



HAL
open science

Riemannian Metric Learning: Closer to You than You Imagine

Samuel Gruffaz, Josua Sassen

► To cite this version:

Samuel Gruffaz, Josua Sassen. Riemannian Metric Learning: Closer to You than You Imagine. 2025. ⟨hal-05111657⟩

HAL Id: hal-05111657

<https://hal.science/hal-05111657v1>

Preprint submitted on 13 Jun 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Riemannian Metric Learning: Closer to You than You Imagine

Samuel Gruffaz ^{*1} and Josua Sassen ^{†1}

¹Université Paris-Saclay, ENS Paris-Saclay, Centre Borelli, F-91190 Gif-sur-Yvette, France.

Abstract

Riemannian metric learning is an emerging field in machine learning, unlocking new ways to encode complex data structures beyond traditional distance metric learning. While classical approaches rely on global distances in Euclidean space, they often fall short in capturing intrinsic data geometry. Enter Riemannian metric learning: a powerful generalization that leverages differential geometry to model the data according to their underlying Riemannian manifold. This approach has demonstrated remarkable success across diverse domains, from causal inference and optimal transport to generative modeling and representation learning. In this review, we bridge the gap between classical metric learning and Riemannian geometry, providing a structured and accessible overview of key methods, applications, and recent advances. We argue that Riemannian metric learning is not merely a technical refinement but a fundamental shift in how we think about data representations. Thus, this review should serve as a valuable resource for researchers and practitioners interested in exploring Riemannian metric learning and convince them that it is closer to them than they might imagine—both in theory and in practice.

1 Introduction

In recent decades, machine learning research has focused on developing vector-based representations for various types of data, including images, text, and time series [22]. Learning a meaningful representation space is a foundational task that accelerates research progress, as exemplified by the success of Large Language Models (LLMs) [182]. A complementary challenge is learning a distance function (defining a metric space) that encodes aspects of the data’s internal structure. This task is known as *distance metric learning*, or simply *metric learning* [20]. Metric learning methods find applications in every field using algorithms relying on a distance such as the ubiquitous k-nearest neighbors classifier: Classification and clustering [195], recommendation systems [89], optimal transport [45], and dimension reduction [116, 186]. However, when using only a global distance, the set of available modeling tools to derive computational algorithms is limited and does not capture the intrinsic data structure.

Hence, in this paper, we present a literature review of *Riemannian metric learning*, a generalization of metric learning that has recently demonstrated success across diverse applications, from causal inference [51, 59, 147] to generative modeling [100, 111, 170]. Unlike metric learning, *Riemannian* metric learning does not merely learn an embedding capturing distance information, but estimates a Riemannian metric characterizing distributions, curvature, and distances in the dataset, *i.e.* the Riemannian structure of the data. This naturally leads to modeling tools inspired by Riemannian geometry, such as geodesics [100, 111, 170] and Ricci curvature [59],

*samuel.gruffaz@ens-paris-saclay.fr

†josua.sassen@ens-paris-saclay.fr

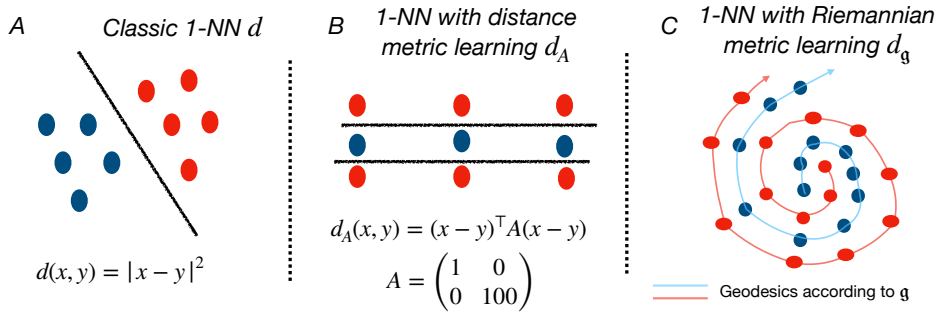


Figure 1: On the left, a classic 1-NN using euclidean distance can separate red and blue points. In the middle, a 1-NN using distance metric learning can separate the points on the anisotropic grid, but not a classic 1-NN. On the right, only a 1-NN using Riemannian metric learning can achieve the separation of spiral data.

further expanding the scope of geometric approaches for addressing modeling and representation learning problems. Thus, we argue that Riemannian metric learning is a fundamental task in representation learning, enriching applications by leveraging geometric theory—much like generative modeling has benefited heavily from probability theory [88].

To recount the story from the beginning, let us begin with the principles of metric learning. Metric learning [20] is an integral, though often implicit, aspect of machine learning. Whether consciously or not, most practitioners in the field are already engaging with it—or will eventually. The concept has a long history, dating back to early methods like the k -nearest neighbors (k -NN) classifier [42, 64, 164]. As the term ‘nearest’ suggests, a measure of distance is at the heart of k -NN algorithms as depicted in Figure 1. Without prior knowledge, the distance is often selected as the Euclidean distance. However, this choice is uninformative of the input and output distribution. That is why Atkeson et al. [10] suggested to learn the distance by learning a symmetric positive definite matrix $G \in \mathbb{R}^{d \times d}$ such that $\mathbf{d}_G(x, y) = \sqrt{(y-x)^\top G (y-x)}$. It turns out that learning $G = L^\top L$ is equivalent to learning a linear transformation $x \mapsto Lx$ of the input space as $\mathbf{d}_G = \mathbf{d}(L \cdot, L \cdot)$. The matrix G can be learned, for example, in an unsupervised manner using Principal Component Analysis (PCA) or in a supervised manner using Linear Discriminant Analysis (LDA). Learning G is valuable for two key reasons: First, it assigns *global relative weights to the different coordinates* when computing distances. Second, it allows us to account for *global correlations* in the dataset. However, G needs to vary to account for *heterogeneous local geometries* within the dataset as shown in Figure 1. This linear framework was then extended with non-linear transformations $x \mapsto \phi(x)$ by using kernels or neural networks [52, 66, 198] such that the distance is $\mathbf{d}(\phi(\cdot), \phi(\cdot))$, which reduces to learning a feature function ϕ as already pointed out.

Going one step further, metric learning can be embedded in a more general framework called Riemannian metric learning [143], in the same way that Galilean relativity is embedded by General relativity [57]. In Riemannian metric learning, the input space $(\mathcal{X}, \mathfrak{g})$ is a Riemannian manifold, *i.e.* a kind of smooth geometric shape (see Figure 2) with a Riemannian metric $\mathfrak{g}: \mathcal{X} \rightarrow \mathbb{R}^{d \times d}$ that allows to locally measure lengths and angles via a matrix field. From this local geometric structure, one can define a global distance between any $x, y \in \mathcal{X}$ as

$$\mathbf{d}(x, y) = \inf_{\gamma \in C^1([0,1], \mathcal{X}), \gamma(0)=x, \gamma(1)=y} \int_0^1 \sqrt{\dot{\gamma}(t)^\top \mathfrak{g}(\gamma(t)) \dot{\gamma}(t)}. \quad (1.1)$$

Minimizing curves of (1.1) are called metric geodesics and can be understood as generalizations of straight lines to manifolds as illustrated in Figure 1. If the Riemannian metric is constant and the manifold is the Euclidean space, then geodesics are straight lines and we recover the distance metric learning framework from above. However, Riemannian metric learning is not just a generalization of distance metric learning with non-linear transformations; it also gives new

mathematical objects useful for modeling. For instance, geodesics can express trajectories for longitudinal modeling and interpolation [100, 109, 123], parallel transport (defined later) can be used in transfer learning [73, 172], and the metric volume can be used for modeling probability distribution and sampling [44, 47]. All these mathematical objects are related to each other and have been extensively studied in Riemannian Geometry [50, 166]. Riemannian metric *learning* leverages the wealth of this literature to offer both flexible and interpretable modeling tools.

As all these mathematical objects depend on the Riemannian metric g , it is worthwhile to learn it by minimizing an objective function. First, as mentioned previously, the metric was chosen constant in an Euclidean space [46, 195, 203] and the objective function asked to reduce the distance between similar points and increase the distance between dissimilar points. The optimization was often convex and was solved using gradient descents, semi-definite programming or Support Vector Machine solvers [19, 185]. Then, people tried to learn metrics on other types of manifolds such as the Grassmanian, the manifold of SPD matrices [208], and the manifold of probability distributions [45]. Furthermore, new parameterizations of the metric have been suggested to handle more complex geometry based on locally constant approximations [148], kernels [5, 72], and neural networks [34, 52, 99]. The optimization is highly non-convex and relies sometimes on sophisticated techniques [23, 44, 170], although often standard accelerated gradient descent techniques prove to be sufficient.

Contributions. There exists a wide range of reviews on *distance* metric learning [19, 52, 66, 108, 169, 186, 198]: Yang and Jin [198] provided the first, then Bellet et al. [19] and Kulis et al. [108] offered an interesting classification of metric learning methods, Wang and Sun [186] focused on dimension reduction methods while Duan et al. [52] addressed deep learning methods, Suárez et al. [169] distinguished themselves from the others by their comprehensive benchmark, and finally Ghojogh et al. [66] exposed the different mathematical formalism exhaustively. However, none of them tackles *Riemannian* metric learning. Recently, Riemannian metric learning has impacted popular research topics from generative modeling [100, 111, 170] to causal inference [51, 59, 147]. It was crucial to improve the state-of-the-art considerably both in performance and interpretability. That is why, considering the diversity of publications in Riemannian metric learning and their rising impact, we dedicate our review to this topic:

- We present Riemannian metric learning as a generalization of distance metric learning and state its general framework.
- We motivate this framework by highlighting the main applications arising from different research communities, demonstrating the wealth and interest of the Riemannian formalism.
- We offer a classification of Riemannian metric learning methods according to their metric parametrizations, related computational methods, objectives, and optimization methods. The key points of distance metric learning will also be recalled in this context to give a complete picture.
- We propose promising theoretical and practical research directions for Riemannian metric learning.

Structure. First, we recall preliminary concepts on the essential tools and properties of Riemannian geometry in Section 2; readers familiar with these notions may safely skip this section. At the same time, we introduce the main manifolds relevant to the data used in Riemannian metric learning in Section 2.1, along with the manifold of Riemannian metrics, to provide a broader perspective on the subject. Second, the formal optimization problem associated with Riemannian metric learning is presented in Section 3.1 to define the topic rigorously. This is followed by an overview of its numerous applications in Section 3.2, highlighting the practical benefits and

key techniques of Riemannian formalism. Third, we discuss how to perform Riemannian metric learning in Section 4, classifying the methods based on their metric parameterization (Section 4.1) and associated computational approaches (Section 4.1.3), their objectives (Section 4.2), and their optimization strategies (Section 4.3). This classification helps clarify the various technical challenges that arise when designing Riemannian metric learning methods. Finally, before concluding in Section 6, we outline key research directions in Riemannian metric learning in Section 5, which may inspire both theorists and practitioners in the applied mathematics community.

Acknowledgments. This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 101034255.

1.1 Notation

We denote by $\mathcal{P}(X)$ the power set of a set X , integer ranges by $[k : l] = \{k, \dots, l\} \subset \mathcal{P}(\mathbb{Z})$ and $[l] = [1 : l]$ with $k, l \in \mathbb{N}$. The canonical scalar product of two vectors $x = (x_i)_{i \in [d]}, y = (y_i)_{i \in [d]} \in \mathbb{R}^d$ is denoted by $x^\top y = \sum_{i=1}^d x_i y_i$. The Euclidean norm is denoted by $\|\cdot\| = \sqrt{\cdot^\top \cdot}$. Denote by $C^k(\mathbb{U}, \mathbb{R}^d)$ the set of all k times continuously differentiable functions from \mathbb{U} to \mathbb{R}^d . When $\mathbb{U} \subset \mathbb{R}$, we denote by $\dot{\gamma}$ the time derivative of $\gamma \in C^1(\mathbb{U}, \mathbb{R}^d)$. The space of square symmetric and symmetric positive definite matrices are denoted by $\mathcal{S}^d = \{A \in \mathbb{R}^{d \times d} : A^\top = A\}$ and $\mathcal{S}_{++}^d = \{A \in \mathcal{S}^d : x^\top A x > 0 \text{ for any } x \in \mathbb{R}^d \setminus \{0\}\}$ respectively. The square identity matrix in dimension $d \times d$ is denoted by I_d . The determinant and the trace of a matrix $A \in \mathbb{R}^{d \times d}$ are denoted by $\text{tr}(A)$ and $\det(A)$ respectively. For any $(x_i)_{i \in [m]} \in \mathbb{R}^m$, $\text{diag}((x_i)_{i \in [m]})$ denotes the diagonal matrix whose diagonal elements are $(x_i)_{i \in [m]}$, arranged from the top-left. For any $f \in C^1(\mathbb{U}, \mathbb{R}^m)$ with $\mathbb{U} \subset \mathbb{R}^d$, df is the differential of f . For any $f \in C^1(\mathbb{U}, \mathbb{R})$ with $\mathbb{U} \subset \mathbb{R}^d$, ∇f is the gradient of f , but can also denote the covariate derivative depending on the context. For any $j \in [d]$ and any $\psi \in C^1(\mathbb{U}, \mathbb{R}^m)$ with $\mathbb{U} \subset \mathbb{R}^d$, $\partial_j \psi \in \mathbb{R}^m$ denotes the directional derivative of ψ in direction e_j , where e_j is the j -th element of the canonical basis. When a random variable $X \in E$ follow a probability distribution P on a measurable space (E, \mathcal{B}) , we write $X \sim P$. $\mathcal{U}([0, 1])$ denotes the uniform probability distribution on $[0, 1]$. If P_1 and P_2 are two probability distribution on (E, \mathcal{B}_1) and (F, \mathcal{B}_2) respectively, then $P_1 \otimes P_2$ their product measure.

2 A Primer on Riemannian Geometry

In metric learning, the core element is naturally the metric — a measure of distance within the data space \mathcal{X} , which can be tailored to various objectives. We begin by introducing the formal definition of a metric space:

Definition 1 (Metric space). A metric space consists of a set \mathcal{X} and a function $\mathbf{d} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ called the metric, such that

1. $\mathbf{d}(x, y) = 0$ if and only if $x = y$,
2. $\mathbf{d}(x, y) = \mathbf{d}(y, x)$, and
3. $\mathbf{d}(x, z) \leq \mathbf{d}(x, y) + \mathbf{d}(y, z)$.

On Euclidean spaces, i.e. $\mathcal{X} = \mathbb{R}^d$, the standard metric is the *Euclidean distance*, given by $\mathbf{d}(x, y) = \|x - y\|_{\mathbb{R}^d} = \sqrt{\sum_{i=1}^d (x_i - y_i)^2}$. However, depending on the application, alternative metrics such as the *Manhattan distance*, defined as $\mathbf{d}(x, y) = \sum_{i=1}^d |x_i - y_i|$, are also employed.

Given some objective function, the metric also can be defined as the solution of an optimization problem, parametrizing the metric by $G \in \mathcal{S}_{++}^d$ as $\mathbf{d}_G(x, y) = \sqrt{(x - y)^\top G (x - y)}$ for vectors $x, y \in \mathbb{R}^d$. This specific parametric form has been the dominant approach in most works on *distance metric learning*. However, such metrics are constant across the entire space and may fail




Figure 2: Illustration of tangent spaces, geodesics and the Exponential map. The curve $\gamma_{x,v}$ is the geodesic starting from $x \in \mathcal{M}$ with velocity $v \in T_x \mathcal{M}$.

to capture the inherent nonlinear structure of the data, similar to how linear PCA may not provide a meaningful approximation in the presence of complex data structures. This limitation has motivated researchers to explore manifold learning methods.

In this work, we aim to approach metric learning from a Riemannian perspective. Therefore, we will briefly introduce key concepts from *Riemannian geometry* necessary to understand the remainder of this review. This section is essentially pedagogical and readers already familiar in Riemannian geometry can safely skip this part. For a more comprehensive treatment, we refer the reader to standard textbooks such as [112].

Manifolds. A *differentiable manifold* $\mathcal{M} \subset \mathbb{R}^{d_a}$ is formally defined in Appendix A or in [112, Chapter 2], where d_a denotes the dimension of the ambient space. For a first reading, and to simplify notation¹, a *differentiable manifold* of dimension $d < \infty$ can be understood as the image of a global, surjective, and smooth parametrization $\psi: U \subset \mathbb{R}^d \rightarrow \mathcal{M}$, where U is an open subset of \mathbb{R}^d . In general, we would need more than one parametrization and all concepts presented hereafter generalize to this scenario. Readers unfamiliar with this concept may safely think of a differentiable manifold as a smoothly curved, nonlinear object, such as a curved surface in three-dimensional space as depicted in Figure 2. Note that in all the following a function $f: \mathcal{M} \rightarrow \mathcal{Y} \subset \mathbb{R}^m$ is called smooth or infinitely differentiable if $f \circ \psi: U \rightarrow \mathcal{Y}$ is so, which offers a valid definition since U is an open set of \mathbb{R}^d .

With this understanding, we can examine a local linearization of this curved space to apply tools from Euclidean vector spaces, which leads to the concept of a tangent space:

Definition 2 (Tangent space). *The tangent space $T_p \mathcal{M}$ of \mathcal{M} at $p \in \mathcal{M}$ is defined as*

$$T_p \mathcal{M} := \{\dot{\gamma}(0) \mid \gamma \in C^1(-\varepsilon, \varepsilon, \mathcal{M}), \gamma(0) = p, \varepsilon > 0\}, \quad (2.1)$$

which is a vector space of dimension d . If $\psi: U \subset \mathbb{R}^d \rightarrow \mathcal{M}$ is a parametrization with $\psi(\xi) = p$ for some $\xi \in U$, then $\partial_{\xi_i} \psi$ for $i = 1, \dots, d$ is a basis of $T_p \mathcal{M}$, called the canonical basis.

The dimension of the manifold is thus given by the dimension of possible local displacements in each point. In the case of two-dimensional surfaces, the tangent space is simply the plane that is tangent to the surface at a given point as illustrated in Figure 2.

Example 3. *The Euclidean space \mathbb{R}^d is a manifold of dimension d . For any $x \in \mathbb{R}^d$, the tangent space is isomorphic to the Euclidean space $T_x \mathbb{R}^d \simeq \mathbb{R}^d$ itself. Despite its simplicity, $\mathcal{M} = \mathbb{R}^d$ is still interesting in many applications when endowed with a spatially varying Riemannian metric.*

A non-trivial manifold prevalent in statistics is the manifold of symmetric positive definite (SPD) matrices

$$\mathcal{S}_{++}^d := \left\{ M \in \mathbb{R}^{d \times d} \mid M = M^T \text{ and } x^T M x > 0 \text{ for any } x \in \mathbb{R}^d \setminus \{0\} \right\}. \quad (2.2)$$

¹All concepts introduced below are compatible with changes of parametrization in overlapping regions and can therefore be generalized to manifolds requiring multiple parametrizations.

For any $M \in \mathcal{S}_{++}^d$, the tangent space $T_M \mathcal{S}_{++}^d$ is isomorphic to the space of symmetric matrices \mathcal{S}^d . Thus, \mathcal{S}_{++}^d is a manifold of dimension $d(d+1)/2$.

More manifold examples and insights related to computations and modeling will be given in Section 2.1.

Riemannian Structure. Next, we aim to measure distances on our curved space, which requires introducing additional geometric structure. Specifically, we will define this structure locally on tangent spaces and then use it to derive global notions of distance and related concepts across the entire manifold. As previously discussed, a scalar product allows us to measure lengths and angles. Therefore, we begin by defining a scalar product for each tangent space:

Definition 4 (Riemannian manifold). *Let \mathcal{M} be a d -dimensional differentiable manifold. A Riemannian metric on \mathcal{M} is a family of bilinear, symmetric and positive-definite forms $\mathfrak{g}_p: T_p \mathcal{M} \times T_p \mathcal{M} \rightarrow \mathbb{R}$ smoothly varying with $p \in \mathcal{M}$, in the sense that for a parametrization $\psi: U \subset \mathbb{R}^d \rightarrow \mathcal{M}$ the map $\xi \mapsto \mathfrak{g}_{ij}(\xi) := \mathfrak{g}_{\psi(\xi)}(\partial_i \psi(\xi), \partial_j \psi(\xi))$ is a smooth function on U . A manifold equipped with a Riemannian metric is called a Riemannian manifold.*

We can also interpret this inner product as a smooth matrix field $\mathfrak{g}: \mathcal{M} \rightarrow \mathbb{R}^{d \times d}$ on the manifold by selecting a basis for each tangent space. In the case that a global parametrization of the manifold exists, this corresponds to a smooth matrix field $\mathfrak{g}: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ defined on the parameter domain, representing a locally varying notion of distance. Since $(\mathfrak{g}_{ij})_{ij}$ is an invertible matrix in $\mathbb{R}^{d \times d}$, it has an inverse $\mathfrak{g}^{-1} \in \mathbb{R}^{d \times d}$, which we denote as $(\mathfrak{g}^{kl})_{kl}$, satisfying $\mathfrak{g}_{ij} \mathfrak{g}^{jk} = \delta_{ik}$.

Example 5. *The Euclidean scalar product on \mathcal{S}_{++}^d , i.e. $\mathfrak{g}_M(W, V) = \text{tr}(W^\top V)$ for any $M \in \mathcal{S}_{++}^d$ and $W, V \in \mathcal{S}^d$, yields a Riemannian metric. However, this does not penalize displacements toward symmetric matrices with nonpositive eigenvalues. Hence, it does not take the structure of the manifold into account. An alternative choice of Riemannian metric penalizing singularities is the affine-invariant metric $\mathfrak{g}_M(W, V) = \text{tr}(M^{-1} W M^{-1} V)$ which implies that $\mathfrak{g}_M(V, V)$ diverges towards infinity when M converges to the boundary of \mathcal{S}_{++}^d .*

Distance. Returning to our objective of learning a measure of distance, we discuss now how such metric gives rise to a Riemannian distance $\mathbf{d}_{\mathfrak{g}}$ on the entire manifold. First, one can define the length of a smooth path $\gamma: [0, 1] \rightarrow \mathcal{M}$ on a Riemannian manifold $(\mathcal{M}, \mathfrak{g})$ in terms of the metric as

$$\mathcal{L}(\gamma) := \int_0^1 \sqrt{\mathfrak{g}_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt. \quad (2.3)$$

This path length is invariant to reparametrization of the curve [112, Lemma 6.1]. Furthermore, the path energy is defined as

$$\mathcal{E}(\gamma) := \int_0^1 \mathfrak{g}_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t)) dt, \quad (2.4)$$

which is not independent of the parametrization. By the Cauchy-Schwarz inequality, one directly sees

$$\mathcal{L}(\gamma) \leq \sqrt{\mathcal{E}(\gamma)} \quad (2.5)$$

and equality holds if and only if $\mathfrak{g}_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))$ is constant along the path, which is the case for *geodesics* as we will see later. Given the length of paths, we can define the distance between two arbitrary points as the shortest possible length of a path connecting them, which leads to the definition of a *Riemannian distance* [112, Lemma 6.2]:

Definition 6 (Riemannian distance). *Let $(\mathcal{M}, \mathfrak{g})$ be a Riemannian manifold and $x, y \in \mathcal{M}$ two points on it. Then the Riemannian distance between x and y is defined as*

$$\mathbf{d}_{\mathfrak{g}}(x, y) := \inf_{\substack{\gamma \in C^1([0,1], \mathcal{M}) \\ \gamma(0)=x, \gamma(1)=y}} \mathcal{L}(\gamma). \quad (2.6)$$




Figure 3: Effect of the metric on the shape of the minimizing curve in (2.6). The effort required for displacement around a point can be modeled in physics as an energy input, which is defined as the metric norm of the local velocity. When retrieving your lost ball from the sea, following a straight trajectory wastes energy for nothing.

Shortest paths can have different shapes depending on the metric as illustrated in Figure 3.

Example 7. When $\mathcal{M} = \mathcal{S}_{++}^d$ and \mathfrak{g} is the Euclidean scalar product, then $\mathbf{d}_{\mathfrak{g}}$ is the usual Frobenius distance. When \mathfrak{g} is the affine-invariant metric defined in Example 5, the Riemannian distance also has a closed form expression $\mathbf{d}_{\mathfrak{g}}(M_1, M_2) = \left\| \log \left(M_1^{-\frac{1}{2}} M_2 M_1^{-\frac{1}{2}} \right) \right\|$ with \log being the matrix logarithm and $\|\cdot\|$ an Euclidean metric on symmetric matrices. Note that $\mathbf{d}_{\mathfrak{g}}(I, \Sigma)^2 = \sum_{i=1}^d \log(\sigma_i)^2$ with (σ_i) the eigenvalues of Σ and thus $\mathbf{d}_{\mathfrak{g}}(I, \Sigma)^2$ diverges towards the infinity as Σ becomes singular.

Reformulating the classical distance metric learning framework with these objects amounts to learning a constant metric \mathfrak{g} defined on a manifold, which corresponds to Euclidean space. In contrast, *Riemannian* metric learning provides a more general framework, accommodating curved spaces and metrics that vary based on position within the space. By learning a Riemannian metric, one gains access to additional geometric objects derived from the metric, which we introduce next.

Remark 8. Note that not all distances on manifolds are Riemannian distances. For instance, for any injective function $f: \mathcal{M} \rightarrow \mathcal{Y}$, $\mathbf{d}_f(\cdot, \cdot) \triangleq \|f(\cdot) - f(\cdot)\|$ is a valid distance on \mathcal{M} . However, a metric \mathfrak{g} such that $\mathbf{d}_f = \mathbf{d}_{\mathfrak{g}}$ does not exist necessarily. For example, if $f = \text{Id}$ maps the unit circle \mathcal{S}^1 to \mathbb{R}^2 , the induced distance \mathbf{d}_f will approximate the intrinsic distance on the circle only when x and y are sufficiently close. For distant points on the circle, the map-induced distance will underestimate the true geodesic distance.

However, one can derive a Riemannian metric from a given distance on the manifold by considering its Hessian. The resulting Riemannian metric gives rise to a Riemannian distance that does not necessarily agree with the initially provided one. In the example of the circle above, the Hessian would yield to the usual Euclidean scalar product on the tangent space and thus the usual distance on the circle. The initially provided distance, however, does yield a local approximation of the Riemannian one and this observation was used by Rumpf and Wirth [150] to derive time-discretization of Riemannian calculus.

In order to find optimality constraints related to (2.6), we need to compute variations of \mathcal{L} with respect to γ which thus implies second order derivatives $\ddot{\gamma}$. However, defining second order derivatives via a difference quotient $\lim_{dt \rightarrow 0} (\dot{\gamma}(t+dt) - \dot{\gamma}(t))/dt$ is not straightforward on manifolds. The vectors $\dot{\gamma}(t+dt) \in T_{\gamma(t+dt)}\mathcal{M}$ and $\dot{\gamma}(t) \in T_{\gamma(t)}\mathcal{M}$ belong to two different spaces as depicted in Figure 2 and thus cannot be subtracted a priori. Thus, we introduce a way to do so next.

Covariant Derivative ∇ . The covariant derivative ∇ is a generalization of the notion of second order derivative and applies to smooth *vector field*. A smooth vector field V is a map that assigns

to every point on the manifold a point in its tangent space, *i.e.* $V: \mathcal{M} \rightarrow T\mathcal{M}$ such that $V(x) \in T_x\mathcal{M}$ for any $x \in \mathcal{M}$ and which is infinitely differentiable. The space of smooth vector fields is denoted by $C^\infty(\mathcal{M}, T\mathcal{M})$. The partial derivatives of the parametrization can be considered as vector fields $\partial_i\psi: \mathcal{M} \rightarrow T\mathcal{M}$ that form a basis of the tangent space. Hence, any vector field V can be written as $V(x) = \sum_{i=1}^d V_i(x)\partial_i\psi(x)$ for any $x \in \mathbb{R}^d$ where $V_i: \mathcal{M} \rightarrow \mathbb{R}$ are the coordinates in the basis $(\partial_i)_{i \in [d]}$. If we aim to differentiate a vector field V , we will end up differentiating $\partial_i\psi$ and thus ψ will be differentiated twice.

We define a covariant derivative ∇ as a map $C^\infty(\mathcal{M}, T\mathcal{M}) \times C^\infty(\mathcal{M}, T\mathcal{M}) \rightarrow C^\infty(\mathcal{M}, T\mathcal{M})$ which satisfies intuitive derivation rules as exposed in the following example or in Definition 18 given in Appendix A or in [112, Chapter 4, p.50]. Intuitively, $\nabla_W V$ describes how the vector field V changes when moving in direction W . In this way, a covariant derivative ∇ of a vector field $V = \sum_{i=1}^d V_i\partial_i\psi$ in direction of $W = \sum_{i=1}^d W_i\partial_i\psi$ denoted by $\nabla_W V$ can be decomposed as

$$\nabla_W V = \sum_{i=1}^d \nabla_W (V_i\partial_i\psi), \quad (\mathbb{R} - \text{linearity}) \quad (2.7)$$

$$= \sum_{i=1}^d \nabla_{\sum_{k=1}^d W_k\partial_k\psi} (V_i\partial_i\psi) \quad (2.8)$$

$$= \sum_{i=1}^d \sum_{k=1}^d W_k \nabla_{\partial_k\psi} (V_i\partial_i\psi), \quad (C^1 - \text{linearity}) \quad (2.9)$$

$$= \sum_{i=1}^d \sum_{k=1}^d W_k [\nabla_{\partial_k\psi} (V_i)\partial_i\psi + V_i \nabla_{\partial_k\psi} (\partial_i\psi)], \quad (\text{Product derivation rule}) \quad (2.10)$$

where $\mathbb{R} - \text{linearity}$, $C^1 - \text{linearity}$ and the product derivation rule can be used because of the definition of ∇ . Thus, to define a covariant derivative, everything reduce to define $\nabla_{\partial_k\psi} (V_i)$ and $\nabla_{\partial_k\psi} (\partial_i\psi)$ for any $k, i \in [d]$. Since V_i is a scalar function and to be consistent with the usual notion of directional derivative, $\nabla_{\partial_k\psi} (V_i)$ is defined as $\nabla V_i^\top \partial_k\psi$ where ∇V_i is the usual gradient. However, there is a priori not a unique way to define $\nabla_{\partial_k\psi} (\partial_i\psi)(x) \in T_x\mathcal{M}$ for any $k, i \in [d]$ and $x \in \mathcal{M}$. The intuition would be to set $\nabla_{\partial_k\psi} (\partial_i\psi)(x)$ as the orthogonal projection of $\partial_k\partial_i\psi(x)$ on $T_x\mathcal{M}$ for any $x \in \mathcal{M}$. This is valid only when the metric \mathfrak{g}_x in $T_x\mathcal{M}$ is consistent with the Euclidean metric of the ambient space \mathbb{R}^{d_a} , *i.e.* \mathfrak{g}_x is the restriction of the ambient scalar product to $T_x\mathcal{M}$. For more general metric \mathfrak{g} , we define the covariant derivative using the basis decomposition for any $i, j \in [d]$,

$$\nabla_{\partial_i\psi} \partial_j\psi = \sum_{k=1}^d \Gamma_{ij}^k \partial_k\psi \quad (2.11)$$

where the coordinates are called the *Christoffel symbols* $(\Gamma_{ij}^k: \mathcal{M} \rightarrow \mathbb{R})_{(i,j,k) \in [d]^3}$ which are smooth scalar functions. There exists a unique way to define the Christoffel symbols to have symmetry constraints² $\Gamma_{ij}^k = \Gamma_{ji}^k$ and a bilinear rule of derivation according to the metric \mathfrak{g} , *i.e.*

$$\nabla_Z \mathfrak{g}(V, W)(x) = \mathfrak{g}_x(\nabla_Z V(x), W(x)) + \mathfrak{g}_x(V(x), \nabla_Z W(x)) \quad (2.12)$$

for any $x \in \mathcal{M}$ and $Z, V, W \in C^\infty(\mathcal{M})$. This defines the Levi-Cevita covariant derivative ∇ also called the Levi-Cevita connection [112, Theorem 5.4]. The following property gives a concrete expression.

Proposition 9. *Let $\mathfrak{g}^{-1} = (\mathfrak{g}^{ij})_{(i,j) \in [d]^2}$ be the inverse of \mathfrak{g} as before. Then, the Levi-Cevita covariant derivative has the following representation of the Christoffel symbols:*

$$\Gamma_{ij}^k = \frac{1}{2} \sum_{l=1}^d \mathfrak{g}^{lk} (\partial_i \mathfrak{g}_{jl} - \partial_l \mathfrak{g}_{ij} + \partial_j \mathfrak{g}_{li}). \quad (2.13)$$

²Also called torsion-free.




Figure 4: Illustration of the parallel transport along the curve γ_{x,v_1} (already represented in Figure 2).

The Levi-Cevita connection ∇ is particularly suitable since we will see below that locally minimizing curves γ of the energy (2.6) verify $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$ [112, Corollary 6.7]. In other words, a local minimizing curve γ has a vanishing second order derivative according to the geometry induced by the metric g . See Appendix A for an infinite dimensional extension of the covariant derivative which is relevant to analyze diffeomorphism groups or Wasserstein space geometry.

Parallel Transport and Geodesics. With the covariant derivative, we can introduce the notion of parallel vector fields and the related concept of parallel transport [112, Chapter 4, p.60].

Proposition 10 (Parallel transport). *Let $\gamma: I \rightarrow \mathcal{M}$ be a curve. A vector field $V: I \rightarrow T\mathcal{M}$ along γ is called parallel if $\nabla_{\dot{\gamma}}V(t) = 0$ for all $t \in I$. For $t_0 \in I$, $V_0 \in T_{\gamma(t_0)}\mathcal{M}$, there exists a unique parallel vector field $V: I \rightarrow T\mathcal{M}$ with $V(t_0) = V_0$ following for any $k \in [d]$ the equation*

$$\dot{V}_k + \sum_{i,j=1}^d \Gamma_{i,j}^k \dot{\gamma}_i V_j = 0. \quad (2.14)$$

Furthermore, the map $\mathfrak{T}_{\gamma(t_0) \rightarrow \gamma(t)}: T_{\gamma(t_0)}\mathcal{M} \rightarrow T_{\gamma(t)}\mathcal{M}$, $V_0 \mapsto V(t)$ is a linear isomorphism.

The parallel transport of a vector along a curve is depicted in Figure 4. Finally, the concept of ‘straight lines’ in Euclidean spaces can be generalized to manifolds using *geodesics*. A geodesic γ transports its own velocity vector in a parallel manner [112, Chapter 4, p.58].

Definition 11 (Geodesic). *A curve $\gamma: I \rightarrow \mathcal{M}$ is called geodesic if it solves the geodesic equation, i.e. the second order ordinary differential equation*

$$\nabla_{\dot{\gamma}}\dot{\gamma}(t) = 0 \quad (2.15)$$

for any $t \in I$.

Intuitively, this means that a geodesic experiences no acceleration along the direction of the manifold, and its curvature is entirely determined by the curvature of the manifold itself. This property underscores why geodesics are regarded as generalizations of straight lines.

Example 12. *When using the affine invariant metric on S_{++}^d , the geodesic $\gamma_{V,\Sigma}$ in direction $V \in S^d$ with starting point $\Sigma \in S_{++}^d$ is $t \mapsto \Sigma^{1/2} \exp(t\Sigma^{-1/2}V\Sigma^{-1/2})\Sigma^{1/2}$. In case of $\Sigma = \text{Id}$, it reduces to the evaluation of the matrix exponential [144]. The parallel transport of a vector $W \in T_{\Sigma}S^d$ at time $s = 0$ to time t along the geodesic $\gamma_{V,\Sigma}$ is $\mathfrak{T}_{\gamma_{V,\Sigma}(0) \rightarrow \gamma_{V,\Sigma}(t)}W = \exp(tV\Sigma^{-1}/2)W \exp(t\Sigma^{-1}V/2)$ [159, Lemma 3].*

Remark 13 (Variational definition). *A common way to explain geodesics is to describe them as “shortest paths”. Here, we aim to make this connection between geodesics and distances more precise. In fact, every geodesic is a local minimizer of the path length \mathcal{L} . However, it is not necessarily a global minimizer. A well-known example of this can be found on a sphere, where geodesics correspond to great circles. For any two points on the sphere, there is always a great circle passing*

through them. The two resulting arcs between the points are both geodesics, but only the shorter of the two arcs represents a global minimizer of the path length, providing the distance between the points. Conversely, not every minimizer of the path length is a geodesic, as the curve may not be parametrized by arc length as the path length is invariant under reparametrization. Therefore, minimizers of the path length that satisfy the constraint of being parametrized by arc length are indeed geodesics. Moreover, minimizers of the path energy \mathcal{E} are always geodesics. Since path energy is not invariant under reparametrization, minimizing the path energy inherently ensures that the path is parametrized by arc length. Additionally, as noted earlier, minimizers of the path energy also minimize the path length, reinforcing the connection between the two. This has been exploited to define a variational time-discretization often used for shape spaces [79, 83, 150].

Remark 14 (Hamiltonian flow). Another way to characterize geodesics is through techniques from geometric mechanics. The second-order geodesic ODE (2.15) can be reformulated as a Hamiltonian system consisting of two first-order ODEs. To achieve this, we introduce the Hamiltonian for any $x, p \in \mathcal{M} \times T\mathcal{M}$,

$$H(x, p) = \frac{1}{2} p^T \mathfrak{g}_x^{-1} p \quad (2.16)$$

and the geodesic equation can be written as the corresponding Hamiltonian flow determined by

$$\frac{d}{dt} x(t) = \frac{\partial}{\partial p} H(x(t), p(t)), \quad \frac{d}{dt} p(t) = -\frac{\partial}{\partial x} H(x(t), p(t)). \quad (2.17)$$

From this system, the geodesic ODE (2.15) is recovered through substitution. This approach to the geodesic ODE has been widely used in the shape space literature using the flow of diffeomorphisms [69]. In our context, it opens up another approach to parametrize the metric structure on the manifold by parametrizing the corresponding Hamiltonian H . This is related to other problems in physics described via Hamiltonian systems [40].

Based on any of these definitions, one can introduce the exponential map [112, Proposition 5.7], which “shoots” geodesics in prescribed directions as illustrated in Figure 2.

Definition 15 (Exponential map). Let $\gamma: [0, 1] \rightarrow \mathcal{M}$, be the solution of $\nabla_{\dot{\gamma}(t)} \dot{\gamma}(t) = 0$ for initial data $\gamma(0) = p$ and $\dot{\gamma}(0) = V$ if it exists. The Riemannian exponential map $\text{Exp}_p: T_p\mathcal{M} \rightarrow \mathcal{M}$ is defined as $\text{Exp}_p(V) = \gamma(1)$.

This map is called the exponential map because of its link with the exponential on matrices when using the affine-invariant metric on \mathcal{S}_{++}^d as detailed in Example 12.

In general, one can show that the exponential map is locally bijective, *i.e.* there exists an $\eta > 0$, such that $\text{Exp}_p: B_\eta(0) \rightarrow \text{Exp}_p(B_\eta(0))$ is a bijection denoting by $B_\eta(0)$ the centered open ball of radius η . This follows from the local uniqueness of solutions of the geodesic equation. The image $U_p := \text{Exp}_p(B_\eta(0))$ is called a *normal neighborhood* [112, Lemma 5.10] of p and, in this neighborhood, we can define the Riemannian logarithm as inverse of the exponential map:

Definition 16 (Logarithm). The inverse operator of the exponential map is called the Riemannian logarithm $\log_p: U_p \rightarrow T_p\mathcal{M}$, where U_p denotes the normal neighborhood of p .

See Figure 5 for an illustration and Figure 6 for an application.

Integration. Finally, to consider probability distributions and statistics on Riemannian manifold, we discuss how to integrate on them [151, Section II.5]. A metric \mathfrak{g} introduces a natural volume form $\text{dvol}_{\mathfrak{g}}$ on the manifold \mathcal{M} through its determinant. That is, for a smooth function $f: \mathcal{M} \rightarrow \mathbb{R}$, we obtain the integral

$$\int_{\mathcal{M}} f(x) \text{dvol}_{\mathfrak{g}}(x) := \int_U f(\psi(\xi)) \sqrt{|\det \mathfrak{g}_{\psi(\xi)}|} \text{d}\xi, \quad (2.18)$$

where the manifold is parametrized using $\psi: U \subset \mathbb{R}^d \rightarrow \mathcal{M}$. In case multiple functions are necessary to parametrize \mathcal{M} , one defines a partition of unity [151, Section I.2.1 and p.62] corresponding to the cover given by the images of the parametrizations and uses it to define the integral by combining the integrals defined for each image as above.

2.1 Some Typical Manifolds

Now that we have introduced the structure of Riemannian manifolds, we will take a look at some typical examples of such manifolds and how they come up in the metric learning context. Progressing from more common and intuitive manifolds to increasingly technical and less familiar ones, we conclude with a general discussion on tractability issues and the benefits of employing such a geometric framework. Finally, the manifold of Riemannian metrics is introduced as a complex yet promising structure—one that must be tamed for the advancement of Riemannian metric learning.

Euclidean Space. Our first example is the most elementary one and, at the same time, by far the most prevalent: the Euclidean space \mathbb{R}^d . It fulfills all of our requirements for a (Riemannian) manifold using the identity as parametrization. Nevertheless, one can put a variety of metrics on the Euclidean space, which make it an interesting space to study in metric learning. Furthermore, it can be used to model a wide range of data, where typical examples in our context include images [2, 34, 35, 175], Electroencephalography (EEG) [5], vector representation of text [204], and tabular data such as financial data or sensor measurements in engineering or natural sciences [14]. All the following examples could also be considered in Euclidean space, however, depending on the application we might have additional structure useful to exploit.

Symmetric Positive Definite Matrices. When studying manifolds in a differential geometry context, the first examples beyond the Euclidean space are usually curved geometric objects such as spheres or tori. However, when working with data, matrix manifolds, *i.e.* sub-manifolds of matrices $\mathbb{R}^{d \times d}$, are more common and thus our next examples. The first such is the manifold of symmetric positive definite (SPD) matrices (2.2). It forms an open submanifold of the space of matrices and, more precisely, is an open convex cone, *i.e.* it is closed under linear interpolation and multiplication with positive scalars. Symmetric positive definite matrices come up regularly as covariance matrices, e.g. in the analysis of images [94, 97, 183, 201, 208] and EEG data [41], as well as for action recognition from motion sensors [77]. Furthermore, fields of SPD matrices appear as data in diffusion tensor imaging [144, 159]. There are three commonly used Riemannian structures on the SPD manifold: (i) *Euclidean* structures as, for examples, induced by the Frobenius distance $\mathbf{d}(M_1, M_2) := \|M_1 - M_2\|_F$, (ii) *affine-invariant* structures induced by $\mathbf{d}(M_1, M_2) := \left\| \log \left(M_1^{-\frac{1}{2}} M_2 M_1^{-\frac{1}{2}} \right) \right\|$ with \log being the matrix logarithm and $\|\cdot\|$ an Euclidean metric on symmetric matrices, and ultimately (iii) *Log-Euclidean* structures induced by $\mathbf{d}(M_1, M_2) := \|\log M_1 - \log M_2\|$ using the same elements [4]. We repeat here some basic facts about these structures and refer interested reader, for example, to [144] for more details.

The Euclidean structures are particularly easy to work with computationally, but they largely disregard the specific structure of SPD matrices, e.g. “it is hard to prevent nonpositive eigenvalues using it. In contrast, affine-invariant structures push symmetric matrices with nonpositive eigenvalues at an infinite distance, but incur high computational costs as it typically leads to an extensive use of matrix inverses, square roots, logarithms, and exponentials. Log-Euclidean metrics were introduced to achieve a compromise between these two properties. Through parametrizing \mathcal{S}_{++}^n over the Euclidean space \mathcal{S}^n of symmetric matrices computations become simply Euclidean in the logarithmic domain, limiting the number of necessary evaluations of the matrix logarithm and exponential yet the results are similar to the affine-invariant case [4]. Furthermore, the Log-Euclidean setting yields a Lie group structure in SPD matrices, *i.e.* turns it into a space that is both a manifold and a group where multiplication and inversion are smooth mappings. We will later discuss in Section 4.1 how to parametrize learnable metrics on SPD manifolds using these various structures.

The Special Orthogonal Group. Next, we turn to more geometric matrix manifolds. The first of which is the *special orthogonal group*

$$SO(n) := \{Q \in \mathbb{R}^{n \times n} \mid Q^T Q = \text{Id}_n, \det Q = 1\}, \quad (2.19)$$

which is a Lie group, *i.e.* also equipped with a smooth manifold structure. It is one of the two connected components of the *orthogonal group* $O(n) := \{Q \in \mathbb{R}^{n \times n} \mid Q^T Q = \text{Id}_n\}$ which consists of all distance-preserving linear transformations of \mathbb{R}^n . The tangent space of $SO(n)$ (and also $O(n)$) at the identity is the Lie algebra of skew-symmetric matrices $T_{\text{Id}}SO(n) = \mathfrak{so}(n) := \{A \in \mathbb{R}^{n \times n} \mid A^T = -A\}$ and the tangent space at an arbitrary element Q is given by the transformation of this space under Q , *i.e.* $T_Q SO(n) = Q \cdot T_{\text{Id}}SO(n) := \{QA \mid A \in \mathfrak{so}(n)\}$. Since these tangent spaces are also matrix spaces, to obtain a Riemannian metric one uses the Euclidean metric on matrices and obtains $\mathfrak{g}_Q(QA, QB) = \mathfrak{g}_{\text{Id}}(A, B) = \frac{1}{2} \text{tr} A^T B$ ³. In dimensions two and three, $SO(n)$ consists of rotations around points or lines respectively and is therefore also called the *rotation group*. The (special) *Euclidean group* consists of its extension to all distance-preserving affine transformations, *i.e.* also considering translations: $SE(n) := SO(n) \times \mathbb{R}^n$. Concerning applications, products of $SO(n)$ or $SE(n)$ are commonly used to model the rotation of joints, e.g. of skeleton models or rigs in computer animation [109] or of robotic arms [16, 32].

The Stiefel Manifold And The Grassmannian. Our next geometric matrix manifold is the *Stiefel manifold* [54]

$$\text{St}_k(\mathbb{R}^n) := \left\{ M \in \mathbb{R}^{n \times k} \mid M^T M = \text{Id}_k \right\}, \quad (2.20)$$

which consists of orthonormal bases of k -dimensional subspaces of \mathbb{R}^n . It is a compact submanifold of the space of matrices with dimension $nk - \frac{1}{2}k(k+1)$ and, furthermore, can be constructed also as a quotient manifold by observing $\text{St}_k(\mathbb{R}^n) \cong O(n)/O(n-k)$. Its tangent space is given by

$$T_M \text{St}_k(\mathbb{R}^n) \cong \left\{ V \in \mathbb{R}^{n \times k} \mid V^T M = -M^T V \right\} \quad (2.21)$$

and via the quotient manifold construction one obtains the so-called *canonical* metric on it as $\mathfrak{g}_M(V, W) = \text{tr} V^T (\text{Id} - \frac{1}{2} M M^T) W$. There is also the *Euclidean* metric on the Stiefel manifold which is obtained by restricting the ambient scalar product on matrices to the tangent space, *i.e.* $\mathfrak{g}_M(V, W) = \text{tr} V^T W$. The relationship between the induced geodesic distances of these two metrics was recently investigated by Mataire et al. [130].

In the metric learning context, the Stiefel manifold is typically used to construct the *Grassmannian* $\text{Gr}_k(\mathbb{R}^n)$, which disregards the concrete basis and only considers the linear subspaces, *i.e.*

$$\text{Gr}_k(\mathbb{R}^n) := \{ \mathcal{U} \mid \mathcal{U} \text{ is a } k\text{-dimensional linear subspace of } \mathbb{R}^n \}. \quad (2.22)$$

That means it can be described as a further quotient manifold by $\text{Gr}_k(\mathbb{R}^n) \cong \text{St}_k(\mathbb{R}^n)/O(k)$, where the action of $O(k)$ is by change of the orthonormal basis of the subspace. We focus here on another description, namely via projection matrices. For this one identifies a linear subspace with the unique orthogonal projection onto it and obtains

$$\text{Gr}_k(\mathbb{R}^n) \cong \{ P \in \mathbb{R}^{n \times n} \mid P = P^2 = P^T, \text{tr} P = k \}, \quad (2.23)$$

which entails

$$T_P \text{Gr}_k(\mathbb{R}^n) \cong \{ V \in \mathbb{R}^{n \times n} \mid V = P V + V P \} \quad (2.24)$$

for the tangent space. Grassmannians occur, for example, in the context of video based face recognition [75, 78, 92, 93]. Via the quotient construction, we obtain again a metric that descends from the Stiefel manifold and is simply the Euclidean metric for this description of the tangent space, *i.e.* $\mathfrak{g}_P(V, W) = \text{tr} V^T W$. The corresponding Riemannian distance $\mathbf{d}_{\mathfrak{g}}$ is also called the *Grassmanian distance*. There is a range of further distances constructed on the Grassmannian, which are not-necessarily Riemannian distances. The most popular of these distances is the so-called *projection distance* $\mathbf{d}(P_1, P_2) := \|P_1 - P_2\|_F$ which is thus the Euclidean distance of matrices applied to the projections. It yields a computationally cheap local approximation of the Grassmanian distance, *i.e.* its differentiation yields the same metric as shown e.g. by Harandi

³The factor $\frac{1}{2}$ is for convenience when defining metrics on the following matrix manifolds.

et al. [76]. This projection distance has been used by a number of approaches to distance/metric learning on Grassmannians such as [75, 93]. Geometrically, the projection and the Grassmanian distance can also be described using the principal angles between the involved subspaces, which is also true for other proposed distances. For a list see, for example, [199, Table 2]. For a more general introduction into the geometry of the Grassmanian and the Stiefel manifold, see, for example, [21].

The Space of Immersions. The idea of the Grassmannian can be generalized by considering nonlinear, curved submanifolds of \mathbb{R}^n instead of linear subspaces. To this end, one considers a base manifold M with dimension smaller than n . Then one studies the space of all its immersions $\text{Imm}(M, \mathbb{R}^n)$ into \mathbb{R}^n of sufficient regularity. It can be shown that $\text{Imm}(M, \mathbb{R}^n)$ is indeed a manifold under certain assumptions [87], albeit an infinite-dimensional one⁴. When considering closed planar curves, *i.e.* $\text{Imm}(S^1, \mathbb{R}^2)$, this manifold can be used, *e.g.*, to model the boundaries of cells in microscopy images. Furthermore, surfaces immersed into \mathbb{R}^3 are used in computer graphics to model characters or similar and in medical imaging to model the shape of anatomical objects. The tangent space of $\text{Imm}(M, \mathbb{R}^n)$ consists of vector fields along the immersed surface on which various metrics on the space of immersions have been proposed. Popular choices include Sobolev metric of vector fields [15], metrics derived from elastic energies [83], and finally metrics that descend from metrics on the *space of diffeomorphisms* [3]. For the latter, one considers the action of diffeomorphisms on immersions by deforming the ambient space and introduces a metric on the Lie group of diffeomorphisms of sufficient regularity that is compatible with this action.

The Group of Diffeomorphisms. This concept can also be applied to other spaces on which we have an action of diffeomorphisms. A very common example for this is the *space of images*, which is usually formalized as square integrable functions on a square domain, *i.e.* $L^2([0, 1]^n, \mathbb{R}^k)$, where the outputs are color channels or similar. Here, one considers the left-action of diffeomorphisms (*i.e.* precomposition with the inverse) and a right-invariant metric on diffeomorphisms (*e.g.* a suitable Sobolev metric) giving rise to the celebrated LDDMM (Large Deformation Diffeomorphic Metric Mapping) framework [65, 163, 177, 181] which is a popular framework in the medical imaging community and remains an active field of research.

The Space of Probability Measures. Another space that is often used to model imaging data is the *space of probability measures*, *i.e.* the set of all probability measures on a given manifold M . To introduce a distance (which yields a Riemannian metric) on this space, one turns to *optimal transport*. Here, one aims to find the most efficient transformation between two probability measures in a certain sense. Given two measures μ and ν , one aspires to find the joint probability distribution π on the product space $M \times M$, which has marginals μ and ν and minimizes $\int_{M \times M} c(x, y) d\pi(x, y)$, where the function $c: M \times M \rightarrow \mathbb{R}$ specifies the cost of moving a unit of mass from one place to another. It is typically derived from the (Riemannian) distance on the base manifold \mathcal{M} . This gives rise to the *p-Wasserstein distance* on probability measures defined by

$$W_p(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \left(\int_{M \times M} \mathbf{d}_M(x, y)^p d\pi(x, y) \right)^{1/p}, \quad (2.25)$$

where $\Pi(\mu, \nu)$ as the set of all couplings as introduced above. This distance turns the space of probability measures into a formal Riemannian manifold.⁵ The field of optimal transport originates from studying economic problems and has since been revealed to have deep connections to convexity, partial differential equations, and statistics that have helped advancing several mathematical fields. More recently, optimal transport has established itself as an important

⁴We do not discuss this here but the ideas from Section 2 can be generalized to infinite-dimensional spaces, see Appendix A.

⁵The technical details on the geometry of the space of probability measure are more complicated and go beyond the scope of this review.

applied tool primarily in mathematical imaging and machine learning. For a more detailed introduction on this subject which also introduces other formulations and makes many of these connections concrete, we refer to the books by Santambrogio [152] and by Peyré and Cuturi [145]. In the metric learning context, one typically learns a distance on the base manifold M to obtain a metric on the space of probability measures via the Wasserstein distance. This has been applied, for example, to learn models for bird migration and single-cell RNA data [157], for image classification models [45, 96, 197], for color transfer [84], causal inference [59] and for domain adaptation [103, 206]. Other metrics have been used on the space of probability measures [172] to address specific application, but they are mostly designed by hand.

Tractability. In general, when we leave \mathbb{R}^d and instead work on one of the above Riemannian manifolds (or others), methods and algorithms become more computationally expensive. They usually require (optimization) algorithms that are tailored to these manifolds. Such algorithms have garnered considerable interest in recent years due to their applicability to a wide range of problems beyond metric learning and there exist standard text books such as [25]. For the matrix manifolds introduced above, the computational burden is often more acceptable since the operations needed in optimization—such as the exponential and logarithm—can be evaluated in closed form and considerable research effort has been put into making these efficient. However, for infinite dimensional spaces, such as spaces of shapes or measures, we often have to first pay the prize of a suitable discretization and then Riemannian operations entail costly numerical algorithms. For example, in the case of immersions of surfaces in three-dimensional space one often works with triangle meshes and images are discretized as piece-wise constant functions.

Geometric Prior. Nevertheless, considering data to lie on Riemannian manifolds allows us to encode additional structure of the data which is particularly interesting when working in data sparse regimes. In this case, this additional structure often helps to improve generalization. For example, van Veldhuizen [180] showed that using a curved, hyperspherical latent space within Variational Autoencoders improved performances compared to a flat Euclidean one. In the context of metric learning, we of course aim to imprint (some of) the structure of the data onto the metric. By providing a Riemannian metric on an Euclidean space, we also obtain a curved space and, by moving points to infinite distance, we can even encode some topology in the metric [41]. However, in some cases, having the manifold structure already as a strong prior can facilitate the learning process and improve performance, especially concerning generalization. These two aspects combined yield an intricate decision between how much structure to prescribe explicitly via the manifold and how much to learn via the metric which has to be made on a case-by-case basis and will be illuminated more in the following sections.

The Riemannian Manifold of Riemannian Metrics. As a *mise en abyme*, before ending this section, the manifold of Riemannian metrics should not be omitted [53, 67]. To avoid technical details, we assume that the base manifold $\mathcal{M} \subset \mathbb{R}^d$ is compact and we denote by $\mathcal{G}_{\mathcal{M}}$ the manifold of Riemannian metrics on \mathcal{M} . Since Riemannian metric can be seen as a smooth \mathcal{S}_{++}^d -valued function on \mathcal{M} , the tangent space $T_{\mathfrak{g}}\mathcal{G}_{\mathcal{M}}$ at $\mathfrak{g} \in \mathcal{G}_{\mathcal{M}}$ is isomorphic to the space of smooth, symmetric, matrix \mathcal{S}^d -valued functions on \mathcal{M} . The metric on $\mathcal{G}_{\mathcal{M}}$ will thus depend on our choice of metric on \mathcal{S}_{++}^d , where affine-invariant metric is arguably one of the most natural choices [144]. Then we extend this metric by integration on the manifold⁶, for any $\mathfrak{g} \in \mathcal{G}_{\mathcal{M}}$ and smooth functions $\mathfrak{h}, \mathfrak{k}: \mathcal{M} \rightarrow \mathcal{S}^d$

$$\mathbf{G}_{\mathfrak{g}}(\mathfrak{h}, \mathfrak{k}) := \int_{\mathcal{M}} \text{tr}(\mathfrak{g}_x^{-1} \mathfrak{h}_x \mathfrak{g}_x^{-1} \mathfrak{k}_x) \text{dvol}_{\mathfrak{g}}(x) \quad (2.26)$$

Note that the integration relies on the metric volume $\text{vol}_{\mathfrak{g}}$ in such way that $\mathbf{G}_{\mathfrak{g}}$ is not a trivial extension of the affine-invariant metric. However, the smoothness of Riemannian metrics does not appear explicitly in $\mathbf{G}_{\mathfrak{g}}$. Even if working in the manifold $\mathcal{G}_{\mathcal{M}}$ might appear challenging, it turns

⁶The integral is finite because the manifold is compact and the integrand is continuous.

out that a closed-formed expression is available for the geodesics which is given in [67, Theorem 3.2] even when the base manifold \mathcal{M} is not compact. Nevertheless, the geodesics are expensive to compute. While this structure seems natural, it is not yet used to perform Riemannian metric learning as we will point out in Section 5.

3 The Riemannian Metric Learning Problem

We first give this abstract definition of Riemannian metric learning, before detailing its different applications.

3.1 The Abstract Problem

If a *Riemannian metric learning* method is used, at some point, the dataset $X = (x_i)_{i \in [N]}$ has to be generated according to an underlying metric \mathfrak{g}^* related to a manifold \mathcal{M} , *i.e.* for any $i \in [N]$, x_i is the realization of $X_i \sim Q_{\mathfrak{g}^*}^i$ where $(Q_{\mathfrak{g}^*}^i)_{i \in [N]}$ are probability laws on the input space \mathcal{X} and $\mathcal{M} \subset \mathcal{X}$. Note that X_i is not necessarily in \mathcal{M} because of possible additive noise ε such that $X_i = M_i + \varepsilon$ where $M_i \in \mathcal{M}$. The dataset X can have a more complex structure, but we assume that $X = (x_i)_i$ here for the sake of simplicity and more complex datasets will be discussed in Section 4.2

Optimizing on the manifold $\mathcal{G}_{\mathcal{M}}$ of all possible infinitely differentiable Riemannian metric on \mathcal{M} can be tractable in dimension 1 [121, Chapter 4-5], but most of times it is not. Thus, a family of parametrized Riemannian metrics $\mathcal{G}_{\theta} \subset \mathcal{G}_{\mathcal{M}}$ is chosen. The wealth of Riemannian metric learning as an optimization problem arise with the use of different mathematical objects, to name the most commons: the geodesics distance $\mathbf{d}_{\mathfrak{g}}$, the exponential map $\text{Exp}_{\mathfrak{g}}$, the parallel transport $\mathfrak{T}_{\mathfrak{g}}$, the metric volume $\text{vol}_{\mathfrak{g}}$. Therefore, denoting by \mathfrak{L} the objective function, the abstract form of Riemannian metric learning can be stated as follow :

$$\operatorname{argmin}_{\mathfrak{g} \in \mathcal{G}_{\theta} \subset \mathcal{G}_{\mathcal{M}}} \mathfrak{L}(\mathcal{F}(\mathfrak{g}), X), \quad \mathcal{F}(\mathfrak{g}) = (\mathbf{d}_{\mathfrak{g}}, \text{Exp}_{\mathfrak{g}}, \mathfrak{T}_{\mathfrak{g}}, \text{vol}_{\mathfrak{g}}). \quad (3.1)$$

This optimization problem is hard for two reasons:

1. The manifold of Riemannian metrics $\mathcal{G}_{\mathcal{M}}$ has a specific structure [67] (see (2.26)). They are not only vectors or constrained-matrices in \mathcal{S}_{++} , they are \mathcal{S}_{++} -valued smooth functions with a manifold input space \mathcal{M} . $\mathcal{G}_{\mathcal{M}}$ is both constrained and infinite dimensional. This makes it challenging to find suitable discretizations of it, which will be further discussed in Section 4.1.
2. The mathematical objects depending on the metric $\mathcal{F}(\mathfrak{g})$ have, in general, not a closed form formula and their dependence on \mathfrak{g} is not linear. Thus, the optimization problem is not convex except when restricting to constant Riemannian metrics.

However, *Riemannian metric learning* not only offers inescapable problems, but also yield new opportunities for representation learning to leverage a large number of data in an unsupervised manner as we will see in Section 4. Before we *how* to learn a Riemannian, we will explain *why* by looking at the applications in the next section.

3.2 Applications: Why Learn a Metric?

First, we will present the applications related to the *distance metric learning* literature (Section 3.2.1), then the applications related to *Riemannian metric learning* (Section 3.2.2). When the learned distance $\mathbf{d}_{\mathfrak{g}}$ is induced by a non-linear Riemannian metric \mathfrak{g} , the application is presented in Section 3.2.2 instead of Section 3.2.1. We aim to showcase the richness of the Riemannian metric learning approach and to highlight the importance of each Riemannian object of interest, namely $(\mathbf{d}_{\mathfrak{g}}, \text{Exp}_{\mathfrak{g}}, \mathfrak{T}_{\mathfrak{g}}, \text{vol}_{\mathfrak{g}})$. Given that Riemannian metric learning is inherently technical and complex, it is crucial to ensure its relevance in the context of the application.

3.2.1 Distance Metric Learning

Learning the metric of a metric space (see Definition 1) can be seen as a foundational task which has its most prominent use in distance-based regression, classification (k-NN), clustering and dimension reduction. Instead of specifying a distance according to your own prior, you can learn relevant information from the dataset and introduce it in your distance.

Regression, Classification and Clustering. On Euclidean data, the paper [195] popularized metric learning for clustering and classification using nine dataset from the UC Irvine repository⁷ (UCI) by extending k-means [1]. This classical repository is composed principally of tabular dataset having less than 3000 samples and 10 classes with low dimensional inputs ($d < 100$). Then, it was followed by [37, 46, 70, 71, 102, 161, 185, 192, 203] outperforming each other chronologically regarding accuracy and time of computations on UCI datasets, Handwritten text recognition (MNIST⁸, USPS⁹) and Face recognition datasets such as Yales [17], PubFig [138], LFW [91].

A unified benchmark was performed in the review [169] and all methods are implemented in the related library pydml [168]. Interestingly, one of the last paper [38] improving on these methods suggest to learn a curved coordinate system to perform metric learning, which is close to Riemannian metric learning. This method presented was then used successfully for recommendation systems in [205].

These methods were then extended to other types of data such as multimodal data [131, 194, 204], sets [207], bag of instances [74] and to other contexts such as Multitask learning [74, 139, 197]. The applications were face, object, and artist recognition in images, videos, and audio. Using metric learning within multitask kNN classifiers [74] is particularly appealing since the number of parameters does not depend on the number of classes contrary to multitask-SVM [58].

Then, a new line of works arose in computer vision. Covariance matrices and subspaces features arising from images and videos were embedded in the Grassmanian [75, 78, 92, 93, 208] and the SPD manifold [77, 94, 183, 201, 208] and distances on these non-trivial manifolds were learned. Benchmarks were performed on objects recognition task using the ETH-80 datasets [113], on material recognition using the Flickr material dataset [162], UIUC material dataset [115], and on face recognition using the Youtube celebrities [104] and Youtube Face dataset [193]. One of the latest papers in this direction, [208], offered a general method working both on the SPD and Grassmanian manifold and further improved accuracy and computation time.

As mentioned in [101, 149], deep metric learning methods have widened the scope of applications and improved on the state-of-the-art on different task of computer vision [135]: visual tracking [90], face recognition [39], person re-identification [85], labeling radiography [2], 3D objects retrieval [27, 82], semantic instance segmentation [61], semantic edge detection [31], pan-sharpening [196], volumetric image recognition [189], crowdedness regression [188]. The methods apply also in audio and natural language processing with topic classification [56], music ranking [125], and speaker recognition [187] using so-called prototypical networks [165]. However, the success of deep metric learning methods was debated by Musgrave et al. [136], who raise issues about the lack of clarity and fairness in the experiments.

Dimension Reduction. The task of dimension reduction [132, 186], also called *manifold learning*, is used to visualize high dimensional datasets or to pre-process data by removing unnecessary components. Most methods such as Multidimensional Scaling (MDS) [43] or Stochastic neighborhood embedding (SNE) [86] reduce the dimension while preserving global or local distances from the raw dataset. That is why, manifold learning methods can benefit from learning a distance.

Moreover, if the distance metric learning method uses a projection matrix L , such as in [37,

⁷<https://archive.ics.uci.edu>

⁸<https://yann.lecun.com/exdb/mnist/>

⁹<https://paperswithcode.com/dataset/usps>




Figure 5: Illustration of a normal and a global coordinate chart for feature representations.

46, 70, 71, 102, 161, 185, 192], the eigenvectors related to the largest eigenvalues of L can also be used directly to project the dataset in small dimensions as explained in the review dedicated to this topic [186].

There are also methods employing *Riemannian metric learning* for dimension reduction which are given in the next sub-section.

3.2.2 Riemannian Metric Learning

As discussed in Section 3.1, a Riemannian metric can be used for modeling through different objects: the distance \mathbf{d}_g , the exponential map Exp_g related to geodesics γ , the parallel transport \mathfrak{T}_g , the metric volume vol_g or even sometimes the Ricci curvature of the manifold R_g [59, 114]. In the following, we clarify and exemplify why these objects can be useful for representation learning, Hamiltonian learning, temporal longitudinal modeling, distribution trajectory inference, sampling and generative modeling, domain adaptation, and finally causal inference.

Representation Learning. The first paper we found with “Riemannian metric learning” in its title was by Peltonen et al. [143] in 2004 for exploratory analysis and dimension reduction. They argued that, although Riemannian metric learning increased the computational burden, they used it since global distance would not have preserved the topology of the data space, which they deemed important for identifying subcategories of the classes within a dataset [143, p.3]. Other papers also used principles of Riemannian metric learning, without saying it, for instance Isomap [173, 2000] is arguably one of the most famous. In Isomap [173], the geodesic distance \mathbf{d}_g is defined according to the k -nearest neighbors graph, which can cause robustness issues [12]. Recently, a variant of Isomap using parallel transport \mathfrak{T} proposed by Budninskiy et al. [28] revealed to be more robust. Therefore, representations can be learned through Riemannian metric learning, not only by using distances but also by using connections on tangent bundles.

Another line of work, learn representations via *normal coordinate charts* [116] using the logarithm and exponential map. If a general metric g is learned, a point cloud $(x_i)_{i \in [N]}$ on the related manifold \mathcal{M} can be represented by their logarithms $(\log_p(x_i))_{i \in [N]}$ taken at their Fréchet mean p defined as

$$p = \underset{q \in \mathcal{M}}{\text{argmin}} \sum_{i=1}^N \mathbf{d}_g(q, x_i)^2 \quad . \quad (3.2)$$

Then, any standard method of dimension reduction (e.g. PCA) can be applied to these representations $(\log_p(x_i))_{i \in [N]}$ [65, 80, 116, 181]. This strategy can thus transform points $(x_i)_i$ on a Riemannian manifold \mathcal{M} , where there is not necessarily a group structure, to vectors $(\log_p(x_i))_{i \in [N]}$ in a vector space $T_p \mathcal{M}$, where there is a group structure. This opens the toolbox of Euclidean statistical methods, which is appealing to analyze complex objects such as *shapes* [179] as depicted in Figure 6 with mouse respiratory cycles.

Simpler than a normal coordinate chart is a global coordinate chart ψ as used in Section 2 and implicitly in autoencoders [13] and variational autoencoders [105]. Here, the observations




Figure 6: On the left (A), an example taken from [65] (with permission) of normal coordinate chart representing shapes of respiratory cycles related to two mouse genotypes: colq and wt. On the right (B), an example taken from [171] (licensed under [CC BY 4.0](#)) of latent representation where the metric encodes uncertainty information related to an ensemble of decoders ψ_θ .

$y = \psi_\theta(z)$ are represented by low dimensional latent variables $z \in \mathcal{Z}$. This idea was further developed in [8, 34, 48, 49, 60, 63, 99] to enhance training and interpretability of the representation through visualization of geometric features, which might also encode task-related uncertainty [171, 191] as depicted in Figure 6. Once a latent space and its geometric structure $(\mathcal{Z}, \mathfrak{g})$ are learned, different applications involving Riemannian distance, geodesics, or simply features representation can be performed. Geodesics in the latent space are especially appealing for temporal longitudinal modeling and distribution trajectory inference as detailed in the next sections.

Hamiltonian Learning. Before presenting the applications related to geodesics modeling, we include a discussion on the broader context of Hamiltonian learning. In physics, trajectories are often analyzed using Hamiltonians, which describe preserved energy along the trajectory. The Hamiltonian related to an object at position q and momentum p is denoted by $\mathcal{H}(q, p)$ and decomposes as $\mathcal{H}(q, p) = V(q) + p^\top \mathfrak{g}(q)^{-1} p / 2$ with $V(q)$ being the potential energy at position q and $p^\top \mathfrak{g}(q)^{-1} p / 2$ the kinetic energy depending on the metric of the space. Recently, deep learning methods were suggested [40, 176] to model the underlying Hamiltonian as a neural network and infer it from observed trajectories. This is depicted on the left of Figure 7.

Neural Hamiltonians connect with Riemannian metric learning since geodesics are Hamiltonian trajectories for the special case of $V \equiv 0$ as mentioned in Remark 14. When the kinetic energy is fixed within Neural Hamiltonians, *i.e.* \mathfrak{g} constant, we only need to analyze the potential V to get insight about the state space, which can enhance the interpretability of the Hamiltonian learning [167]. In the other way around, if V is removed or fixed, learning the metric \mathfrak{g} shed light not only on states q , but also on displacement p . It is not only about where you are, but where you want to go. In the next paragraph, we illustrate these ideas with Longitudinal modeling and Trajectory inference.

Temporal Longitudinal Modeling. In various scientific fields, a key objective is to deduce the dynamics of an underlying system from noisy data. This is particularly challenging in areas such as biomedicine, where only limited longitudinal data is available—such as when tracking health indicators or studying disease progression [81]. For example, in Alzheimer’s disease follow-up, psychometric assessment and brain imaging are available at different times depending on the patient [128].

In temporal longitudinal modeling¹⁰, the same individual i is observed with features (y_i^j) at different time points (t_i^j) which can be irregularly sampled. The goal is to recover the underlying process $t \rightarrow Y_i(t)$ such that $Y_i(t_i^j) = y_i^j + \varepsilon_i^j$ with ε_i^j an i.i.d. noise. One approach is to model the individual trajectories Y_i as geodesics, *i.e.* for any $t \in [0, 1]$, $Y_i(t) = \text{Exp}_{q_i}(tv_i)$ with

¹⁰This is not the general case of longitudinal modeling. In longitudinal modeling the outcome variable $Y(x)$ might depends on a covariate x which is not a scalar as timepoints.

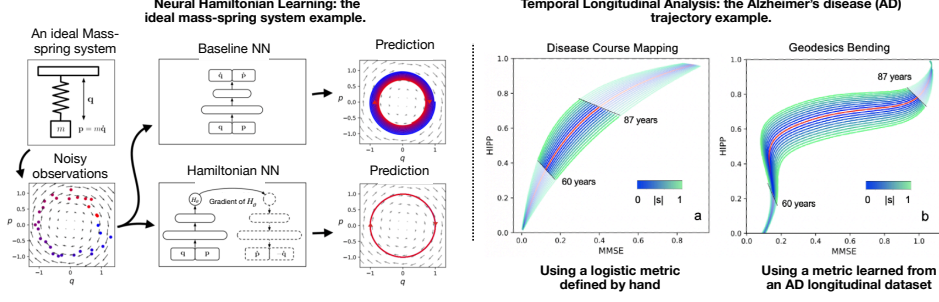


Figure 7: On the left, an example of Hamiltonian learning taken from [40] (licensed under [CC BY 4.0](#)). Neural Hamiltonian learning is compared to a naive neural networks encoding the system by its differential equation. Note that predicted baseline trajectories are thicker because of the larger uncertainty of estimation on the derivatives compared to the Hamiltonian. On the right, an example taken from [72] (with permission) where the geodesics are learned to mimic AD patient trajectories. HIPP is the cerebral volume of the Hippocampus and MMSE refers to the Mini Mental Status Examination score. The scores were reversed and normalized in order to be increasing in $[0, 1]$. The learned exponential shape in figure b is consistent with clinical observations, but is not present when considering a hand-crafted logistic metric [159] showing the interest of learning the Riemannian metric from scratch. The tubular green and blue curves are the exp-parallelization of the average geodesic in red for different values of space shifts w (3.3), it represents the inter-individual variability of trajectories along the mean trajectory.

$(q_i, v_i) \in \mathcal{M} \times T\mathcal{M}$. The idea is that geodesics generalized straight lines which make them both flexible and simple enough to avoid overfitting on noisy datasets. Geodesics learning was also used to model trajectories on rotations manifold to interpolate animations [109], human or robot motion captures [16, 175], image registration [163], and taxi routes in traffic jams [23, 148]. We should note that geodesic learning can also occur in exotic application such as a two player game where the goal of the first player is to install a surveillance system and the goal of the second is to reach a target area without being detected [134].

In the following, we detail the case of Alzheimer's disease (AD) follow-up to enlighten how Riemannian metric learning can be relevant in a specific application. Geodesics are particularly suited to model AD follow-up [120, 159] since AD is a neurodegenerative disease, *i.e.* there is rarely recovery of patients, which implies that there are no loops in patient trajectories [97]. Moreover, there is an inter-patient variability in the AD trajectories. That means two AD patients can have a similar trajectory direction because of their disease, but their trajectories remain parallel and distinct because of their individual health conditions. This variability can be identified within mixed-effect models by generalizing the concept of parallel straight lines using the exp-parallelization [159]. The exp-parallelization $\text{ExPar}(w)(\gamma, \cdot)$ of a curve $\gamma: \mathbb{I} \subset \mathbb{R} \rightarrow \mathcal{M}$ at $t_0 \in \mathbb{I}$ in the direction $w \in T_{\gamma(t_0)}\mathcal{M}$ is defined as,

$$\text{ExPar}(w)(\gamma, t) = \text{Exp}_{\gamma(t)}(\mathfrak{T}_{\gamma, t_0, t}(w)), \quad (3.3)$$

for any $t \in \mathbb{I}$. This shows how parallel transport and the exponential can be combined when optimizing a Riemannian metric. Note that the exp-parallelization of a geodesic is not necessarily a geodesic. An illustration of the transformation is given on the right of Figure 7. Alzheimer's disease does not start at the same age for everyone and does not evolve at the same speed. This generates inter-patient temporal variability. Thus, it is crucial to compose an exp-parallelization with an individual time warp $\psi_i: \mathbb{I} \rightarrow \mathbb{R}$ as $Y_i = \text{ExPar}(w_i)(\text{Exp}_{q_0}(v_0 \cdot), \psi_i)$ with $(q_0, v_0, w_i) \in \mathcal{M} \times T_{q_0}\mathcal{M} \times T_{q_0}\mathcal{M}$. Parameters related to time (ψ_i) and space (q_0, v_0, w_i) are statistically identifiable as long as $\|v_0\| = 1$ since the geodesics are parametrized by arc length (Remark 13). This property is particularly interesting to disentangle time and space variability in the analysis.




Figure 8: Example of trajectory inference with different choices of interpolant. $P_{1/2}^{Euclidean}$ and $P_{1/2}^{Geo}$ are the marginal density at time $t = 1/2$ when the interpolant uses respectively straight lines or geodesics, respectively.

This line of works was thus naturally extended by Riemannian metric learning methods [72, 121, 123, 156]. Here, the goal was to learn the best metric related to the mixed-effect model to represent the patients’ trajectories. The resulting representation of the patient trajectories can then be analyzed to understand the disease progression according to different health conditions. This proved successful to handle multimodality [62, 72, 154].

Distribution Trajectory Inference. In single-cell RNA sequencing (scRNA), cross-sectional data collection is sparse and static due to the high costs and destructive nature of the procedure [126]. As a result, there is a need to create models capable of reconstructing a system’s temporal dynamics (for example, in cellular studies) from observed samples taken at discrete time points. This general problem is known as *trajectory inference* [111]. In trajectory inference, we only access the population distribution $(P^j)_{j \in [T]}$ at different fixed time points $(t^j)_{j \in [T]}$. The goal is to recover the continuous trajectory of the distribution by interpolation.

One can approach this problem via optimal transport [152]: For any $j \in [T - 1]$, a transport map T^j is learned such that $P^{j+1} = T^j \# P^j$ is the pushforward measure of P^j by T^j . When using the Euclidean metric as base metric, the interpolation becomes $P^{j+t} = (tT^j + (1-t)\text{Id}) \# P^j$ for any $t \in [0, 1]$. However, this Euclidean interpolation can put some mass in regions of the space which are unrealistic. A conceptual illustration is given in Figure 8. That is why Scarvelis and Solomon [157] and Huguët et al. [95] proposed to learn a Riemannian metric to be used in the optimal transport problem. In the end, for any x in the observational space, defining $\gamma(x, t)$ as the geodesics between x and $T^j(x)$ for any $j \in [T - 1]$, the interpolation becomes $P^{j+t} = \gamma(\cdot, t) \# P^j$. The method improved the results on a bird migration dataset¹¹ and an scRNA dataset [158].

Flow matching [117] is another option, where the interpolation of distributions is achieved by integrating a flow equation $\frac{dX}{dt} = v_t(X)$. Here, the velocity field is learned in order to match a quantity related to an interpolating function $I: \mathcal{M}^2 \times [0, 1] \rightarrow \mathcal{M}$. This function is a prior on the shape of trajectory we want after integrating the flow equation. Most of the time, I is chosen to be linear because of its simplicity. However, as for Euclidean optimal transport interpolation, this choice is not optimal since the interpolation might leave the data manifold as depicted in Figure 8. Again, a Riemannian metric can be learned from the dataset and geodesics can be used to define the desired interpolant as $I(x, y, t) \approx \text{Exp}_x(t \text{Log}_x(y))$ using neural network approximations [100, 170]. Kapusniak et al. [100] and Sun et al. [170] successfully used Riemannian metric learning within flow matching for single-cell trajectory inference on the CITE-seq and Multiome datasets from a NeurIPS competition [29], where both methods greatly improved on the state-of-the-art.

Sampling and Generative Modeling. More generally, Riemannian metric learning finds also applications in sampling [34, 35, 44] and generative modeling [5, 8, 63]. In these applications, the metric \mathfrak{g} and its volume $\text{vol}_{\mathfrak{g}}$ are particularly involved. In [68], the authors demonstrate that sampling algorithms such as Hamiltonian Monte Carlo (HMC) [24] can be improved by using a

¹¹<https://birdsoftheworld.org/bow/species/snogoo/cur/movement>




Figure 9: On the left, a Riemannian manifold structure is learned [170] (licensed under [CC BY 4.0](#)) using embryoid data and data augmentation is performed by sampling according to the metric volume. On the right, synthetic proteins are generated in [47] (licensed under [CC BY 4.0](#)) by interpolating a protein with high hydrophobic moment and low charge x and another one with a low hydrophobic moment and high charge y .

Riemannian metric related to the distribution. Chadebec et al. [34] built on this idea by learning the metric on a latent space in a variational autoencoders to perform data augmentation for brain MRI [35]. Learning more structure in the latent space is particularly interesting [35, 180] when datasets are small and high dimensional—such as MRI datasets. Sun et al. [170] suggest to sample uniformly according to the volume vol_g in the Riemannian latent space in order to avoid selection bias due to observations using the unadjusted Langevin algorithm as depicted on the left of Figure 9. Cui et al. [44] remind us that the convergence guarantee of Langevin dynamics to their invariant distribution is related to Poincaré inequalities. Thus, they optimize the Poincaré constant to accelerate convergence by learning a Riemannian metric on the state space. Therefore, Riemannian metric learning can be a tool to improve sampling efficiency motivated by convergence results.

In an unsupervised learning context, Arvanitidis et al. [5] suggested to learn a Gaussian mixture on a Riemannian manifold adapted to a sleep stages dataset from Physionet. They adapted the geometry of the Gaussian in order to avoid holes in the data distribution. Using autoencoders and variational autoencoders, several papers [6, 8, 9, 47, 63] have underlined that learning the Riemannian structure of the latent space improves uncertainty estimation and interpretability.

For instance, de Kruiff et al. [47] they learn an isometric map between the latent space and the observational space in order to preserve distance related to biological properties of observed proteins. By using the Riemannian structure of their latent manifold, they can generate new proteins with specific biological properties with simple geodesic interpolation [47, Figure 5] as depicted on the right of Figure 9. A score-based variant of the method is given in [49] and its theoretical underpinnings are discussed in more detail in [48].

Domain Adaptation and Transfer Learning. The main goal of domain adaptation and transfer learning is the following to generalize a learned statistical model from a source distribution P_S to a target distribution P_T using weak information. For instance, in the case of classification, the target features are available but not their labels [11].

Recently, optimal transport has become a key tool in domain adaptation. For classification, the target features are transported on source features by minimizing the global cost of displacement and then the statistical model estimated on the source domain can be applied. The cost of displacement, *i.e.* the ground metric, can be learned to improve performance [30, 45, 84, 96]. This way domain adaptation benefits from distance metric learning [103, 206].

Furthermore, domain adaptation can also benefit from Riemannian metric learning using parallel transport. For instance, in cardiac motion modeling, when a motion is estimated as a geodesic passing through q_0 in the source domain, we can transport this trajectory to a point q_1

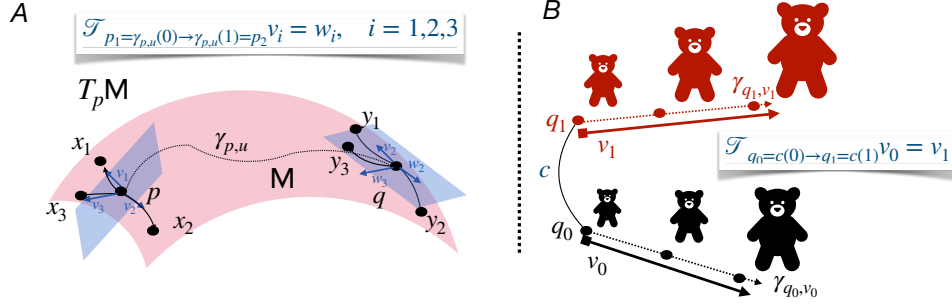


Figure 10: On the left (A), the domain adaptation is performed by transporting the normal coordinates from p_T to p_S . On the right (B), the growth of the black bear is transferred to the red bear using parallel transport in a manifold encoding bear shapes and colors.

in the target domain by transporting its velocity in parallel along the geodesic going from q_0 to q_1 [73]. This is illustrated on the right of Figure 10 with bear shapes instead of cardiac motions for the sake of simplicity.

Tam et al. [172] apply a similar reasoning when working on a Riemannian space of probability distribution. Points in the source and target domains can be represented using normal coordinates at their Karcher means p_S and p_T (3.2), respectively. Then the normal representation of the target features can be transported to the source domain using the parallel transport following a trajectory going from p_T to p_S . This is illustrated on the left of Figure 10. In practice, to reduce the computational burden, Tam et al. [172] learn a statistical model (linear regression, PCA) on the normal representations in the source domain, then they transport the statistical model to the target domain. Crucially, this transport is independent of the number of data points [172].

More generally, when one learns a regressor used for a downstream task, it is important to adapt the loss of your regression according to the task you target. This idea was developed by Bansal et al. [14] using Riemannian metric learning and successfully tested on portfolio management applications.

Causal Inference. Causal inference is a growing subject in the machine learning community [26, 98, 160]. Recently, a link was drawn between causal inference [140] and Riemannian geometry [51]. In causal inference [141], the outcome variable Y is represented as the product of covariates X and treatment T , such that $Y = f(X, T, U)$ where U is a random exogenous factor. The goal is then estimate to the causal effect of the treatment. For instance, Y could indicate whether the patient has a headache, X the age and the diet of the patient, and T could indicate whether the patient has taken some pain medication. Here, we might aim to measure the effect of the pain medication on the headache. Most of the time, the treatment effect is estimated with matching methods: we hire a group of patients which will take the treatment, then we hire a control group of patients which will not take the treatment. However, we select the second group such that each patient in the first group is paired with a patient in the control group having similar covariates. By doing so, the covariate effect is approximately removed and we can directly observe the treatment effect by computing the difference of average between the two groups.

The causal mechanism gives structural equations (e.g. $Y = f(X, T, U)$) which can induce a geometric structure on the data [51]. Thus, when manipulating causal inference data, we should respect this structure. This idea was developed by Pradier and González [147], who learn a Riemannian metric to compute the Riemannian distance between patient covariates X before applying matching methods: the patients are paired if the distance between their covariates is small enough. This method revealed to be competitive on the synthetic version of the Infant Health Development Program Dataset introduced by [124] and the Lalonde dataset [110]. Similar ideas have been applied to generate counterfactual explanations [142] by sampling in a Riemannian latent space.

When the treatment is given to patients in a community where their health conditions influence each other, the measurement of the treatment effect is complicated further. Recently, Farzam et al. [59] have drawn a link between the curvature of the graph related to the community and the treatment effect estimation. They even suggest to use the Ricci curvature related to the underlying manifold to adapt the metric of the graph before performing the estimation.

In summary, since the geometric point of view on structural equations is relevant [51, 59], causal inference benefits from Riemannian metric learning by using the Riemannian distance [147] or more complicated objects such as the Ricci curvature [59].

4 The Mathematical Problem: How to Learn a Metric?

Now that we have introduced a range of applications for metric learning, we will turn back to its technical aspects, *i.e.* how to make the abstract problem from Section 3.1 more concrete and how to solve it. While the choice of the objective function \mathcal{L} is at the crux of the problem (Section 4.2), the choice of the parametrization (Section 4.1) will deeply influence its flexibility and its computational tractability. Interestingly, the optimization methods related to Riemannian metric learning (Section 4.3) are less specific to the problem than one would expect, the discussion on this issue is continued in Section 5.

4.1 Parametrization of the Metric

The first aspect we discuss is how to parametrize the learned metric, or, in other words, which ansatz space $\mathcal{G}_\theta \subset \mathcal{G}_\mathcal{M}$ to select for the learning problem (3.1). As highlighted in Section 3.1, the manifold of metrics $\mathcal{G}_\mathcal{M}$ on a manifold \mathcal{M} is infinite-dimensional. Thus, to perform optimization, we must restrict our focus to a finite-dimensional subspace. Over the years, various approaches have been proposed for this purpose, and we will summarize and critically discuss the most significant ones.

To achieve this, we broadly categorize the approaches into two categories: explicit and implicit. In *explicit* approaches, a map $\mathcal{M} \rightarrow \mathbb{R}^{d \times d}$ is directly learned. Conversely, *implicit* approaches involve learning an alternative structure—such as a distance function, an embedding and so on—on the manifold, which still induces a Riemannian metric.

At the end of this section, we discuss the computational challenges related to the estimation of the Riemannian objects $\mathcal{F}(\mathfrak{g}) = (\mathbf{d}_\mathfrak{g}, \text{Exp}_\mathfrak{g}, \mathfrak{T}_\mathfrak{g}, \text{vol}_\mathfrak{g})$ related to the metric.

4.1.1 Explicit Methods

Choosing an explicit expression for the metric can be viewed as parametrizing a function that maps to the space of symmetric positive-definite (SPD) matrices \mathcal{S}_{++}^n . This task parallels regression methods, where the function can be chosen as constant, piecewise constant, or built using kernels or neural networks.

Constant Metric. The first—and most straightforward—explicit approach is also one of the most prevalent in the metric learning literature: selecting a *constant* metric across the entire manifold. In the case of an Euclidean space $\mathcal{M} = \mathbb{R}^d$, this reduces to learning a single SPD matrix G . This approach has been widely used in distance metric learning [37, 46, 70, 71, 74, 89, 139, 161, 185, 192, 194, 195, 204, 207]. It is typically framed as learning a *Mahalanobis distance*:

$$\mathbf{d}(x, y) = \sqrt{(x - y)^T G (x - y)} = \|A(x - y)\|, \quad (4.1)$$

where $G = A^T A$.

The idea of a constant metric has also been extended to other types of manifolds. A common extension is to matrix manifolds, where the Frobenius inner product can be weighted,

for example, via expressions like $\text{tr}(V^T G W)$ or $\text{tr}(G V^T W)$, or by adapting the distance functions introduced in Section 2.1. Such approaches have been applied to SPD matrix manifolds [41, 94, 183, 208] and Grassmannian manifolds [75, 78, 208].

Another common extension targets homogeneous spaces, such as $SO(n)$, where the metric at any point on the manifold can be obtained by deforming a reference metric defined at the identity element. In this case, one only needs to learn a suitable metric at the reference point with the usual diffeomorphic action propagating it to the entire manifold. This idea has been successfully applied to products of $SO(3)$ by Kushner et al. [109] and to finite distributions [45, 96, 103, 197].

While this parametrization is computationally efficient, it has a significant limitation: it fails to capture the inhomogeneous geometry of more complex manifolds and data.

Piecewise Constant Metric. This prompted researchers to explore *spatially varying metrics*. A natural first step in this direction is to consider *piecewise constant metrics*, which correspond to a non-conforming ansatz space, as the resulting metric is not globally smooth.

The most common approach for constructing piecewise constant metrics involves defining a metric at selected key points $c_i \in \mathcal{M}$ and extending it to the entire manifold using nearest-neighbor interpolation. This effectively assigns a constant metric within each Voronoi cell. Such an approach has been employed in studies such as [80, 148].

An alternative method involves discretizing the manifold with a (simplicial) grid and assigning a constant metric to each grid cell. This strategy, inspired by finite element methods, was explored by Cui et al. [44] and Qiu and Dai [148], who applied it from a computational perspective. However, discretizing high dimensional manifolds using a grid quickly yields computational challenges.

While these methods are conceptually simple and offer flexibility, they do not yield a smooth Riemannian metric due to discontinuities at the boundaries of Voronoi cells. The lack of smoothness at these singularities may limit their theoretical rigor and practical applicability.

Kernel Metric. A conforming approach, meaning one that yields a *smooth* metric, can be achieved through *smooth interpolation*. This is often accomplished using kernel-based or radial basis function interpolation. In this framework, one selects or optimizes key points $c_i \in \mathcal{M}$ along with corresponding metrics \mathfrak{g}_i defined on the tangent spaces $T_{c_i} \mathcal{M}$. The metric at any point $x \in \mathcal{M}$ is then obtained as the weighted sum

$$\mathfrak{g}_x = \sum_i \mathfrak{g}_i k(x, c_i), \quad (4.2)$$

where $k(x, c_i)$ is a kernel function that defines the interpolation weights. This method has been studied extensively in the literature [39, 80, 121, 155, 156, 163], with Gaussian kernels being a popular choice, typically of the form $k(x, y) = e^{-\frac{\|x-y\|^2}{\sigma^2}}$. As with constant metrics, this interpolation approach can be extended to matrix manifolds and homogeneous spaces by leveraging their specific geometric structures.

While this method provides a flexible and moderately interpretable way to define Riemannian metrics, it may face scalability issues in high-dimensional spaces due to the computational complexity of kernel evaluations and the potential curse of dimensionality.

Neural Metric. Last but not least, an approach to explicitly parametrize the metric, which is gaining increasing popularity, involves the use of *neural networks*. Thus far, these methods have primarily been applied in Euclidean spaces, where a key challenge is ensuring that the output of the network defines an SPD matrix.

For example, Arvanitidis et al. [8] address this challenge by restricting the metric to a scaled identity matrix of the form $\lambda(x)I_d$, where the scaling factor $\lambda(x)$ is also tied to the data density. In contrast, Scarvelis and Solomon [157] and Bansal et al. [14] propose a more general parametrization of the metric as $G_\theta = A_\theta^T A_\theta + \eta I_d$, where $A_\theta: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ is a neural network and η is a small positive constant to ensure positive definiteness.

Another neural network-based approach takes inspiration from the Hamiltonian formulation of Riemannian geometry (see Remark 14) by directly parametrizing the Hamiltonian, as explored in [40, 176]. This perspective allows for learning the underlying geometric structure through dynamics.

Finally, Chadebec et al. [34, 35], Scarvelis and Solomon [157], van Veldhuizen [180] combine neural networks with interpolation techniques and variational autoencoder latent distribution. In their approach, the metric is first parametrized at key points using a neural network of the form $G_i = A_i A_i^T + \eta I_d$, and then extended across the manifold via Gaussian kernel interpolation.

While neural network-based methods benefit from the wealth of ongoing research and tools in deep learning, they remain computationally expensive and often lack interpretability compared to more traditional approaches.

Hand-crafted Metric. Since learning a Riemannian metric completely can be computationally expensive and given that some problems can be solved without precise metric estimation, several studies propose to define the Riemannian metric manually based on the desired geometric properties.

In many cases, the primary objective is to ensure that geodesics remain close to an observational dataset $\mathcal{D} = (x_i)_i$, as real-world observations often lie on an unknown manifold $\widetilde{\mathcal{M}} \subset \mathbb{R}^d$. To achieve this, Arvanitidis et al. [5, 6] and Kapusniak et al. [100] suggest defining a metric \mathfrak{g} on \mathbb{R}^d as

$$\mathfrak{g}_x = (\text{diag}(h(x)) + \epsilon I_d)^{-1}, \quad h(x) = \left(\sum_{i=1}^N (x_i^j - x^j)^2 \exp\left(-\frac{\|x - x_i\|^2}{2\sigma^2}\right) \right)_{j \in [d]}, \quad (4.3)$$

for any point $x \in \mathbb{R}^d$ with coordinates $x = (x^j)_{j \in [d]}$ where $\epsilon > 0$ and $\sigma > 0$ are hyperparameters to be tuned. With this choice, when x is far from the dataset \mathcal{D} , $h(x)$ will be small because of the exponential term contrary to the metric’s minimal eigenvalue since $\mathfrak{g}_x \approx \epsilon^{-1} I_d$. Thus, when x is far from the dataset \mathcal{D} , it induces large energies of displacement as the sketched in Figure 3. To increase the contrast, the hyperparameters are chosen such that $h(x) \approx 1$ when x is in the dataset. This method has been shown to be effective in tasks such as flow matching, where geodesics play a critical role [100].

Hand-crafted metrics are also relevant when the related manifold is infinite dimensional such as the space of probability measures [172] or the group of diffeomorphism [65, 181] to avoid computational issues. In the same way, even in finite dimension, when data is scarce and belongs to a constrained manifold, hand-crafted metrics are still relevant [97, 180]. Designing a metric by hand is also performed in implicit methods, but they distinguish themselves from explicit form by learning other objects used in the process.

4.1.2 Implicit Methods

In implicit methods, the metric is not learned as a matrix field $\mathcal{M} \rightarrow \mathbb{R}^{d \times d}$ but derived from different objects. Compared to explicit metric parametrizations, implicit methods can be more computationally efficient while still offering significant flexibility.

Pullback Metric. One of the most popular implicit approaches is the use of a *pullback metric*. In this framework, a (smooth) map $f: \mathcal{M} \rightarrow \mathcal{N}$ is learned from the original manifold \mathcal{M} to another Riemannian manifold \mathcal{N} and the metric on \mathcal{M} is induced by the pullback, *i.e.*

$$\mathfrak{g}_x(v, w) = \mathfrak{g}_{f(x)}(\mathbf{d}_x f \cdot v, \mathbf{d}_x f \cdot w), \quad (4.4)$$

where $\mathbf{d}_x f$ denotes the differential of f at point x . This pullback metric is valid as long as f is locally injective. The map f can take various forms depending on the application. In the simplest case, f is a linear map, which brings us back to the scenario of Mahalanobis distance learning. More commonly, nonlinear forms of f are used including kernel methods [72, 92, 131, 175] and neural networks [63, 99, 123, 206].

A frequent application of pullback metrics involves autoencoder architectures [16, 99, 142, 170, 171, 175, 191, 200] where the metric can be induced by either the encoder or the decoder depending on the task requirements. When $f \sim P$ is random or not unique in the context, relying on the expected value $\mathbb{E}_{f \sim P}(d_x f)$ instead of $d_x f$ is an interesting extension account for the probabilistic nature of the space geometry [60, 171, 175]. Arvanitidis et al. [9] designed the metric according to a learned prior in the latent space and not using the pullback to gain robustness while still having a similar structure.

While some method directly explicitly define the pullback metric [72, 99, 123, 175], others, such as [11, 27, 31, 39, 52, 56, 61, 77, 82, 92, 93, 95, 102, 131] focus on learning a local approximation of the distance using the map-induced distance

$$\mathbf{d}_f(x, y) = \|f(x) - f(y)\|. \quad (4.5)$$

Note that \mathbf{d}_f is only a local approximation of the geodesic distance, as it does not account for the full geometry of the manifold as explained in Remark 8.

In some cases, the pullback metric is replaced with a *pushforward metric*, where a diffeomorphism $\tilde{f}: \mathcal{N} \rightarrow \mathcal{M}$ is learned instead. This approach induces a pushforward structure in both the input and output spaces [47, 48, 49, 72]. An advantage of pushforward metrics is that key Riemannian objects such as the exponential map, the logarithm map, and parallel transport often have closed-form expressions. These formulas are detailed in [48].

This approach is computationally feasible, quite flexible and interpretable as long as the function f shares these property.

Curvilinear Coordinate System. Instead of directly modeling the smooth map f , Chen et al. [38] and Zhou et al. [206] propose learning a *curvilinear coordinate system*. This approach corresponds to identifying the principal geodesics of the target metric. The distance between points is then computed by projecting them onto these principal geodesics—*i.e.* the curvilinear coordinate system—yielding a local approximation of the Riemannian distance.

This original idea has shown success in distance metric learning [38]. However, its application to general Riemannian metric learning remains limited, as this would require the computation of more complex geometric objects, such as parallel transport, to fully exploit the underlying manifold structure.

Graph-based Approach. The final implicit method for encoding a metric that we consider is the *graph-based approach*. In this approach, it is assumed that sparse measurements or approximations of distances on the manifold are encoded through the edges of a graph. Given this graph, the distance between points can be approximated by computing the shortest paths on the graph, which in turn induces a metric [23, 59, 84, 114, 143].

Such graph-based methods are commonly used to learn distances in a non-parametric manner, based on the assumption that the dataset is isometrically embedded in an ambient space. The nearest-neighbor graph is then computed leading to algorithms like Isomap [173] and its extensions such as [28].

4.1.3 Computing Riemannian Objects

Once one has parametrized the Riemannian metric \mathfrak{g} , one has to compute the Riemannian objects of interest $\mathcal{F}(\mathfrak{g}) = (\mathbf{d}_{\mathfrak{g}}, \text{Exp}_{\mathfrak{g}}, \mathfrak{T}_{\mathfrak{g}}, \text{vol}_{\mathfrak{g}})$ as explained in Section 3. Here, we provide a few elaborations on how to do this.

Constant Metric. For the case, that $\mathfrak{g} \equiv G$, where G is a symmetric positive definite matrix and we are working on Euclidean space, the associated distance $\mathbf{d}_{\mathfrak{g}}$ corresponds to the Mahalanobis distance. In this case, geodesics are straight lines, the parallel transport reduces to the identity map, and the metric volume remains constant at $\sqrt{|\det(A)|}$. This again highlights the lack of nonlinear structure provided by a constant metric.

Non-constant Metric. For nonlinear explicit forms, no universally accepted gold standard currently exists. The metric volume $\text{vol}_{\mathfrak{g}}$ requires computing a determinant at each point $x \in \mathcal{M}$. The exponential map $\text{Exp}_{\mathfrak{g}}$ and the parallel transport $\mathfrak{T}_{\mathfrak{g}}$ can be approximated by numerically integrating the corresponding second-order differential equations, see Section 2.

However, in high-dimensional settings, alternative strategies are often preferable for estimating parallel transport. In infinite-dimensional spaces, Schild's ladder [55, 119] and the Pole ladder [127] provide approximations of parallel transport, particularly when evaluating the exponential map and logarithm is computationally efficient, as is the case for Lie groups. When computing the logarithm is not tractable, the fanning scheme offers an alternative approach approximating parallel transport using only the exponential map [122, 202]. Its computational complexity scales as d^2 where d is the dimension.

The computation of the distance $\mathbf{d}_{\mathfrak{g}}(x, y)$ involves solving the minimization problem (1.1), which is equivalent to estimating a shortest geodesic given the endpoints $x, y \in \mathcal{M}$. This task is particularly challenging, and three general approaches can be employed for approximation:

1. **Shooting method.** The exponential map is computed by numerically integrating the geodesic equation and the distance is estimated as

$$\mathbf{d}_{\mathfrak{g}}(x, y) = \sqrt{\mathfrak{g}_x(\log_x(y), \log_x(y))} \quad (4.6)$$

by approximating the logarithm map:

$$\log_x(y) = \operatorname{argmin}_{v \in T_x \mathcal{M}} L(\text{Exp}_x(v), y), \quad (4.7)$$

where L is a loss function often called data-attachment term measuring the discrepancy between the points. Notably, this method also yields the logarithm $\log_x(y)$, which can be useful for representation tasks such as computing a Karcher mean. Although relatively straightforward to implement, its computational cost can be significant due to the integration of the geodesic equation and the need to differentiate its solution.

2. **Geodesic discretization.** Another approach involves computing (shortest) geodesics by discretizing in time the curve γ and its associated path energy, as defined in (2.4). This leads to

$$\left(\hat{\gamma} \left(\frac{k}{N} \right) \right)_{k \in [N]} = \operatorname{argmin}_{(x_k)_{k \in [N]}} \sum_{k=0}^N \tilde{\mathbf{d}}_{\mathfrak{g}}(x_k, x_{k+1})^2, \quad (4.8)$$

where $\tilde{\mathbf{d}}_{\mathfrak{g}}(x_k, x_{k+1})^2$ is a local approximation of the squared distance, e.g. $\tilde{\mathbf{d}}_{\mathfrak{g}}(x_k, x_{k+1})^2 \approx \mathfrak{g}_{x_k}(x_{k+1} - x_k, x_{k+1} - x_k)$. Here, N must be sufficiently large to ensure a realistic solution [7, 83] and the convergence of this approach for N going to infinity was investigated by Rumpf and Wirth [150]. The minimization can be performed in various ways. A recent fast and robust approach suggests using Gaussian processes [7] combined with fixed-point iterations to avoid the computational overhead of Jacobian processing. While such methods may be less precise than the shooting method, in some cases shooting methods are not feasible especially in high dimensions.

3. **Neural network regression.** When only a rough approximation of the distance or geodesic is required, a neural network $\phi_{\eta}: \mathcal{M} \times \mathcal{M} \times [0, 1] \rightarrow \mathcal{M}$ parameterized by η can be trained to produce geodesics between given points. Concretely, in [100, 170], geodesics are parameterized as

$$\gamma_{\eta}(x, y, t) = (1-t)x + ty + t(1-t)\phi_{\eta}(x, y, t), \quad (4.9)$$

where η is estimated by minimizing a stochastic versions of the path energy (2.4):

$$\eta \mapsto \mathbb{E}_{x, y \sim P_{\mathcal{D}} \otimes P_{\mathcal{D}}, t \sim \mathcal{U}([0, 1])} \left(\mathfrak{g}_{\gamma_{\eta}(x, y, t)}(\dot{\gamma}_{\eta}(x, y, t), \dot{\gamma}_{\eta}(x, y, t)) \right), \quad (4.10)$$

using stochastic gradient descent. Here, $P_{\mathcal{D}}$ represents the empirical distribution associated with noisy observations $\mathcal{D} = (x_i)_i$ on the manifold.

In practical applications, the metric is typically hand-crafted as in (4.3) to ensure that geodesics remain close to the manifold approximated by the dataset. This method has proven to be both practical and sufficiently accurate to outperform state-of-the-art techniques in trajectory inference [100, 170]. Additionally, it has been extended using splines in the context of optimal transport [146].

Pullback Metric With Diffeomorphisms. As mentioned earlier, if the pullback metric (4.4) is defined via a diffeomorphism f on the ambient space [47, 48, 49, 72], then all objects in $\mathcal{F}(\mathfrak{g})$ admit a closed-form expression as presented in [48]. This modeling approach is particularly useful when working in the observational space. However, estimating f becomes challenging in high-dimensional settings. For this reason, variational autoencoders (VAEs) are often preferred in high-dimensional spaces, although they are not diffeomorphic mappings, since the latent space has a lower dimension than the observational space.

Graph-based Metrics. When using a graph-based structure, geodesics are estimated via shortest path algorithms such as Dijkstra’s algorithm [23, 134, 143] or approximations based on the heat equation [84]. The same holds for distance computations. Parallel transport can be computed using ad hoc methods based on local Principal Component Analysis (PCA) [28]. Curvature is approximated using graph-equivalent definitions [59, 114].

4.2 Objectives

After having discussed the parametrization of the learned metric, we now turn to another central aspect of making the abstract problem (3.1) concrete: the choice of objective. It is, of course, the central ingredient to define *which* Riemannian metric we want to obtain in the end restricted by the space of possible metrics defined above and the minima achievable by the chosen optimization algorithm discussed in the next section.

The objective is usually chosen with some application of the metric in mind. We will thus first classify the objective broadly into two main categories: objectives based on regression and objectives based on classification. This follows, on the one hand, the very common classification of machine learning problems and, on the other hand, also mirrors the historical development of metric learning outlined in the introduction. There are some objectives that do not fit perfectly into these two categories and we will discuss them in a third subsection afterwards.

4.2.1 Classification

As explained before, the original motivation leading to the introduction of distance metric learning was to improve the performance of k -nearest neighbor classification. As the term *nearest* suggests, the distance between samples is a central ingredient and, thus, a wide range of objectives have been suggested that are based on this application. They can be carried over to Riemannian metric learning via the geodesic distance derived from the metric, which is the way will formulate them here. We will focus on a subset of important examples which will not be exhaustive as this would go beyond the scope of our review. For a more detailed overview on these losses, we refer to existing reviews focused on distance metric learning, such as the one by Kaya and Bilge [101].

Contrastive. The first set of objectives based on classification we consider are usually referred to as *contrastive losses* [37, 38, 46, 52, 61, 70, 74, 75, 77, 78, 80, 92, 93, 94, 102, 189, 191, 196]. Here, one assume two sets as input: a set $S_p \subset \mathcal{X} \times \mathcal{X}$ of examples that should be classified as similar (e.g. have the same label) and another set $S_n \subset \mathcal{X} \times \mathcal{X}$ of examples that should be classified as different. These sets, especially the latter, are typically assumed to contain “challenging” examples for a classification algorithm, *i.e.* such examples that are close to an envisioned decision

boundary (in the Euclidean norm). Hence, the objective in metric learning is to achieve a clearer separation of these examples, *i.e.* to *maximize* the distance for elements of S_n , while keeping similar points close, *i.e.* to *minimize* the distance for elements of S_p . Various different objectives have been proposed that formalize this idea, for example, as

$$\mathcal{L}(\mathbf{g}) = \sum_{(p,p') \in S_p} \mathbf{d}_{\mathbf{g}}(p, p') - \sum_{(q,q') \in S_n} \mathbf{d}_{\mathbf{g}}(q, q') \quad (4.11)$$

in [92] or as

$$\mathcal{L}(\mathbf{g}) = \sum_{(p,p') \in S_p} \mathbf{d}_{\mathbf{g}}(p, p') + \sum_{(q,q') \in S_n} \mathbf{d}_{\mathbf{g}^{-1}}(q, q'), \quad (4.12)$$

in [208] where the inverse of the (local) metric is considered for negative pairs. In a slightly different direction, also placing constraints on the distances for one set, e.g. demanding that the distance for all examples in S_p is below a certain threshold, while optimizing the distance for the other has been considered [183]. This amounts to replacing one of the sums in the above objectives with an indicator function which is zero if a pair fulfills the imposed threshold and is infinity otherwise. Various other variation of this idea have been proposed and we refer, for example, to [101] for a more complete overview.

Triplet. Another set of objectives, that is typically referred to as *triplet losses* [2, 31, 39, 52, 56, 82, 85, 89, 90, 118, 187, 205], assumes that similar information is given in the form triplets $S_t \subset \mathcal{X} \times \mathcal{X} \times \mathcal{X}$. Here, each triplet (q, q_p, q_n) contains an anchor point q , a positive example q_p to be classified as similar, and a negative example q_n to be classified as different. Based on these triplets, most commonly

$$\mathcal{L}(\mathbf{g}) = \sum_{(q, q_p, q_n) \in S_t} \mathbf{d}_{\mathbf{g}}(q, q_p) + \max(0, 1 + \mathbf{d}_{\mathbf{g}}(q, q_p) - \mathbf{d}_{\mathbf{g}}(q, q_n)) \quad (4.13)$$

is used as objective, e.g. in [192]. It encourages the positive example q_p to be close while pushing the negative example (also called impostor) q_n *further* away than q_p by a margin (which is one here). The idea behind this objective is to ensure a “clean” neighborhood for the anchors, *i.e.* such that every point with a different label is farther away than points with the same label as the anchor. Thus a (k-)nearest neighbor classifier would correctly classify the anchor based on its neighborhood. Again, over time, variations of this objective have been proposed and we refer, for example, to [101] for a more complete overview.

Others. Since classification-based objectives have been the main types of objective considered for distance metric learning, naturally many more have been developed over the course of time and we want to mention a few ideas motivating these other objectives. One interesting direction is to achieve unsupervised or at least semi-supervised objectives instead of the supervised approaches discussed above. Another direction is to model more non-local interaction between points, e.g. considering not just triplets but larger tuples or entire cluster of points, primarily to speed-up the convergence of descent algorithms, which otherwise might need a lot of iterations due to only local changes affected in each step [71, 131, 188]. A common challenge that most of these approaches share with the one we have explained is the need to sample points from the dataset in a specific way [101, 148]. For example, for the triplet loss, we need the anchors to be sufficiently densely sampled, and for each anchor, sufficiently many negative examples available to control their neighborhood [192].

4.2.2 Regression

In this second category of objectives, we assume to be given a more geometric type of supervision to learn the desired metric, *i.e.* observations of some continuous object derived from the metric instead of a classification. This could be scalar quantities, as e.g. the distance between points, geometric quantities such as trajectories on the manifold, or even just the dataset itself

describing its distribution. Below, we will discuss examples for these different types of supervisions. This second category of objectives is usually more closely connected to Riemannian metric learning in the literature, although, as argued before, this is not necessary conceptually.

Distribution. The most fundamental form of supervision is given just by (potentially noisy) samples $\mathcal{O} \subset \mathcal{M}$ on the manifold, which can be thought of as observations of its density. Thus, this first set of objectives broadly aims to learn a metric such that the corresponding volume form matches this density, which can be turned into concrete objectives in various ways.

To achieve this goal, Arvanitidis et al. [5] construct a diagonal metric where distances are smaller in areas with more samples and larger far away from them using the inverse of the local diagonal covariance matrices. They show empirically that a corresponding Riemannian normal distribution then reasonably approximates the data density. In [6], they extend this idea to an expected pull-back metric on the latent space of variational autoencoders, where the variance of the decoder is optimized to be large away from the given samples. Finally, in [8], they optimize a metric to give constant distances in dense areas while growing when moving away from them and combine this with the variance estimation of [6] to pull the metric back to the latent space of a variational autoencoder. This first part of this approach is also used by Kapusniak et al. [100].

Another approach is followed by Chadebec et al. [34, 35], who train a variational autoencoder using the common ELBO loss to approximate the data density. However, the distribution on the latent space is assumed to be given via a learned Riemannian metric and the sampling to evaluate the objective is performed using Riemannian Hamiltonian Monte Carlo. A similar idea is followed by Kalatzis et al. [99], who use a Riemannian Brownian motion as prior for the distribution in the latent space using a pulled-back metric. Toth et al. [176] and Sauty and Durrleman [156] combine the idea of approximating the data distribution using a variational autoencoder with approximating given trajectories as will be explained below.

Distance. Starting from distance observations, we assume having insights into the global geometric structure of our manifold and aim to derive a Riemannian structure matching it. Concretely, we assume to have a dataset $\mathcal{O} \subset \mathcal{M} \times \mathcal{M} \times \mathbb{R}_{\geq 0}$ consisting of pairs of points (p, q) on the manifold along with a (possibly noisy) observation d of the distance between them. Then the objective usually turns into a basic least squares estimation, *i.e.*

$$\mathcal{L}(\mathfrak{g}) = \sum_{(p,q,d) \in \mathcal{O}} |d - \mathbf{d}_{\mathfrak{g}}(p, q)|^2. \quad (4.14)$$

If one has a notion of density of the samples associated with them, *i.e.* some form of weights, these can of course also be integrated into the objective. This approach to regress distances has, for example, been investigated in [47, 48, 49, 63, 148].

Trajectories. For our last set of regression-based objectives, we assume that we are not just given observations of distance but of entire trajectories on the manifold which we want to model as geodesics under our learned metric. To this end, we assume that we are given K trajectories $(x_k^j)_{k \in [K], j \in [T_k]}$ at variable time points $(t_k^j)_{k=1 \in [K], j \in [T_k]}$ where T_k is the number of observations in trajectory k . Based on this data, we can introduce a formal objective for the regression problem via the set of all geodesics $\mathcal{T}_{\mathfrak{g}}$ for a given metric \mathfrak{g} as

$$\mathcal{L}(\mathfrak{g}) = \sum_{k=1}^K \inf_{\gamma \in \mathcal{T}_{\mathfrak{g}}} \sum_{j=1}^{T_k} \|x_k^j - \gamma(t_k^j)\|^2. \quad (4.15)$$

Since optimizing over the set of all geodesics for a given metric is rarely achievable this formulation is mostly of educational value. Hence, a central challenge is how to parametrize a sufficiently expressive subset of geodesic that allows to approximate a objective of the form (4.15) in a tractable manor.

To this end, Louis et al. [123] consider an autoencoder setup where the metric on the data manifold is given by the push-forward of the Euclidean metric on the latent space (cf. Sec-

tion 4.1). This means geodesics on the data manifold are given by the images of linear trajectories in the latent space, which are much more tractable to work with. Furthermore, they assume that all observed trajectories are parallel to each other (and a coordinate axis), which means that each trajectory is described by its constant speed, its starting time relative to the observations, and its offset to the coordinate axis. These quantities are then optimized together with the parameters of the autoencoder determining the Riemannian metric considering an objective of the form (4.15) and an additional regularization on them. This idea has been extended by Sauty and Durrleman [155] to variational autoencoders also modeling the data distribution (see above). Here, the regression objective is obtained by comparing points in the latent space whereas Louis et al. [123] compare points in the original data space after decoding.

This idea can be generalized to nonlinear *exp-parallel curves* [159] (see Section 3.2.2), which are not necessarily geodesics. Sauty and Durrleman [156] employ this idea while directly learning a metric on the data manifold. Thus, they approximate the geodesic and the parallel curves using a Hamiltonian shooting approach. In contrast, Gruffaz et al. [72] obtain the metric using a diffeomorphic deformation of a reference manifold (usually the Euclidean space) with known closed forms for the geodesic and parallel curves.

Aiming for a larger candidate set, Toth et al. [176] consider general geodesics in the latent space of a variational autoencoder parametrized by their starting point and direction which are used in a Hamiltonian shooting. Finally, Kushner et al. [109] assume that they are given coarsely sampled discrete trajectories and aim to upsample them using geodesics. Hence, they parametrize candidate geodesics by tangent vectors at given points on the manifold.

4.2.3 Others

In this last subsection, we discuss some works performing Riemannian metric learning with objectives that do not fall into the aforementioned categories.

Optimal Transport. As explained in Section 2.1, one can learn a metric on the space of probability measures on a manifold \mathcal{M} by learning a Riemannian metric on this base manifold. Hence, a range of works as focused on objectives for Riemannian metric learning that incorporates optimal transport based on it.

Considering the classification of histograms, Cuturi and Avis [45] adapt their objective from the contrastive losses in distance metric learning and arrive at the loss $\mathcal{L}(\mathfrak{g}) = \sum_{(\mu, \nu) \in S_p} W_2(\mu, \nu) - \sum_{(\mu, \nu) \in S_n} W_2(\mu, \nu)$, where W_2 is the Wasserstein distance induced by \mathfrak{g} as defined in (2.25). In contrast, Zhou et al. [206] minimize the Wasserstein distance between subsets of data points (representing different source domains in a domain adaption setup) while combining it with a triplet loss (see Section 4.2.1) over all data points. Likewise, Kerdoncuff et al. [103] optimize the Wasserstein distance between two sets of data points, one with labels and one without, with an additional penalty on the transport plan preventing two differently labeled points being sent to the same target. They also combine this with classification-based losses on the points with known labels.

Aiming towards the regression of measures, Scarvelis and Solomon [157] consider given pairs $(\mu_k, \nu_k)_k$ of probability distributions that are assumed to be consecutive times steps of (discrete) trajectories and aim to minimize their Wasserstein distance, *i.e.* $\mathcal{L}(\mathfrak{g}) = \sum_k W_1(\mu_k, \nu_k)$, where W_1 depends on \mathfrak{g} as before. To prevent the norm of the metric becoming arbitrarily small, they include an additional regularization penalizing the norm of the inverse metric tensor along linear interpolations between samples of the given distributions. Similarly, Heitz et al. [84] assume that the supervision for the metric learning comes in form of complete discrete trajectories of measures that should be well-approximated by optimal transport regression. Hence, they adapt the objective (4.15) to interpolations of measures with fixed given endpoints for the trajectories. This means they have to solve optimal transport interpolation problems to evaluate their objective for which they use entropy-regularized optimal transport [145].

Let us note that these objectives based on optimal transport could also be considered as Riemannian metric learning on the space of probability measure with the objectives introduced above. For example, the loss of Cuturi and Avis [45] is simply a contrastive loss from before with the Wasserstein distance. However, we chose to consider these objectives for the base metric based on optimal transport. This is because, to the best of our knowledge, there is thus far no viable approach to directly learn a metric on the space of probability measures forgoing the base metric on the one hand and, on the other hand, these objectives are sometimes combined with objectives that do not involve optimal transport.

Optimization. Finally, a few works focus on the use of the Riemannian metric in optimization applications and design objectives accordingly. For example, Bansal et al. [14] learn the metric using a bilevel optimization problem, where the learned metric is used in the lower level problem to assess the prediction accuracy of a regression model while the overall objective (*i.e.* upper level problem) is to optimize the performance of this model on a different task. Such problems arise, e.g., in reinforcement learning when the overall objective measures the quality of the model’s downstream value predictions while the regression problem concerns next state prediction.

In another direction, Cui et al. [44] minimize the Poincaré constant implied by the Riemannian metric for a given distribution on the manifold. The resulting metric is then used in Riemannian Langevin dynamics to achieve efficient sampling from the distribution as its convergence rate is given via the Poincaré constant.

4.3 Optimization Methods

Having chosen the set of metrics over which to optimize and the objective to be minimized, we have to also choose a practical algorithm to find an approximate solution. Over time, numerous approaches have been proposed for this and we give an overview of them in this last section on the technical aspects of Riemannian metric learning. Since the chosen optimization approach often amounts to the application of a standard algorithm, we will keep the overview brief and give some remarks about common things to note when applying the various algorithms to the metric learning problem.

Closed Form Formulas. The best-case scenario from a performance point of view is that we have a known closed form of the minimizer of the learning problem given the training data, e.g. in form of an eigendecomposition. This is, for example, the case for certain approaches using a constant metric and classification-based losses as introduced in the previous section as the problems are convex with a known global solution. This has been used in [75, 78, 80, 203, 208] among others. For other approaches, such as [5], this is the case as the metric is defined primarily to follow the density of the data and constructed from the beginning with a formula in mind. The metric can also be known if it is constructed via the pull-back along an otherwise constructed map, e.g. an autoencoder, which was used in [63, 175]. However, the number of problems that admit a known closed form solutions is small and especially it is rarely true for non-linear problems.

Gradient Descent Methods. Thus, in the absence of closed form solutions, most works turn towards iterative descent algorithms to solve the optimization problem. Here, the most fundamental approach is to apply gradient descent or its stochastic extension which was used in [14, 38, 70, 71, 102, 103, 139, 161, 195, 205, 206]. To speed-up the convergence of the descent, some works, such as [8, 34, 35, 100, 121, 154, 155, 156, 157, 176], use acceleration techniques such as Nesterov’s acceleration or the Adam algorithm. Others [84, 163] use L-BFGS to approximate the Hessian of the objective and achieve faster convergence.

These iterative methods typically face two challenges: (i) computing the gradient, and (ii) ensuring that the learned metric is positive definite for explicit parametrizations. The first challenge traditionally required manual computation of the gradient of the objective which has been mostly replaced by the usage of automatic differentiation frameworks in more recent approaches. Furthermore, the objective might not admit a classical gradient, e.g. if it involves solving an op-

timization problem as in the case of optimal transport-based objectives. Thus, some works explicitly use subgradient methods [45, 102, 139, 194]. The second challenge is often solved via the parametrization of the metric (see Section 4.1) and thus mostly comes up in the optimization for constant metrics. In these cases, it is typically [70, 195] solved by projected gradient descent, *i.e.* alternating a gradient step with projecting the metric matrix onto cone of positive definite matrices.

Structured Optimization. Instead of applying acceleration techniques, another range of works focuses more on the concrete structure of the considered optimization problem. For example, [46, 94, 183] apply Bregman iterations to the problem consisting of a convex objective with convex constraints. Alternatively, some problems can be reformulated as a semidefinite program which allows to use corresponding specialized optimization techniques, e.g. ones developed to train SVMs. This was used in [131, 185, 192, 207] for example. Furthermore, some methods, e.g. [92, 204], split the parameters of the metric into different subsets and then alternate optimizing over these subset, which is usually the case when for one subset a more structured optimization can be applied.

Another subset notably uses the Riemannian structure of the set of SPD matrices itself in the optimization by turning to Riemannian optimization methods. These methods use the Riemannian gradient of the objective in the tangent space of the current iterate and then apply an approximation of the exponential map to perform line search along it or a derived descent direction. This has been applied in [77, 93, 201] to avoid degeneracies of the metric and speed-up the convergence compared to methods such as projected gradient descent.

All of these methods require sufficient insight into the structure of the problem, either in terms of the objective or in terms of the parametrization of the metric. This is the reason why they have been applied mostly in approaches considering constant metrics or kernel methods. The general theory of Riemannian metric learning is thus far not developed enough to produce structured methods regardless of the chosen parametrization which can also applied, e.g., to neural network-based approaches, see also Section 5.

5 Future Directions

This section emphasizes promising avenues for future research in Riemannian metric learning, including:

- **Statistical and Optimization Guarantees for Nonlinear Metric Learning.** Developing rigorous statistical and optimization frameworks to provide theoretical guarantees and performance bounds for Riemannian metric learning methods.
- **Geometry-Aware Optimization Methods.** Designing optimization techniques that account for the intrinsic geometric properties of the space of Riemannian metrics, thereby improving convergence and efficiency.
- **Enhanced Software and Methodological Tools.** Creating new computational tools and algorithms aimed at improving the tractability, scalability, and practical usability of Riemannian metric learning approaches.
- **Discovering New Applications.** As Riemannian metric learning tools become more accessible, their applications will expand. In particular, leveraging Ricci curvature could enhance optimization, generalization, and sampling efficiency.

5.1 Theoretical Research Avenues

Theoretical researcher in statistics, optimization and geometry are challenged to contribute to Riemannian metric learning.

Statistical Guarantees for Nonlinear Metric Learning. Theoretical advancements in the field of metric learning remain relatively limited compared to the volume of empirical studies in the field. However, there is a notable recent trend toward theoretical and statistical investigations, particularly in the context of preference systems and metric learning [18, 33, 36, 106, 129, 184, 190].

In [18, 184], the authors establish foundational results, addressing robustness [18] and sample complexity [184] for learning Mahalanobis distances under various loss functions. Building on these works, Mason et al. [129] refine the sample complexity analysis by focusing on low-dimensional metrics, particularly under regularization constraints on the learned matrix and when employing a triplet loss.

More recently, Canal et al. [33], Chen et al. [36], Wang et al. [190] have explored the intersection of Mahalanobis metric learning and user preference modeling. In these studies, the objective is to jointly learn a Mahalanobis metric and user embeddings (u_j) in the space, where a user’s preference for x_1 over x_2 is encoded as $\|x_1 - u\| < \|x_2 - u\|$. Initial theoretical results were presented in [33] and later refined by [190], which introduced low-rank assumptions to improve the learning process. Furthermore, [36] extended these ideas to pluralistic alignment in fine-tuning large language models (LLMs), further motivating theoretical advancements in the field.

The next significant challenge lies in addressing the non-linear case of metric learning, such as using Riemannian metrics parameterized by neural networks or non-linear pushforward maps. Progress in this direction would benefit from the development of prior distributions and low-rank assumptions tailored to Riemannian metrics. For instance, Arvanitidis et al. [9] propose learning a prior in the latent space of a variational auto-encoder to derive a Riemannian metric that better accounts for latent data distribution. However, this approach is not yet well-supported by theoretical guarantees.

In the same way, some theory related to the optimization problem exists, but mainly for Mahalanobis distance metric learning [19, 185, 203], scarcely for Riemannian metric learning problems concerning the geodesic distance [23, 30, 148] and nearly nothing for the other Riemannian objects ($\text{Exp}_g, \mathfrak{T}_g, \text{vol}_g$) [121, Chapter 4-5]. Theoretical results on the optimization related to these objects would greatly enhance the understanding of the topic.

A deeper understanding of the space of Riemannian metrics [53] is essential for constructing relevant parameterizations [178] and for advancing the design of general priors. Moreover, the geometric structure of these spaces plays a crucial role in the design of effective optimization trajectories, underscoring the importance of integrating geometry into future methodological developments.

Geometry-Aware Optimization Methods. Numerous Riemannian optimization methods [107, 153] have been developed to account for the geometric structure of the space associated with the learnable parameters. Incorporating geometry into optimization can be compared to employing a Newton method rather than a first-order gradient descent as it leverages additional information about the underlying manifold.

However, these methods rely on the Riemannian manifolds being sufficiently simple to facilitate the implementation of key mathematical operations, such as projections onto the tangent space, the exponential map, and parallel transport. To the best of our knowledge, these operations are not yet widely available, cheap, or standardized for the space of Riemannian metrics [53, 67]. Further advancements on the choice of the metric on \mathcal{G}_M are anticipated to address these limitations, in the same way that the Log-Euclidean metric [4] improves the numerical tractability over the affine-invariant metric while still offering a useful structure.

5.2 Applied Research Avenues

Numerical and modeling challenges are central to achieve the widespread adoption of Riemannian metric learning, paving the way for more applied research directions.

Enhanced Software and Methodological Tools. Several libraries currently exist for statistics and optimization on Riemannian manifolds, such as Geomstats [133] and Geoopt [107], as well as libraries for distance metric learning, including pydml [168] and PyTorch Metric Learning [137], particularly in the context of deep learning. However, to date, no comprehensive library specifically tailored for Riemannian metric learning has been developed.

This gap is understandable, given the wide diversity of use cases in Riemannian metric learning, which complicates the design of general-purpose algorithms. Nonetheless, as outlined in Section 3, the associated optimization problem is inherently general. To be more specific, a library specialized in the computation of the Riemannian objects according to the metric parametrization $\mathcal{F}(\mathfrak{g}) = (\mathbf{d}_{\mathfrak{g}}, \text{Exp}_{\mathfrak{g}}, \mathfrak{T}_{\mathfrak{g}}, \text{vol}_{\mathfrak{g}})$ is to be waited. It is reasonable to anticipate that future research will focus on making these methods scalable (Section 4.1.3) akin to advancements observed in deep learning.

Discovering New Applications. As specialized and user-friendly libraries for Riemannian metric learning become available, the range of potential applications will naturally expand. Moreover, as the conceptual framework of Riemannian geometry becomes better understood, practitioners are likely to identify relevant applications in their respective fields. In particular, the Ricci curvature remains insufficiently explored and utilized, as highlighted in [59, 114]. Since the curvature of loss functions serves as an indicator of a model’s generalization capacity [174], Riemannian metric learning could contribute to improvements in both optimization and generalization, much as it has already proven beneficial for sampling [44].

6 Conclusion

All in all, this paper has highlighted the richness and potential of the Riemannian metric learning framework across various applications (Section 3.2). As we have seen, distance metric learning has helped to increase noticeably the performance of classification methods in the past. We begin to see similar effects of *Riemannian* metric learning in a range of applications, improving the performance in, e.g., generative modeling, trajectory inference, and causal inference. Since we have argued that tools descending from a Riemannian metric are applicable in many contexts in a flexible manner, we anticipate this list of successful applications will grow further as more researchers discover the possibilities of Riemannian metric learning. Thus, we hope that this review provides them with an accessible and motivating introduction and we encourage readers to explore new ways to leverage it in their respective fields.

Despite its promise, Riemannian metric learning still faces significant computational and theoretical challenges (Section 5). While there is substantial theoretical knowledge about the infinite-dimensional Riemannian manifold of Riemannian metrics, it is not connected to any statistical theory or similar for the metric learning problem. From a methodological perspective, research in distance metric learning is already well-established, providing a solid foundation upon which Riemannian metric learning could build (Sections 3.2.1 and 4.2.1). Providing a solid theoretical framework along with flexible computational tools will undoubtedly be demanding, but also offers exciting opportunities to further the adoption and impact.

In conclusion, we firmly believe that Riemannian metric learning represents an exciting new frontier in representation learning—an emerging Eldorado with vast, unexplored potential in both theory and practice.

References

- [1] Mohiuddin Ahmed, Raihan Seraj, and Syed Mohammed Shamsul Islam. The k-means algorithm: A comprehensive survey and performance evaluation. *Electronics*, 9(8):1295, 2020.
- [2] Mauro Annarumma and Giovanni Montana. Deep metric learning for multi-labelled radiographs. In *Proceedings of the 33rd annual ACM symposium on applied computing*, pages 34–37, 2018.
- [3] Sylvain Arguillère, Michael I. Miller, and Laurent Younes. Diffeomorphic surface registration with atrophy constraints. *SIAM Journal on Imaging Sciences*, 9(3):975–1003, 2016. doi: 10.1137/15M104431X.
- [4] Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. 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.
- [5] Georgios Arvanitidis, Lars K Hansen, and Søren Hauberg. A locally adaptive normal distribution. *Advances in Neural Information Processing Systems*, 29, 2016.
- [6] Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the curvature of deep generative models. In *International Conference on Learning Representations*, 2018.
- [7] Georgios Arvanitidis, Søren Hauberg, Philipp Hennig, and Michael Schober. Fast and robust shortest paths on manifolds learned from data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1506–1515. PMLR, 2019.
- [8] Georgios Arvanitidis, Søren Hauberg, and Bernhard Schölkopf. Geometrically Enriched Latent Spaces. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 631–639. PMLR, 2021.
- [9] Georgios Arvanitidis, Bogdan Georgiev, and Bernhard Schölkopf. A prior-based approximate latent riemannian metric. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 4634–4658. PMLR, 2022.
- [10] Christopher G Atkeson, Andrew W Moore, and Stefan Schaal. Locally weighted learning. *Lazy learning*, pages 11–73, 1997.
- [11] Mahsa Baktashmotlagh, Mehrtash T Harandi, Brian C Lovell, and Mathieu Salzmann. Unsupervised domain adaptation by domain invariant projection. In *Proceedings of the IEEE international conference on computer vision*, pages 769–776, 2013.
- [12] Mukund Balasubramanian and Eric L Schwartz. The isomap algorithm and topological stability. *Science*, 295(5552):7–7, 2002.
- [13] Dor Bank, Noam Koenigstein, and Raja Giryes. Autoencoders. *Machine learning for data science handbook: data mining and knowledge discovery handbook*, pages 353–374, 2023.
- [14] Dishank Bansal, Ricky TQ Chen, Mustafa Mukadam, and Brandon Amos. Taskmet: Task-driven metric learning for model learning. *Advances in Neural Information Processing Systems*, 36, 2024.

- [15] Martin Bauer, Philipp Harms, and Peter W. Michor. Sobolev metrics on shape space of surfaces. *Journal of Geometric Mechanics*, 3(4):389–438, 2011. ISSN 1941-4897. doi: 10.3934/jgm.2011.3.389.
- [16] H. Beik-Mohammadi, S. Hauberg, G. Arvanitidis, G. Neumann, and L. Rozo. Learning riemannian manifolds for geodesic motion skills. In *Robotics: Science and Systems XVII*, July 2021. doi: 10.15607/RSS.2021.XVII.082. URL <https://www.roboticsproceedings.org/rss17/p082.pdf>.
- [17] Peter N. Belhumeur, Joao P Hespanha, and David J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on pattern analysis and machine intelligence*, 19(7):711–720, 1997.
- [18] Aurélien Bellet and Amaury Habrard. Robustness and generalization for metric learning. *Neurocomputing*, 151:259–267, 2015.
- [19] Aurélien Bellet, Amaury Habrard, and Marc Sebban. A survey on metric learning for feature vectors and structured data. *arXiv preprint arXiv:1306.6709*, 2013.
- [20] Aurélien Bellet, Amaury Habrard, and Marc Sebban. *Metric learning*. Morgan & Claypool Publishers, 2015.
- [21] Thomas Bendokat, Ralf Zimmermann, and P-A Absil. A grassmann manifold handbook: Basic geometry and computational aspects. *Advances in Computational Mathematics*, 50(1):6, 2024.
- [22] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [23] Fethallah Benmansour, Guillaume Carlier, Gabriel Peyré, and Filippo Santambrogio. Derivatives with respect to metrics and applications: subgradient marching algorithm. *Numerische Mathematik*, 116:357–381, 2010.
- [24] Michael Betancourt. A conceptual introduction to Hamiltonian Monte Carlo. *arXiv preprint arXiv:1701.02434*, 2017.
- [25] Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023. doi: 10.1017/9781009166164. URL <https://www.nicolasboumal.net/book>.
- [26] Jennie E Brand, Xiang Zhou, and Yu Xie. Recent developments in causal inference and machine learning. *Annual Review of Sociology*, 49(1):81–110, 2023.
- [27] Michael M. Bronstein, Alexander M. Bronstein, Fabrice Michel, and Nikos Paragios. Data fusion through cross-modality metric learning using similarity-sensitive hashing. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3594–3601, 2010. doi: 10.1109/CVPR.2010.5539928.
- [28] Max Budninskiy, Gloria Yin, Leman Feng, Yiyong Tong, and Mathieu Desbrun. Parallel transport unfolding: A connection-based manifold learning approach. *SIAM Journal on Applied Algebra and Geometry*, 3(2):266–291, 2019.
- [29] Daniel Burkhardt, Jonathan Bloom, Robrecht Cannoodt, Malte D Luecken, Smita Krishnaswamy, Christopher Lance, Angela O Pisco, and Fabian J Theis. Multimodal single-cell integration across time, individuals, and batches. *NeurIPS Competitions*, 2022.

- [30] Giuseppe Buttazzo, Andrea Davini, Ilaria Fragala, and Fabricio Macia. Optimal riemannian distances preventing mass transfer. *Journal für die reine und angewandte Mathematik*, 2004. doi: 10.1515/crll.2004.077.
- [31] Shulian Cai, Jiabin Huang, Xinghao Ding, and Delu Zeng. Semantic edge detection based on deep metric learning. In *2017 International Symposium on Intelligent Signal Processing and Communication Systems (ISPACS)*, pages 707–712. IEEE, 2017.
- [32] Sylvain Calinon. Gaussians on riemannian manifolds: Applications for robot learning and adaptive control. *IEEE Robotics & Automation Magazine*, 27(2):33–45, 2020. doi: 10.1109/MRA.2020.2980548.
- [33] Gregory Canal, Blake Mason, Ramya Korlakai Vinayak, and Robert Nowak. One for all: Simultaneous metric and preference learning over multiple users. In *NeurIPS*, 2022.
- [34] Clément Chadebec, Clément Mantoux, and Stéphanie Allasonnière. Geometry-aware hamiltonian variational auto-encoder. *arXiv preprint arXiv:2010.11518*, 2020.
- [35] Clément Chadebec, Elina Thibeau-Sutre, Ninon Burgos, and Stéphanie Allasonnière. Data augmentation in high dimensional low sample size setting using a geometry-based variational autoencoder. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):2879–2896, 2022.
- [36] Daiwei Chen, Yi Chen, Aniket Rege, and Ramya Korlakai Vinayak. PAL: Pluralistic alignment framework for learning from heterogeneous preferences. In *NeurIPS 2024 Workshop on Behavioral Machine Learning*, 2024.
- [37] Shuo Chen, Chen Gong, Jian Yang, Xiang Li, Yang Wei, and Jun Li. Adversarial metric learning. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 2021–2027, 2018.
- [38] Shuo Chen, Lei Luo, Jian Yang, Chen Gong, Jun Li, and Heng Huang. Curvilinear distance metric learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [39] Xianhong Chen, Liang He, Can Xu, and Jia Liu. Distance-dependent metric learning. *IEEE Signal Processing Letters*, 26(2):357–361, 2019.
- [40] Zhijie Chen, Mingquan Feng, Junchi Yan, and Hongyuan Zha. Learning neural Hamiltonian dynamics: A methodological overview. *arXiv preprint arXiv:2203.00128*, 2022.
- [41] Marco Congedo, Alexandre Barachant, and Rajendra Bhatia. Riemannian geometry for eeg-based brain-computer interfaces; a primer and a review. *Brain-Computer Interfaces*, 4(3):155–174, 2017.
- [42] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.
- [43] Trevor F Cox and Michael AA Cox. *Multidimensional scaling*. CRC press, 2000.
- [44] Tiangang Cui, Xin Tong, and Olivier Zahm. Optimal Riemannian metric for Poincaré inequalities and how to ideally precondition Langevin dynamics. *arXiv preprint arXiv:2404.02554*, 2024.
- [45] Marco Cuturi and David Avis. Ground metric learning. *The Journal of Machine Learning Research*, 15(1):533–564, 2014.

- [46] Jason V Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S Dhillon. Information-theoretic metric learning. In *Proceedings of the 24th international conference on Machine learning*, pages 209–216, 2007.
- [47] Friso de Kruiff, Erik Bekkers, Ozan Öktem, Carola-Bibiane Schönlieb, and Willem Diepeveen. Pullback flow matching on data manifolds. *arXiv preprint arXiv:2410.04543*, 2024.
- [48] Willem Diepeveen. Pulling back symmetric Riemannian geometry for data analysis. *arXiv preprint arXiv:2403.06612*, 2024.
- [49] Willem Diepeveen, Georgios Batzolis, Zakhar Shumaylov, and Carola-Bibiane Schönlieb. Score-based pullback Riemannian geometry. *arXiv preprint arXiv:2410.01950*, 2024.
- [50] Manfredo Perdigao Do Carmo and J Flaherty Francis. *Riemannian geometry*, volume 2. Springer, 1992.
- [51] Ricardo Dominguez-Olmedo, Amir-Hossein Karimi, Georgios Arvanitidis, and Bernhard Schölkopf. On data manifolds entailed by structural causal models. In *International Conference on Machine Learning*, pages 8188–8201. PMLR, 2023.
- [52] Yueqi Duan, Wenzhao Zheng, Xudong Lin, Jiwen Lu, and Jie Zhou. Deep adversarial metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2780–2789, 2018.
- [53] David G Ebin. On the space of riemannian metrics. *Bulletin of the American Mathematical Society*, 74(5):1001–1003, 1968.
- [54] Alan Edelman, Tomás A Arias, and Steven T Smith. The geometry of algorithms with orthogonality constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.
- [55] JP EHLERS, FAE Pirani, and A Schild. The geometry of free fall and light propagation, [in:] general relativity-papers in honour of jl synge, ed. *L. O’Raifeartaigh, Clarendon Pres*, 1972.
- [56] Liat Ein-Dor, Yosi Mass, Alon Halfon, Elad Venezian, Ilya Shnayderman, Ranit Aharonov, and Noam Slonim. Learning thematic similarity metric using triplet networks. In *Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics (ACL), 2018.
- [57] Albert Einstein et al. Zur elektrodynamik bewegter körper. *Annalen der physik*, 17(10): 891–921, 1905.
- [58] Theodoros Evgeniou and Massimiliano Pontil. Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 109–117, 2004.
- [59] Amirhossein Farzam, Allen Tannenbaum, and Guillermo Sapiro. From geometry to causality-ricci curvature and the reliability of causal inference on networks. In *Forty-first International Conference on Machine Learning*, 2024.
- [60] Oluwadamilola Fasina, Guillaume Huguet, Alexander Tong, Yanlei Zhang, Guy Wolf, Maximilian Nickel, Ian Adelstein, and Smita Krishnaswamy. Neural fim for learning fisher information metrics from point cloud data. In *International Conference on Machine Learning*, pages 9814–9826. PMLR, 2023.

- [61] Alireza Fathi, Zbigniew Wojna, Vivek Rathod, Peng Wang, Hyun Oh Song, Sergio Guadarrama, and Kevin P Murphy. Semantic instance segmentation via deep metric learning. *arXiv preprint arXiv:1703.10277*, 2017.
- [62] Nemo Fournier and Stanley Durrleman. A multimodal disease progression model for genetic associations with disease dynamics. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 601–610. Springer, 2023.
- [63] Max F Frenzel, Bogdan Teleaga, and Asahi Ushio. Latent space cartography: Generalised metric-inspired measures and measure-based transformations for generative models. *arXiv preprint arXiv:1902.02113*, 2019.
- [64] Jerome H Friedman et al. Flexible metric nearest neighbor classification. Technical report, Citeseer, 1994.
- [65] Thibaut Germain, Samuel Gruffaz, Charles Truong, Alain Durmus, and Laurent Oudre. Shape analysis for time series. In *Advances in Neural Information Processing Systems*, volume 37, 2024.
- [66] Benyamin Ghogh, Ali Ghodsi, Fakhri Karray, and Mark Crowley. Spectral, probabilistic, and deep metric learning: Tutorial and survey. *arXiv preprint arXiv:2201.09267*, 2022.
- [67] Olga Gil-Medrano and Peter W. Michor. The Riemannian manifold of all Riemannian metrics. *The Quarterly Journal of Mathematics*, 42(1):183–202, 1991. doi: 10.1093/qmath/42.1.183.
- [68] Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 73(2):123–214, 2011.
- [69] Joan Glaunès, Alain Trouvé, and Laurent Younes. Modeling planar shape variation via hamiltonian flows of curves. *Statistics and analysis of shapes*, pages 335–361, 2006.
- [70] Amir Globerson and Sam Roweis. Metric learning by collapsing classes. *Advances in neural information processing systems*, 18, 2005.
- [71] Jacob Goldberger, Geoffrey E Hinton, Sam Roweis, and Russ R Salakhutdinov. Neighbourhood components analysis. *Advances in neural information processing systems*, 17, 2004.
- [72] Samuel Gruffaz, Pierre-Emmanuel Poulet, Etienne Maheux, Bruno Jedynak, and Stanley Durrleman. Learning riemannian metric for disease progression modeling. *Advances in Neural Information Processing Systems*, 34:23780–23792, 2021.
- [73] Nicolas Guigui and Xavier Pennec. Parallel transport, a central tool in geometric statistics for computational anatomy: Application to cardiac motion modeling. In *Handbook of Statistics*, volume 46, pages 285–326. Elsevier, 2022.
- [74] Matthieu Guillaumin, Jakob Verbeek, and Cordelia Schmid. Multiple instance metric learning from automatically labeled bags of faces. In *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part I 11*, pages 634–647. Springer, 2010.
- [75] Jihun Hamm and Daniel D Lee. Grassmann discriminant analysis: a unifying view on subspace-based learning. In *Proceedings of the 25th international conference on Machine learning*, pages 376–383, 2008.

- [76] Mehrtaash Harandi, Conrad Sanderson, Chunhua Shen, and Brian C. Lovell. Dictionary learning and sparse coding on grassmann manifolds: An extrinsic solution. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2013.
- [77] Mehrtaash Harandi, Mathieu Salzmann, and Richard Hartley. Dimensionality reduction on spd manifolds: The emergence of geometry-aware methods. *IEEE transactions on pattern analysis and machine intelligence*, 40(1):48–62, 2017.
- [78] Mehrtaash T Harandi, Conrad Sanderson, Sareh Shirazi, and Brian C Lovell. Graph embedding discriminant analysis on grassmannian manifolds for improved image set matching. In *CVPR 2011*, pages 2705–2712. IEEE, 2011.
- [79] Emmanuel Hartman, Yashil Sukurdeep, Eric Klassen, Nicolas Charon, and Martin Bauer. Elastic shape analysis of surfaces with second-order sobolev metrics: a comprehensive numerical framework. *International Journal of Computer Vision*, 131(5):1183–1209, 2023.
- [80] Søren Hauberg, Oren Freifeld, and Michael Black. A geometric take on metric learning. *Advances in Neural Information Processing Systems*, 25, 2012.
- [81] James A Hay, Lee Kennedy-Shaffer, Sanjat Kanjilal, Niall J Lennon, Stacey B Gabriel, Marc Lipsitch, and Michael J Mina. Estimating epidemiologic dynamics from cross-sectional viral load distributions. *Science*, 373(6552):eabh0635, 2021.
- [82] Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, and Xiang Bai. Triplet-center loss for multi-view 3d object retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1945–1954, 2018.
- [83] Behrend Heeren, Martin Rumpf, Peter Schröder, Max Wardetzky, and Benedikt Wirth. Exploring the geometry of the space of shells. *Computer Graphics Forum*, 33(5):247–256, 2014.
- [84] Matthieu Heitz, Nicolas Bonneel, David Coeurjolly, Marco Cuturi, and Gabriel Peyré. Ground metric learning on graphs. *Journal of Mathematical Imaging and Vision*, 63:89–107, 2021.
- [85] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.
- [86] Geoffrey E Hinton and Sam Roweis. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15, 2002.
- [87] Morris W Hirsch. Immersions of manifolds. *Transactions of the American Mathematical Society*, 93(2):242–276, 1959.
- [88] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [89] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. Collaborative metric learning. In *Proceedings of the 26th international conference on world wide web*, pages 193–201, 2017.
- [90] Junlin Hu, Jiwen Lu, and Yap-Peng Tan. Deep metric learning for visual tracking. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(11):2056–2068, 2015.
- [91] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008.

- [92] Zhiwu Huang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Learning euclidean-to-riemannian metric for point-to-set classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1677–1684, 2014.
- [93] Zhiwu Huang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Projection metric learning on grassmann manifold with application to video based face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 140–149, 2015.
- [94] Zhiwu Huang, Ruiping Wang, Shiguang Shan, Xianqiu Li, and Xilin Chen. Log-euclidean metric learning on symmetric positive definite manifold with application to image set classification. In *International conference on machine learning*, pages 720–729. PMLR, 2015.
- [95] Guillaume Hugué, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.
- [96] Pratik Jawanpuria, Dai Shi, Bamdev Mishra, and Junbin Gao. A Riemannian Approach to Ground Metric Learning for Optimal Transport. *arXiv preprint arXiv:2409.10085*, 2024.
- [97] Seungwoo Jeong, Wonsik Jung, Junghyo Sohn, and Heung-Il Suk. Deep geometric learning with monotonicity constraints for alzheimer’s disease progression. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
- [98] Jean Kaddour, Aengus Lynch, Qi Liu, Matt J Kusner, and Ricardo Silva. Causal machine learning: A survey and open problems. *arXiv preprint arXiv:2206.15475*, 2022.
- [99] Dimitris Kalatzis, David Eklund, Georgios Arvanitidis, and Søren Hauberg. Variational autoencoders with riemannian brownian motion priors. In *Proceedings of the 37th International Conference on Machine Learning, ICML’20*, 2020. doi: 10.5555/3524938.3525407.
- [100] Kacper Kapusniak, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Avishek Joey Bose, and Francesco Di Giovanni. Metric Flow Matching for Smooth Interpolations on the Data Manifold. In *Advances in Neural Information Processing Systems*, volume 37, pages 135011–135042. Curran Associates, Inc., 2024.
- [101] Mahmut Kaya and Hasan Şakir Bilge. Deep metric learning: A survey. *Symmetry*, 11(9): 1066, 2019.
- [102] Dor Kedem, Stephen Tyree, Fei Sha, Gert Lanckriet, and Kilian Q Weinberger. Non-linear metric learning. *Advances in neural information processing systems*, 25, 2012.
- [103] Tanguy Kerdoncuff, Rémi Emonet, and Marc Sebban. Metric learning in optimal transport for domain adaptation. In *International Joint Conference on Artificial Intelligence*, 2021.
- [104] Minyoung Kim, Sanjiv Kumar, Vladimir Pavlovic, and Henry Rowley. Face tracking and recognition with visual constraints in real-world videos. In *2008 IEEE Conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2008.
- [105] Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019.
- [106] Matthäus Kleindessner and Ulrike Luxburg. Uniqueness of ordinal embedding. In *Conference on Learning Theory*, pages 40–67. PMLR, 2014.

- [107] Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geopt: Riemannian optimization in pytorch. *arXiv preprint arXiv:2005.02819*, 2020.
- [108] Brian Kulis et al. Metric learning: A survey. *Foundations and Trends® in Machine Learning*, 5(4):287–364, 2013.
- [109] Sarah Kushner, Vismay Modi, and Nina Miolane. Learning riemannian metrics for interpolating animations.
- [110] Robert J LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American economic review*, pages 604–620, 1986.
- [111] Hugo Lavenant, Stephen Zhang, Young-Heon Kim, and Geoffrey Schiebinger. Toward a mathematical theory of trajectory inference. *The Annals of Applied Probability*, 34(1A), February 2024. doi: 10.1214/23-AAP1969.
- [112] John M Lee. *Riemannian manifolds: an introduction to curvature*, volume 176. Springer Science & Business Media, 2006.
- [113] Bastian Leibe and Bernt Schiele. Analyzing appearance and contour based methods for object categorization. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–409. IEEE, 2003.
- [114] Yangyang Li, Chaoqun Fei, Chuanqing Wang, Hongming Shan, and Ruqian Lu. Geometry flow-based deep riemannian metric learning. *IEEE/CAA Journal of Automatica Sinica*, 10(9):1882–1892, 2023.
- [115] Zicheng Liao, Jason Rock, Yang Wang, and David Forsyth. Non-parametric filtering for geometric detail extraction and material representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 963–970, 2013.
- [116] Tong Lin and Hongbin Zha. Riemannian manifold learning. *IEEE transactions on pattern analysis and machine intelligence*, 30(5):796–809, 2008.
- [117] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.
- [118] Yiding Liu, Kaiqi Zhao, and Gao Cong. Efficient similar region search with deep metric learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1850–1859, 2018.
- [119] Marco Lorenzi, Nicholas Ayache, and Xavier Pennec. Schild’s ladder for the parallel transport of deformations in time series of images. In *Biennial international conference on information processing in medical imaging*, pages 463–474. Springer, 2011.
- [120] Marco Lorenzi, Xavier Pennec, Giovanni B Frisoni, Nicholas Ayache, Alzheimer’s Disease Neuroimaging Initiative, et al. Disentangling normal aging from alzheimer’s disease in structural magnetic resonance images. *Neurobiology of aging*, 36:S42–S52, 2015.
- [121] Maxime Louis. *Computational and statistical methods for trajectory analysis in a Riemannian geometry setting*. PhD thesis, Sorbonne Université, 2019.
- [122] Maxime Louis, Alexandre Bône, Benjamin Charlier, Stanley Durrleman, and Alzheimer’s Disease Neuroimaging Initiative. Parallel transport in shape analysis: a scalable numerical scheme. In *Geometric Science of Information: Third International Conference, GSI 2017, Paris, France, November 7-9, 2017, Proceedings 3*, pages 29–37. Springer, 2017.

- [123] Maxime Louis, Raphaël Couronné, Igor Koval, Benjamin Charlier, and Stanley Durrleman. Riemannian geometry learning for disease progression modelling. In *Information Processing in Medical Imaging: 26th International Conference, IPMI 2019, Hong Kong, China, June 2–7, 2019, Proceedings 26*, pages 542–553. Springer, 2019.
- [124] Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. *Advances in neural information processing systems*, 30, 2017.
- [125] Rui Lu, Kailun Wu, Zhiyao Duan, and Changshui Zhang. Deep ranking: Triplet matchnet for music metric learning. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 121–125. IEEE, 2017.
- [126] Evan Z Macosko, Anindita Basu, Rahul Satija, James Nemesh, Karthik Shekhar, Melissa Goldman, Itay Tirosh, Allison R Bialas, Nolan Kamitaki, Emily M Martersteck, et al. Highly parallel genome-wide expression profiling of individual cells using nanoliter droplets. *Cell*, 161(5):1202–1214, 2015.
- [127] Lorenzi Marco and Xavier Pennec. Parallel transport with pole ladder: Application to deformations of time series of images. In *International Conference on Geometric Science of Information*, pages 68–75. Springer, 2013.
- [128] Razvan V Marinescu, Neil P Oxtoby, Alexandra L Young, Esther E Bron, Arthur W Toga, Michael W Weiner, Frederik Barkhof, Nick C Fox, Stefan Klein, Daniel C Alexander, et al. TADPOLE challenge: Prediction of longitudinal evolution in Alzheimer’s disease. *arXiv preprint arXiv:1805.03909*, 2018.
- [129] Blake Mason, Lalit Jain, and Robert Nowak. Learning low-dimensional metrics. *Advances in neural information processing systems*, 30, 2017.
- [130] Simon Matabaigne, P. A. Absil, and Nina Miolane. Bounds on the geodesic distances on the Stiefel manifold for a family of Riemannian metrics. *arXiv preprint arXiv:2408.07072*, 2024.
- [131] Brian McFee, Gert Lanckriet, and Tony Jebara. Learning multi-modal similarity. *Journal of machine learning research*, 12(2), 2011.
- [132] Marina Meilă and Hanyu Zhang. Manifold learning: What, how, and why. *Annual Review of Statistics and Its Application*, 11, 2024.
- [133] Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, Yann Thanwerdas, Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti, et al. Geomstats: a python package for riemannian geometry in machine learning. *Journal of Machine Learning Research*, 21(223):1–9, 2020.
- [134] Jean-Marie Mirebeau and Johann Dreo. Automatic differentiation of non-holonomic fast marching for computing most threatening trajectories under sensors surveillance. In *International Conference on Geometric Science of Information*, pages 791–800. Springer, 2017.
- [135] Deen Dayal Mohan, Bhavin Jawade, Srirangaraj Setlur, and Venu Govindaraju. Deep metric learning for computer vision: A brief overview. *Handbook of Statistics*, 48:59–79, 2023.
- [136] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A metric learning reality check. In *Computer Vision–ECCV2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 681–699. Springer, 2020.

- [137] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. Pytorch metric learning. *arXiv preprint arXiv:2008.09164*, 2020.
- [138] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.
- [139] Shibin Parameswaran and Kilian Q Weinberger. Large margin multi-task metric learning. *Advances in neural information processing systems*, 23, 2010.
- [140] Judea Pearl. Graphs, causality, and structural equation models. *Sociological Methods & Research*, 27(2):226–284, 1998.
- [141] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3, 2009. ISSN 1935-7516. doi: 10.1214/09-SS057.
- [142] Paraskevas Pegios, Aasa Feragen, Andreas Abildtrup Hansen, and Georgios Arvanitidis. Counterfactual explanations via riemannian latent space traversal. *arXiv preprint arXiv:2411.02259*, 2024.
- [143] Jaakko Peltonen, Arto Klami, and Samuel Kaski. Improved learning of riemannian metrics for exploratory analysis. *Neural Networks*, 17(8-9):1087–1100, 2004.
- [144] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. *International Journal of computer vision*, 66:41–66, 2006.
- [145] Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- [146] Aram-Alexandre Pooladian, Carles Domingo-Enrich, Ricky TQ Chen, and Brandon Amos. Neural Optimal Transport with Lagrangian Costs. In *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence*, volume 244 of *Proceedings of Machine Learning Research*, pages 2989–3003. PMLR, 2024.
- [147] Melanie F Pradier and Javier González. Beyond flatland: A geometric take on matching methods for treatment effect estimation. *arXiv preprint arXiv:2409.05459*, 2024.
- [148] Jiaming Qiu and Xiongtao Dai. Estimating riemannian metric with noise-contaminated intrinsic distance. *Advances in Neural Information Processing Systems*, 36, 2024.
- [149] Karsten Roth, Timo Milbich, Samarth Sinha, Prateek Gupta, Bjorn Ommer, and Joseph Paul Cohen. Revisiting training strategies and generalization performance in deep metric learning. In *International Conference on Machine Learning*, pages 8242–8252. PMLR, 2020.
- [150] Martin Rumpf and Benedikt Wirth. Variational time discretization of geodesic calculus. *IMA Journal of Numerical Analysis*, 35(3):1011–1046, 2015. doi: 10.1093/imanum/dru027.
- [151] Takashi Sakai. *Riemannian geometry*, volume 149. American Mathematical Soc., 1996.
- [152] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55 (58-63):94, 2015.
- [153] Hiroyuki Sato. *Riemannian optimization and its applications*, volume 670. Springer, 2021.
- [154] Benoît Sauty. *Multimodal modelling of Alzheimer’s Disease progression*. PhD thesis, Sorbonne Université, 2023.

- [155] Benoît Sauty and Stanley Durrleman. Longitudinal variational autoencoders learn a riemannian progression model for imaging data. In *GeoMedIA Workshop 2022-Geometric Deep Learning in Medical Image Analysis*, 2022.
- [156] Benoît Sauty and Stanley Durrleman. Riemannian metric learning for progression modeling of longitudinal datasets. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022.
- [157] Christopher Scarvelis and Justin Solomon. Riemannian metric learning via optimal transport. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=v3y68gz-WEz>.
- [158] Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- [159] Jean-Baptiste Schiratti, Stéphanie Allasonnière, Olivier Colliot, and Stanley Durrleman. A bayesian mixed-effects model to learn trajectories of changes from repeated manifold-valued observations. *Journal of Machine Learning Research*, 18(133):1–33, 2017.
- [160] Bernhard Schölkopf. Causality for machine learning. In *Probabilistic and causal inference: The works of Judea Pearl*, pages 765–804. Association for Computing Machinery, 2022.
- [161] Shai Shalev-Shwartz, Yoram Singer, and Andrew Y Ng. Online and batch learning of pseudo-metrics. In *Proceedings of the twenty-first international conference on Machine learning*, page 94, 2004.
- [162] Lavanya Sharan, Ruth Rosenholtz, and Edward Adelson. Material perception: What can you see in a brief glance? *Journal of Vision*, 9(8):784–784, 2009.
- [163] Zhengyang Shen, François-Xavier Vialard, and Marc Niethammer. Region-specific diffeomorphic metric mapping. *Advances in Neural Information Processing Systems*, 32, 2019.
- [164] R Short and Keinosuke Fukunaga. The optimal distance measure for nearest neighbor classification. *IEEE transactions on Information Theory*, 27(5):622–627, 1981.
- [165] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.
- [166] Stefan Sommer, Tom Fletcher, and Xavier Pennec. Introduction to differential and riemannian geometry. In *Riemannian Geometric Statistics in Medical Image Analysis*, pages 3–37. Elsevier, 2020.
- [167] Vincent Souveton, Arnaud Guillin, Jens Jasche, Guilhem Lavaux, and Manon Michel. Fixed-kinetic Neural Hamiltonian Flows for enhanced interpretability and reduced complexity. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 3178–3186. PMLR, 2024.
- [168] Juan Luis Suárez, Salvador García, and Francisco Herrera. pydml: A python library for distance metric learning. *Journal of Machine Learning Research*, 21(96):1–7, 2020.
- [169] Juan Luis Suárez, Salvador García, and Francisco Herrera. A tutorial on distance metric learning: Mathematical foundations, algorithms, experimental analysis, prospects and challenges. *Neurocomputing*, 425:300–322, 2021.

- [170] Xingzhi Sun, Danqi Liao, Kincaid MacDonald, Yanlei Zhang, Chen Liu, Guillaume Huguet, Guy Wolf, Ian Adelstein, Tim GJ Rudner, and Smita Krishnaswamy. Geometry-Aware Generative Autoencoder for Warped Riemannian Metric Learning and Generative Modeling on Data Manifolds. In *The 28th International Conference on Artificial Intelligence and Statistics*, 2025.
- [171] Stas Syrota, Pablo Moreno-Muñoz, and Søren Hauberg. Decoder ensembling for learned latent geometries. In *Proceedings of the Geometry-grounded Representation Learning and Generative Modeling Workshop (GRaM)*, volume 251 of *Proceedings of Machine Learning Research*, pages 277–285. PMLR, 2024.
- [172] Tran Tien Tam, Ines Adouani, and C Samir. *Transfer Learning on the Riemannian Manifold of Probability Measures*. PhD thesis, University of Clermont Auvergne, 2023.
- [173] Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [174] Valentin Thomas, Fabian Pedregosa, Bart Merriënboer, Pierre-Antoine Manzagol, Yoshua Bengio, and Nicolas Le Roux. On the interplay between noise and curvature and its effect on optimization and generalization. In *International Conference on Artificial Intelligence and Statistics*, pages 3503–3513. PMLR, 2020.
- [175] Alessandra Tosi, Søren Hauberg, Alfredo Vellido, and Neil D Lawrence. Metrics for probabilistic geometries. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, UAI’14, pages 800–808. AUAI Press, 2014.
- [176] Peter Toth, Danilo Jimenez Rezende, Andrew Jaegle, Sébastien Racanière, Aleksandar Botev, and Irina Higgins. Hamiltonian Generative Networks. In *International Conference on Learning Representations*, 2019.
- [177] Alain Trouvé. Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision*, 28:213–221, 1998. doi: 10.1023/A:1008001603737.
- [178] Wilderich Tuschmann and David J Wraith. *Moduli spaces of Riemannian metrics*, volume 46. Springer, 2015.
- [179] Marc Vaillant, Michael I Miller, Laurent Younes, and Alain Trouvé. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23:S161–S169, 2004.
- [180] Vivien van Veldhuizen. Geometry-aware latent representation learning for modeling disease progression of barrett’s esophagus. *arXiv preprint arXiv:2303.12711*, 2023.
- [181] Valerio Varano, Stefano Gabriele, Luciano Teresi, Ian L Dryden, Paolo E Puddu, Concetta Torromeo, and Paolo Piras. The tps direct transport: a new method for transporting deformations in the size-and-shape space. *International Journal of Computer Vision*, 124: 384–408, 2017.
- [182] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [183] Raviteja Vemulapalli and David W Jacobs. Riemannian metric learning for symmetric positive definite matrices. *arXiv preprint arXiv:1501.02393*, 2015.
- [184] Nakul Verma and Kristin Branson. Sample complexity of learning mahalanobis distance metrics. *Advances in neural information processing systems*, 28, 2015.

- [185] Faqiang Wang, Wangmeng Zuo, Lei Zhang, Deyu Meng, and David Zhang. A kernel classification framework for metric learning. *IEEE transactions on neural networks and learning systems*, 26(9):1950–1962, 2014.
- [186] Fei Wang and Jimeng Sun. Survey on distance metric learning and dimensionality reduction in data mining. *Data mining and knowledge discovery*, 29(2):534–564, 2015.
- [187] Jixuan Wang, Kuan-Chieh Wang, Marc T Law, Frank Rudzicz, and Michael Brudno. Centroid-based deep metric learning for speaker recognition. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3652–3656. IEEE, 2019.
- [188] Qi Wang, Jia Wan, and Yuan Yuan. Deep metric learning for crowdedness regression. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(10):2633–2643, 2017.
- [189] Xueping Wang and Min Liu. Multi-view deep metric learning for volumetric image recognition. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2018.
- [190] Zhi Wang, Geelon So, and Ramya Korlakai Vinayak. Metric learning from limited pairwise preference comparisons. In *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence*, volume 244 of *Proceedings of Machine Learning Research*, pages 3571–3602. PMLR, 2024.
- [191] Frederik Warburg, Marco Miani, Silas Brack, and Søren Hauberg. Bayesian metric learning for uncertainty quantification in image retrieval. *Advances in Neural Information Processing Systems*, 36, 2024.
- [192] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of machine learning research*, 10(2), 2009.
- [193] Lior Wolf, Tal Hassner, and Itay Maoz. Face recognition in unconstrained videos with matched background similarity. In *CVPR 2011*, pages 529–534. IEEE, 2011.
- [194] Pengtao Xie and Eric P. Xing. Multi-modal distance metric learning. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pages 1806–1812. AAAI Press, 2013.
- [195] Eric Xing, Michael Jordan, Stuart J Russell, and Andrew Ng. Distance metric learning with application to clustering with side-information. *Advances in neural information processing systems*, 15, 2002.
- [196] Yinghui Xing, Min Wang, Shuyuan Yang, and Licheng Jiao. Pan-sharpening via deep metric learning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 145:165–183, 2018.
- [197] Jie Xu, Lei Luo, Cheng Deng, and Heng Huang. Multi-level metric learning via smoothed wasserstein distance. In *IJCAI*, pages 2919–2925, 2018.
- [198] Liu Yang and Rong Jin. Distance metric learning: A comprehensive survey. *Michigan State University*, 2(2):4, 2006.
- [199] Ke Ye and Lek-Heng Lim. Schubert varieties and distances between subspaces of different dimensions. *SIAM Journal on Matrix Analysis and Applications*, 37(3):1176–1197, 2016. doi: 10.1137/15M1054201. URL <https://doi.org/10.1137/15M1054201>.
- [200] Zhandos Yessenbayev and Zhanibek Kozhirkbayev. Use of riemannian distance metric to verify topological similarity of acoustic and text domains. In *International Conference on Artificial Neural Networks*, pages 368–380. Springer, 2024.

- [201] Florian Yger and Masashi Sugiyama. Supervised logeuclidean metric learning for symmetric positive definite matrices. *arXiv preprint arXiv:1502.03505*, 2015.
- [202] Laurent Younes. Jacobi fields in groups of diffeomorphisms and applications. *Quarterly of applied mathematics*, pages 113–134, 2007.
- [203] Pourya Zadeh, Reshad Hosseini, and Suvrit Sra. Geometric mean metric learning. In *International conference on machine learning*, pages 2464–2471. PMLR, 2016.
- [204] Xiaohua Zhai, Yuxin Peng, and Jianguo Xiao. Heterogeneous metric learning with joint graph regularization for cross-media retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, pages 1198–1204, 2013.
- [205] Hangbin Zhang, Raymond K Wong, and Victor W Chu. Curvilinear collaborative metric learning with macro-micro attentions. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021.
- [206] Fan Zhou, Zhuqing Jiang, Changjian Shui, Boyu Wang, and Brahim Chaib-draa. Domain generalization via optimal transport with metric similarity learning. *Neurocomputing*, 456:469–480, 2021.
- [207] Pengfei Zhu, Lei Zhang, Wangmeng Zuo, and David Zhang. From point to set: Extend the learning of distance metrics. In *Proceedings of the IEEE international conference on computer vision*, pages 2664–2671, 2013.
- [208] Pengfei Zhu, Hao Cheng, Qinghua Hu, Qilong Wang, and Changqing Zhang. Towards generalized and efficient metric learning on riemannian manifold. In *IJCAI*, pages 3235–3241, 2018.

A Some Rigorous Aspects of Riemannian Geometry

Manifold. We begin with the basic definitions introducing our objects of study.

Definition 17 (Differentiable manifold). *A differentiable manifold \mathcal{M} of dimension $d < \infty$ is a set together with a family of injective maps $\psi_\alpha: U_\alpha \subset \mathbb{R}^d \rightarrow \mathcal{M}$ of open sets U_α of \mathbb{R}^d into \mathcal{M} such that*

1. $\cup_\alpha \psi_\alpha(U_\alpha) = \mathcal{M}$, and
2. for any pair α, β with $\psi_\alpha(U_\alpha) \cap \psi_\beta(U_\beta) = W \neq \emptyset$, the sets $\psi_\alpha^{-1}(W)$ and $\psi_\beta^{-1}(W)$ are open in \mathbb{R}^d and the map $\psi_\beta^{-1} \circ \psi_\alpha$ is differentiable.

The pair (U_α, ψ_α) with $p \in \psi_\alpha(U_\alpha)$ is called a parametrization of \mathcal{M} at p .

Typically one assumes that the differentiable structure in Definition 17 is maximal with respect to the given conditions. This can be achieved by extending the structure by parametrizations compatible with condition (2).

Definition of Covariant Derivative.

Definition 18. *Let \mathcal{M} be a smooth manifold, V a smooth vector field on \mathcal{M} , and W . The covariant derivative of W with respect to V , denoted by $\nabla_V W$, is defined as a map:*

$$\nabla: C^\infty(\mathcal{M}, T\mathcal{M}) \times C^\infty(\mathcal{M}, T\mathcal{M}) \rightarrow C^\infty(\mathcal{M}, T\mathcal{M}),$$

where $C^\infty(\mathcal{M}, T\mathcal{M})$ is the space of smooth vector fields on \mathcal{M} . A covariant derivative satisfies the following properties:

1. **Linearity:** For any smooth vector fields $Z, V, W \in C^\infty(\mathcal{M}, T\mathcal{M})$ and $\alpha \in \mathbb{R}$,

$$\nabla_Z(V + \alpha W) = \nabla_Z V + \alpha \nabla_Z W.$$

2. **Product Rule:** For any vector field $V, W \in C^\infty(\mathcal{M}, T\mathcal{M})$ and smooth function $f: \mathcal{M} \rightarrow \mathbb{R}$,

$$\nabla_V(fW) = (\nabla_V f)W + f(\nabla_V W).$$

where $\nabla_V f = \nabla f^\top V$ where ∇ at right denotes the usual gradient operator.

3. **C^1 -linearity:** For any smooth function $f, g: \mathcal{M} \rightarrow \mathbb{R}$ and smooth vector fields $Z, V, W \in C^\infty(\mathcal{M}, T\mathcal{M})$,

$$\nabla_{fZ+gW} V = f\nabla_Z V + g\nabla_W V.$$

General definition of covariate derivative in infinite dimension.. One can obtain a coordinate-free formulation, which is also useful for infinite dimensional manifolds, by defining a bilinear operator through $\Gamma(\partial_i \psi, \partial_j \psi) = \sum_{k=1}^d \Gamma_{ij}^k \partial_k \psi$ and realizing that because of Proposition 9 this is the unique bilinear form fulfilling the property of the following definition.

Definition 19 (Christoffel operator). *For $p \in \mathcal{M}$ the Christoffel operator $\Gamma = \Gamma_p$ is a mapping $\Gamma_p: T_p \mathcal{M} \times T_p \mathcal{M} \rightarrow T_p \mathcal{M}$. For $U, V \in T_p \mathcal{M}$ the evaluation $\Gamma_p(U, V)$ is defined implicitly by*

$$\mathfrak{g}_p(\Gamma_p(U, V), W) = \frac{1}{2} \left((D_p \mathfrak{g})(V)(U, W) + (D_p \mathfrak{g})(U)(V, W) - (D_p \mathfrak{g})(W)(U, V) \right). \quad (\text{A.1})$$

for all $W \in T_p \mathcal{M}$

Next, we will use this operator to define a derivative of vector fields. To this end, let $\gamma: [0, 1] \rightarrow \mathcal{M}$ be a smooth curve, $W: [0, 1] \rightarrow T\mathcal{M}$ a vector field along the curve, i.e. defined through $W(t) = \sum_{l=1}^d w_l(t) \partial_l \psi(\gamma(t)) \in T_{\gamma(t)}\mathcal{M}$ for parameters w_l , and lastly let $\dot{W}(t) := \sum_{l=1}^d \dot{w}_l(t) \partial_l \psi(\gamma(t))$ also define a vector field along γ . Then, by projecting the derivative of $W(t)$ onto the tangent space, i.e. ignoring the normal components in (2.11), we get the following definition (stated again in a coordinate-free fashion).

Definition 20 (Covariant derivative). *Let $\gamma: I \rightarrow \mathcal{M}$ be a curve and $W: I \rightarrow T\mathcal{M}$ a vector field along γ . We define the covariant derivative $\frac{D}{dt}W$ of W along γ at $p = \gamma(t)$ for $t \in I$ by*

$$\mathfrak{g}_{\gamma(t)}\left(\frac{D}{dt}W(t), U\right) = \mathfrak{g}_{\gamma(t)}\left(\dot{W}(t) + \Gamma_p(W(t), \dot{\gamma}(t)), U\right) \quad \forall U \in T_{\gamma(t)}\mathcal{M}. \quad (\text{A.2})$$

This covariant derivative is a generalization of the usual Euclidean directional derivative of a vector-valued function to Riemannian manifolds. Indeed, for Riemannian manifolds that are isometrically embedded into Euclidean space (e.g. the case of embedded surface illustrated before) it is exactly the projection of the Euclidean directional derivative onto the tangent space.