we compare the cell group structures from different time points, treating each group as a separate dataset. For LES, we construct the operator \mathbf{W}_ℓ at each time point, $\ell = 1, \dots, 39$, and then compare the descriptors, $F_{LES}^{(\ell)}$, for every pair of time points.

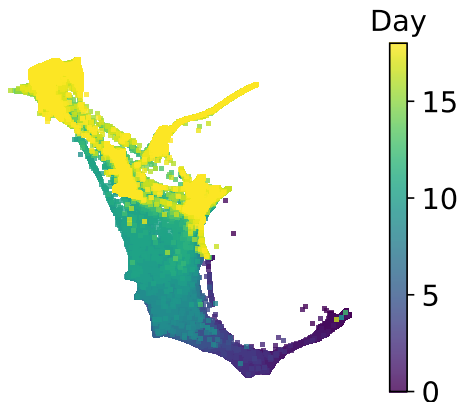


Figure 7. Gene expression data visualization (Schiebinger et al., 2019).

LES hyperparameters. In the gene expression data we set the LES parameters to $K = 500$ and $\gamma = 10^{-5}$, which is slightly larger than the maximal λ_{500} in all time points, to reduce the effect of noise in the data.

D. Identifying Difficult Tasks in Few-Shot Learning: Additional Details and Results

We present details regarding the experiment described in Section 5.3 and summarized in Table 1.

Benchmark datasets. Both CIFAR-FS (Bertinetto et al., 2018) and FC100 (Oreshkin et al., 2018) are few-shot learning benchmarks derived from the CIFAR-100 dataset (Krizhevsky et al., 2009). First the 100 classes are disjointly partitioned into train, validation, and test. CIFAR-FS partitions the classes randomly, while FC100 utilizes the class similarity information (100 classes are grouped into 20 superclasses in CIFAR-100) to reduce the semantic overlap between train and

test classes. Due to the partitioning, FC100 is a more challenging benchmark making it especially interesting for our study. Then the training tasks are obtained by randomly selecting classes and the corresponding images from the set of training classes. Similarly, test tasks are sampled from the set of test classes. All tasks in our experiments are 5-way, i.e. consist of 5 unique classes, following standard practice for these benchmarks.

Feature extractor and test accuracies. For each benchmark and shot setting we train MetaOptNet-SVM following the original implementation (Lee et al., 2019). Each MetaOptNet network consists of a ResNet-12 feature extractor that is shared across tasks during evaluation and a linear support vector machine classification “head” that is trained on k labeled samples, the support (adaptation) data, for each task to compute accuracies, where k denotes the number of “shots” per class. Test accuracies (first row in Table 1) in each setting are computed using 1000 random test tasks, and we report mean and standard deviation.

Evaluating correlations. To evaluate correlations we first sample $n = 10$ tasks from the train task distribution. Corresponding images are then passed through the MetaOptNet feature extractor resulting in a collection of n train task embeddings $\{X^{(\ell)} \in \mathbb{R}^{N_\ell \times d}\}_{\ell=1}^n$, $d = 2560$. To obtain a range of new tasks from easier to more difficult ones, we consider mixtures of train and test classes: only test classes (most difficult), 4 test and 1 train classes, 3 test and 2 train classes, 2 test and 3 train classes, 1 test and 4 train classes, and only train classes (easiest). For each of the six combinations we sample 3 tasks, resulting in 18 new tasks for computing the correlations. For each new task we obtain its embeddings, $X' \in \mathbb{R}^{N' \times d}$, compute accuracy using k (number of shots) random labeled examples per class, and compute a measure of dissimilarity with the train tasks, $\mathcal{A}(\{d(X^{(\ell)}, X')\}_{\ell=1}^n)$. For dataset distance computations, $d(\cdot, \cdot)$, we use all available (unlabeled) images, i.e. $N_\ell = N' = N = 3000 \forall l$ (600 images per each of the 5 classes in a task). We note that in practice this would require collecting unlabeled examples from a new task to evaluate the dissimilarity, however it is a minor limitation due to easy availability of unlabeled data in many application domains.

For a dataset distance that is informative of the task difficulty, we expect dissimilarity values and accuracies to be negatively correlated, i.e., large distances should indicate low accuracies. An informative dataset distance can be used in practice to anticipate difficult tasks before collecting labels and to prevent deployment of potentially poor-performing models. To compare different dataset distances we report Pearson correlation coefficient between accuracies and dissimilarity measures evaluated on 18 new tasks sampled as previously described. We considered three choices of the aggregation function $\mathcal{A}(\cdot)$: mean (Table 1), min (Table 2), and max (Table 3). We repeat each experiment 20 times and report mean and standard deviation of the correlations. Comparing the results across tables, we note that the choice of the aggregation function does not alter the performance much. Our proposed LES is the overall best performing dataset distance regardless of the aggregation function choice.

Run time comparison. In Table 4 we summarize run times for computing the dissimilarity measures based on each of the dataset distances considered in this experiment. We report mean and standard deviation over 360 dissimilarity computations (18 tasks for each of the 20 experiment repetitions). IMD and LES are the fastest, with IMD being slightly faster overall. To compute the dissimilarity measure for a new task we need to evaluate each dataset distance $n = 10$ times, i.e. one evaluation for each of the reference train tasks. LES and IMD are more efficient because they allow pre-computing train task descriptors and reusing them with every new task. Thus, their dissimilarity evaluation run-time is effectively equal to the run-time of computing task descriptor for a new task (the run time of $n = 10$ Euclidean distance computations with the pre-computed train task descriptors is negligible). On the other hand, dataset distances such as OT and GW do not have a notion of task descriptors and require $n = 10$ respective dataset distance evaluations to compute the dissimilarity measure. While GW is comparable to LES in terms of the correlation scores, its runtime might be problematic in practice.

Why does LES work well? We present a hypothesis explaining the strong correlation between LES-based dissimilarity measure and task accuracies. MetaOptNet uses linear SVM for classification, and thus its embeddings need to be linearly separable with respect to classes for good performance. In Figure 8 (left), we present 2-dimensional visualization of MetaOptNet embeddings obtained with Multidimensional Scaling (Mead, 1992) on one of the FC100 5-shot train tasks. While not guaranteed in general without the label information, it is likely that the 5 well-separated clusters we see correspond to the 5 classes from this task. Accuracy on this task is 98.2%, and we hypothesize that if the embeddings on a new task look similar, i.e. have number of clusters corresponding to the number of classes, then it is likely that the corresponding few-shot accuracy will be high. In Figure 8 (center) and (right) we show embeddings of two new tasks with distinct few-shot accuracies: 66.2% for the (center) embeddings without cluster structure, and 96.0% for the (right) embeddings with pronounced cluster structure. Evidently, the geometric structure of the (left) and (right) embeddings is a lot more similar

Table 2. Few-shot learning: MetaOptNet test accuracies and correlations (20 trials) with task dissimilarities using **min**-aggregation.

	FC100 1-shot	FC100 5-shot	FC100 10-shot	CIFAR-FS 1-shot	CIFAR-FS 5-shot	CIFAR-FS 10-shot
Acc.	38.19±0.48%	54.45±0.49%	60.52±0.48%	70.79±0.69%	83.98±0.44%	87.11±0.40%
LES	-0.913 ±0.038	-0.922 ±0.042	-0.913 ±0.046	-0.588 ±0.150	-0.676 ±0.131	-0.597 ±0.203
IMD	-0.356±0.258	-0.399±0.223	-0.359±0.203	-0.116±0.196	-0.011±0.221	-0.044±0.272
OT	0.314±0.264	0.099±0.253	0.077±0.267	0.243±0.205	-0.013±0.283	-0.031±0.277
GW	-0.898±0.036	-0.869±0.090	-0.831±0.166	-0.525±0.309	-0.506±0.272	-0.388±0.284
GS	-0.744±0.080	-0.546±0.241	-0.555±0.245	-0.306±0.248	-0.094±0.230	-0.051±0.257

Table 3. Few-shot learning: MetaOptNet test accuracies and correlations (20 trials) with task dissimilarities using **max**-aggregation.

	FC100 1-shot	FC100 5-shot	FC100 10-shot	CIFAR-FS 1-shot	CIFAR-FS 5-shot	CIFAR-FS 10-shot
Acc.	38.19±0.48%	54.45±0.49%	60.52±0.48%	70.79±0.69%	83.98±0.44%	87.11±0.40%
LES	-0.938 ±0.032	-0.942 ±0.029	-0.931 ±0.029	-0.663±0.167	-0.690 ±0.170	-0.644 ±0.194
IMD	-0.119±0.269	-0.298±0.223	-0.178±0.284	-0.129±0.194	-0.026±0.247	-0.028±0.236
OT	0.754±0.096	0.634±0.124	0.621±0.146	0.432±0.184	0.283±0.231	0.184±0.220
GW	-0.924±0.034	-0.915±0.052	-0.908±0.052	-0.679 ±0.146	-0.660±0.166	-0.610±0.171
GS	-0.705±0.100	-0.517±0.170	-0.517±0.175	-0.407±0.207	-0.073±0.252	-0.005±0.281

than that of the (left) and (center) embeddings, which is captured by LES, but not IMD. Specifically, $LES(\text{left}, \text{center}) = 188.6 \gg LES(\text{left}, \text{right}) = 3.5$; $IMD(\text{left}, \text{center}) = 3.5 < IMD(\text{left}, \text{right}) = 12.6$. LES takes into account cluster structure when comparing datasets, leading to strong correlation between LES-based dissimilarity and few-shot accuracy in our experiments.

LES hyperparameters. In these experiments we set the LES parameters to $K = 500$ and $\gamma = 10^{-6}$. Other values of γ were tested and did not significantly affect the results.

E. Matching Neural Network Layers

In this experiment, we demonstrate the applicability of LES to measuring similarity of NN hidden representations, inspired by Kornblith et al. (2019).

The hidden representations of NN are intrinsically difficult to compare, since they can differ in dimensionality and may be invariant to a variety of transformations, such as permutation of the hidden units. Therefore, the intrinsic properties of the LES distance facilitate the comparison of different layer embeddings.

We train a LeNet convolutional NN on the MNIST training set using two different random seeds (seed 0 and seed 1). The MNIST test set is then passed through the two trained NN and the resulting layer embeddings at the i th layer are saved as $X_i^{(0)}$ and $X_i^{(1)}$, where $i \in \{0, \dots, 4\}$. Here, $i = 0$ denotes the input, and we omit the output of the last layer. We compute the pairwise distances between layers of the two NN, $d(X_i^{(0)}, X_j^{(1)}) \forall i, j \in \{0, \dots, 4\}$, using LES (with $K = 200$ and $\gamma = 10^{-8}$) and IMD. We denote these distance matrices by \mathbf{D}^{LES} and \mathbf{D}^{IMD} . Due to the similar structure of the two NNs, we expect that hidden representations of the layers will be consistent across the random initializations, resulting in a distance matrix with minimal diagonal values. Indeed, \mathbf{D}^{LES} depicts this expected behavior, whereas in \mathbf{D}^{IMD} some of the off-diagonal entries are smaller than the diagonal entries. To demonstrate this different behavior, we define the margin factor α_δ for each layer shift $\delta \in [1, 4]$ by

$$\alpha_\delta = \frac{\min_{j,k \text{ s.t. } |k-j|=\delta} D_{jk}}{\max_i D_{ii}}$$

where D_{jk} denotes the (j, k) th entry in either \mathbf{D}^{LES} or \mathbf{D}^{IMD} , and we set $\alpha_0 = 1$. Namely, we compute the ratio between

Table 4. Few-shot learning: run times (in seconds) for evaluating new task dissimilarity score using $n = 10$ train tasks (360 trials).

	FC100 1-shot	FC100 5-shot	FC100 10-shot	CIFAR-FS 1-shot	CIFAR-FS 5-shot	CIFAR-FS 10-shot
LES	23.3±27.4	20.4±11.8	19.2±11.3	26.2±27.2	20.8±11.6	18.8±10.6
IMD	16.6±11.1	16.1±11.5	24.1±28.6	18.4±11.3	22.9±27.7	16.7±9.9
OT	29.2±10.7	31.0±11.1	40.3±41.3	31.0±10.3	39.3±41.6	28.2±10.3
GW	1569.3±480.4	1591.8±474.5	1716.2±435.7	1458.9±438.9	1539.8±448.6	1579.0±453.5
GS	96.2±7.7	97.2±7.9	99.2±6.5	99.8±7.4	98.3±7.5	97.6±8.6




Figure 8. Projected MetaOptNet FC100 5-shot embeddings on one train (left) and two new tasks (center and right). The corresponding few-shot accuracies are 98.2%, 66.2%, and 96.0%. LES(left, center) = 188.6, LES(left, right) = 0.5; IMD(left, center) = 3.5, IMD(left, right) = 12.6. LES better captures the geometry of the data. Specifically, LES is perceptive of the cluster structure in the embeddings which we hypothesize to be the reason for its success in anticipating difficult tasks in few-shot learning.

the minimal distance of layers shifted by δ and the maximal distance of layers from the same level in the two random initializations. Figure 9 presents the log of the margin factor for LES and IMD as a function of the layer shift, δ . The log-margin factor computed based on the LES distance exhibits the expected behavior (the log-margin factor is greater than 0 ($\log(1)$) for all layer shifts, δ , and increases as $|\delta|$ grows), whereas the IMD plot dips below 0 for $\delta = \pm 1$. Both methods have margin curves that monotonically increase away from $\delta = \pm 1$, indicating that nearby layers are significantly more similar than layers further apart.

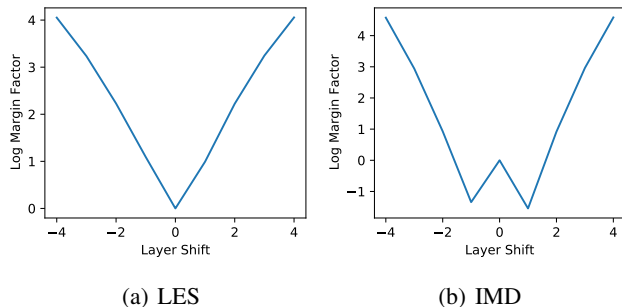


Figure 9. Matching neural network layers for MNIST across random seeds.

To validate these results on a larger scale, we replicate the first experiment in Kornblith et al. (2019). We train 10 instances of an All-CNN-C convolutional NN (Springenberg et al., 2015) with 10 layers (including the logits layer) on CIFAR-10, with different random initializations, reaching average train and test accuracies of 85% and 80%, respectively. We compute the test set embeddings of each layer in each NN instance, and evaluate the accuracy of finding corresponding layers in all other

Table 5. Layer position classification accuracy.

	Identical inputs	Different inputs
LES	96%	95.8%
IMD	85.2%	81%
CKA	97.3%	-

networks based on the minimal distance (maximum similarity) between them. We treat each layer embedding as a separate dataset and compute distances between these datasets using LES, IMD and centered kernel alignment (CKA) (Kornblith et al., 2019), a similarity measure that is closely related to canonical correlation analysis and requires known correspondence between embeddings, i.e. the same input to all NN, unlike LES and IMD. Table 5 presents the layer classification accuracies for the compared distance and similarity measures in two scenarios, “Identical inputs”, in which the layer embeddings of all NN were computed using the same $N = 10000$ test samples, and “Different Inputs”, in which the layer embeddings of each NN were computed based on a random subset of $N \in [7000, 10000]$ test samples. In the latter scenario, the correspondence between the different inputs is unknown, and therefore, CKA cannot be computed. Table 5 depicts that LES is capable of correctly identifying similarities of layer embeddings, obtaining accuracies that are close to CKA even with unknown input correspondence.

These experiments serve as a sanity-check for LES’s applicability to such tasks and demonstrate that LES is successful at ensuring both that corresponding layers are assigned a small distance, and that distances between non-corresponding layers are large. The results highlight the potential of LES for analyzing and understanding different NN architectures, even when trained on different datasets.