# **Manifold Structured Prediction**

Alessandro Rudi •.1 Carlo Ciliberto •.\*.2 Gian Maria Marconi <sup>3</sup> Lorenzo Rosasco <sup>3,4</sup> alessandro.rudi@inria.fr c.ciliberto@imperial.ac.uk gian.maria.marconi@iit.it lrosasco@mit.edu <sup>1</sup>INRIA - Département d'informatique, École Normale Supérieure - PSL Research University, Paris, France. <sup>2</sup>Department of Electrical and Electronic Engineering, Imperial College, London, UK. <sup>3</sup>Università degli studi di Genova & Istituto Italiano di Tecnologia, Genova, Italy. <sup>4</sup>Massachusetts Institute of Technology, Cambridge, USA. • Equal Contribution

### Abstract

Structured prediction provides a general framework to deal with supervised problems where the outputs have semantically rich structure. While classical approaches consider finite, albeit potentially huge, output spaces, in this paper we discuss how structured prediction can be extended to a continuous scenario. Specifically, we study a structured prediction approach to manifold valued regression. We characterize a class of problems for which the considered approach is statistically consistent and study how geometric optimization can be used to compute the corresponding estimator. Promising experimental results on both simulated and real data complete our study.

## **1** Introduction

Regression and classification are probably the most classical machine learning problems and correspond to estimating a function with scalar and binary values, respectively. In practice, it is often interesting to estimate functions with more structured outputs. When the output space can be assumed to be a vector space, many ideas from regression can be extended, think for example to multivariate [21] or functional regression [33]. However, a lack of a natural vector structure is a feature of many practically interesting problems, such as ranking [18], quantile estimation [27] or graph prediction [41]. In this latter case, the outputs are typically provided only with some distance or similarity function that can be used to design appropriate loss function. Knowledge of the loss is sufficient to analyze an abstract empirical risk minimization approach within the framework of statistical learning theory, but deriving approaches that are at the same time statistically sound and computationally feasible is a key challenge. While ad-hoc solutions are available for many specific problems [15, 40, 25, 7], structured prediction [5] provides a unifying framework where a variety of problems can be tackled as special cases.

Classically, structured prediction considers problems with finite, albeit potentially huge, output spaces. In this paper, we study how these ideas can be applied to non-discrete output spaces. In particular, we consider the case where the output space is a Riemannian manifold, that is the problem of manifold structured prediction (also called manifold valued regression [50]). While also in this case ad-hoc methods are available [51], in this paper we adopt and study a structured prediction approach starting from a framework proposed in [13]. Within this framework, it is possible to derive a statistically sound, and yet computationally feasible, structured prediction approach as long as the loss function satisfies suitable structural assumptions [14, 17, 30, 26, 12, 39]. Moreover we can guarantee that the computed prediction is always an element of the manifold.

Our main technical contribution is a characterization of loss functions for manifold structured prediction satisfying such a structural assumption. In particular, we consider the case where the

<sup>\*</sup>Work performed while C.C. was at the University College London.

<sup>32</sup>nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada.

Riemannian metric is chosen as a loss function. As a byproduct of these results, we derive a manifold structured learning algorithm that is universally consistent and corresponding finite sample bounds. From a computational point of view, the proposed algorithm requires solving a linear system (at training time) and a minimization problem over the output manifold (at test time). To tackle this latter problem, we investigate the application of geometric optimization methods, and in particular Riemannian gradient descent [1]. We consider both numerical simulations and benchmark datasets reporting promising performances. The rest of the paper is organized as follows. In Section 2, we define the problem and explain the proposed algorithm. In Section 3 we state and prove the theoretical results of this work. In Section 4 we explain how to compute the proposed algorithm and we show the performance of our method on synthetic and real data.

### 2 Structured Prediction for Manifold Valued Regression

The goal of supervised learning is to find a functional relation between an input space  $\mathcal{X}$  and an output space  $\mathcal{Y}$  given a finite set of observations. Traditionally, the output space is either a linear space (e.g.  $\mathcal{Y} = \mathbb{R}^M$ ) or a discrete set (e.g.  $\mathcal{Y} = \{0, 1\}$  in binary classification). In this paper, we consider the problem of manifold structured prediction [51], in which output data lies on a manifold  $\mathcal{M} \subset \mathbb{R}^d$ . In this context, statistical learning corresponds to solving

$$\underset{f \in \mathcal{X} \to \mathcal{Y}}{\operatorname{argmin}} \mathcal{E}(f) \quad \text{with} \quad \mathcal{E}(f) = \int_{\mathcal{X} \times \mathcal{Y}} \triangle(f(x), y) \,\rho(x, y) \tag{1}$$

where  $\mathcal{Y}$  is a subset of the manifold  $\mathcal{M}$  and  $\rho$  is an unknown distribution on  $\mathcal{X} \times \mathcal{Y}$ . Here,  $\triangle : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$  is a loss function that measures prediction errors for points estimated on the manifold. The minimization is meant over the set of all measurable functions from  $\mathcal{X}$  to  $\mathcal{Y}$ . The distribution is fixed but unknown and a learning algorithm seeks an estimator  $\hat{f} : \mathcal{X} \rightarrow \mathcal{Y}$  that approximately solves Eq. (1), given a set of training points  $(x_i, y_i)_{i=1}^n$  sampled independently from  $\rho$ .

A concrete example of loss function that we will consider in this paper is  $\triangle = d^2$  the squared geodesic distance  $d: \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$  [28]. The geodesic distance is the natural metric on a Riemannian manifold (it corresponds to the Euclidean distance when  $\mathcal{M} = \mathbb{R}^d$ ) and is a natural loss function in the context of manifold regression [50, 51, 19, 24, 22].

### 2.1 Manifold Valued Regression via Structured Prediction

In this paper we consider a structured prediction approach to manifold valued regression following ideas in [13]. Given a training set  $(x_i, y_i)_{i=1}^n$ , an estimator for problem Eq. (1) is defined by

$$\widehat{f}(x) = \underset{y \in \mathcal{Y}}{\operatorname{argmin}} \sum_{i=1}^{n} \alpha_i(x) \bigtriangleup (y, y_i)$$
(2)

for any  $x \in \mathcal{X}$ . The coefficients  $\alpha(x) = (\alpha_1(x), \dots, \alpha_n(x))^\top \in \mathbb{R}^n$  are obtained solving a linear system for a problem akin to kernel ridge regression (see Section 2.2): given a positive definite kernel  $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  [4] over  $\mathcal{X}$ , we have

$$\alpha(x) = (\alpha_1(x), \dots, \alpha_n(x))^\top = (K + n\lambda I)^{-1} K_x$$
(3)

where  $K \in \mathbb{R}^{n \times n}$  is the empirical kernel matrix with  $K_{i,j} = k(x_i, x_j)$ , and  $K_x \in \mathbb{R}^n$  the vector whose *i*-th entry corresponds to  $(K_x)_i = k(x, x_i)$ . Here,  $\lambda \in \mathbb{R}_+$  is a regularization parameter and  $I \in \mathbb{R}^{n \times n}$  denotes identity matrix.

Computing the estimator in Eq. (2) can be divided into two steps. During a *training step* the score function  $\alpha : \mathcal{X} \to \mathbb{R}^n$  is learned, while during the *prediction step*, the output  $\widehat{f}(x) \in \mathcal{Y}$  is estimated on a new test point  $x \in \mathcal{X}$ . This last step requires minimizing the linear combination of distances  $\Delta(y, y_i)$  between a candidate  $y \in \mathcal{Y}$  and the training outputs  $(y_i)_{i=1}^n$ , weighted by the corresponding scores  $\alpha_i(x)$ . Next, we recall the derivation of the above estimator following [13].

### 2.2 Derivation of the Proposed Estimator

The derivation of  $\hat{f}$  in Eq. (2) is based on the following key structural assumption on the loss.

**Definition 1** (Structure Encoding Loss Function (SELF)). Let  $\mathcal{Y}$  be a compact set. A function  $\triangle : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$  is a Structure Encoding Loss Function if there exist a separable Hilbert space  $\mathcal{H}$ , a continuous feature map  $\psi : \mathcal{Y} \to \mathcal{H}$  and a continuous linear operator  $V : \mathcal{H} \to \mathcal{H}$  such that for all  $y, y' \in \mathcal{Y}$ 

$$\triangle(y, y') = \langle \psi(y), V\psi(y') \rangle_{\mathcal{H}}.$$
(4)

Intuitively, the SELF definition requires a loss function to be "bi-linearizable" over the space  $\mathcal{H}$ . This is similar, but more general, than requiring the loss to be a kernel since it allows also to consider distances (which are not positive definite) or even non-symmetric loss functions. As observed in [13], a wide range of loss functions often used in machine learning are SELF. In Section 3 we study how the above assumption applies to manifold structured loss functions, including the squared geodesic distance.

We first recall how the estimator Eq. (2) can be obtained assuming  $\triangle$  to be SELF. We begin by rewriting the expected risk in Eq. (1) as

$$\mathcal{E}(f) = \int_{\mathcal{X}} \left\langle \psi(f(x)), V \int_{\mathcal{Y}} \psi(y) \, d\rho(y|x) \right\rangle_{\mathcal{H}} \, d\rho(x) \tag{5}$$

where we have conditioned  $\rho(y, x) = \rho(y|x)\rho_{\mathcal{X}}(x)$  and used the linearity of the integral and the inner product. Therefore, any function  $f^* : \mathcal{X} \to \mathcal{Y}$  minimizing the above functional must satisfy the following condition

$$f^*(x) = \operatorname*{argmin}_{y \in \mathcal{Y}} \langle \psi(y), Vg^*(x) \rangle_{\mathcal{H}} \quad \text{where} \quad g^*(x) = \int_{\mathcal{Y}} \psi(y) \, d\rho(y|x) \tag{6}$$

where we have introduced the function  $g^* : \mathcal{X} \to \mathcal{H}$  that maps each point  $x \in \mathcal{X}$  to the conditional expectation of  $\psi(y)$  given x. However we cannot compute explicitly  $g^*$ , but noting that it minimizes the expected least squares error

$$\int \|\psi(y) - g(x)\|_{\mathcal{H}}^2 d\rho(x, y) \tag{7}$$

suggests that a least squares estimator can be considered. We first illustrate this idea for  $\mathcal{X} = \mathbb{R}^d$  and  $\mathcal{H} = \mathbb{R}^k$ . In this case we can consider a ridge regression estimator

$$\widehat{g}(x) = \widehat{W}^{\top} x \quad \text{with} \quad \widehat{W} = \operatorname*{argmin}_{W \in \mathbb{R}^{d \times k}} \frac{1}{n} \|XW - \psi(Y)\|_{F}^{2} + \lambda \|W\|_{F}^{2}$$
(8)

where  $X = (x_1, \ldots, x_n)^\top \in \mathbb{R}^{n \times d}$  and  $\psi(Y) = (\psi(y_1), \ldots, \psi(y_n))^\top \in \mathbb{R}^{n \times k}$  are the matrices whose *i*-th row correspond respectively to the training sample  $x_i \in \mathcal{X}$  and the (mapped) training output  $\psi(y_i) \in \mathcal{H}$ . We have denoted  $\|\cdot\|_F^2$  the squared Frobenius norm of a matrix, namely the sum of all its squared entries. The ridge regression solution can be obtained in closed form as  $\widehat{W} = (X^\top X + n\lambda I)^{-1} X^\top \psi(Y)$ . For any  $x \in \mathcal{X}$  we have

$$\widehat{g}(x) = \psi(Y)^{\top} X (X^{\top} X + n\lambda I)^{-1} x = \psi(Y)^{\top} \alpha(x) = \sum_{i=1}^{n} \alpha_i(x) \psi(y_i)$$
(9)

where we have introduced the coefficients  $\alpha(x) = X(X^{\top}X + n\lambda I)^{-1}x \in \mathbb{R}^n$ . By substituting  $\hat{g}$  to  $g^*$  in Eq. (6) we have

$$\widehat{f}(x) = \operatorname*{argmin}_{y \in \mathcal{M}} \left\langle \psi(y), V\left(\sum_{i=1}^{n} \alpha_i(x)\psi(y_i)\right) \right\rangle = \operatorname*{argmin}_{y \in \mathcal{M}} \sum_{i=1}^{n} \alpha_i(x) \bigtriangleup (y, y_i)$$
(10)

where we have used the linearity of the sum and the inner product to move the coefficients  $\alpha_i$  outside of the inner product. Since the loss is SELF, we then obtain  $\langle \psi(y), V\psi(y_i) \rangle = \triangle(y, y_i)$  for any  $y_i$  in the training set. This recovers the estimator  $\hat{f}$  introduced in Eq. (2), as desired.

We end noting how the above idea can be extended. First, we can consider  $\mathcal{X}$  to be a set and  $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  a positive definite kernel. Then  $\hat{g}$  can be computed by kernel ridge regression (see e.g. [47]) to obtain the scores  $\alpha(x) = (K + n\lambda I)^{-1}K_x$ , see Eq. (3). Second, the above discussion applies if  $\mathcal{H}$  is infinite dimensional. Indeed, thanks to the SELF assumption,  $\hat{f}$  does not depend on explicit knowledge of the space  $\mathcal{H}$  but only on the loss function.

We next discuss the main results of the paper, showing that a large class of loss functions for manifold structured prediction are SELF. This will allow us to prove consistency and learning rates for the manifold structured estimator considered in this work.

### **3** Characterization of SELF Function on Manifolds

In this section we provide sufficient conditions for a wide class of functions on manifolds to satisfy the definition of SELF. A key example will be the case of the squared geodesic distance. To this end we will make the following assumptions on the manifold  $\mathcal{M}$  and the output space  $\mathcal{Y} \subseteq \mathcal{M}$  where the learning problem takes place.

**Assumption 1.**  $\mathcal{M}$  is a complete d-dimensional smooth connected Riemannian manifold, without boundary, with Ricci curvature bounded below and positive injectivity radius.

The assumption above imposes basic regularity conditions on the output manifold. In particular we require the manifold to be locally diffeomorphic to  $\mathbb{R}^d$  and that the tangent space of  $\mathcal{M}$  at any  $p \in \mathcal{M}$  varies smoothly with respect to p. This assumption avoids pathological manifolds and is satisfied for instance by any smooth compact manifold (e.g. the sphere, torus, etc.) [28]. Other notable examples are the statistical manifold (without boundary) [3] any open bounded sub-manifold of the cone of positive definite matrices, which is often studied in geometric optimization settings [1]. This assumption will be instrumental to guarantee the existence of a space of functions  $\mathcal{H}$  on  $\mathcal{M}$  rich enough to contain the squared geodesic distance.

**Assumption 2.**  $\mathcal{Y}$  is a compact geodesically convex subset of the manifold  $\mathcal{M}$ .

A subset  $\mathcal{Y}$  of a manifold is geodesically convex if for any two points in  $\mathcal{Y}$  there exists one and only one minimizing geodesic curve connecting them. The effect of Asm. 2 is twofold. On one hand it guarantees a generalized notion of convexity for the space  $\mathcal{Y}$  on which we will solve the optimization problem in Eq. (2). On the other hand it avoids the geodesic distance to have singularities on  $\mathcal{Y}$ (which is key to our main result below). For a detailed introduction to most definitions and results reviewed in this section we refer the interested reader to standard references for differential and Riemannian geometry (see e.g. [28]). We are ready to prove the main result of this work.

**Theorem 1** (Smooth Functions are SELF). Let  $\mathcal{M}$  satisfy Asm. 1 and  $\mathcal{Y} \subseteq \mathcal{M}$  satisfy Asm. 2. Then, any smooth function  $h : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$  is SELF on  $\mathcal{Y}$ .

*Sketch of the proof (Thm. 1).* The complete proof of Thm. 1 is reported in Appendix A. The proof hinges around the following key steps:

Step 1 If there exists an RKHS  $\mathcal{H}$  on  $\mathcal{M}$ , then any  $h \in \mathcal{H} \otimes \mathcal{H}$  is SELF. Let  $\mathcal{H}$  be a reproducing kernel Hilbert space (RKHS) [4] of functions on  $\mathcal{M}$  with associated bounded kernel  $k : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$ . Let  $\mathcal{H} \otimes \mathcal{H}$  denote the RKHS of functions  $h : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$  with associated kernel  $\bar{k}$  such that  $\bar{k}((y, z), (y', z')) = k(y, y')k(z, z')$  for any  $y, y', z, z' \in \mathcal{M}$ . Let,  $h : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$  be such that  $h \in \mathcal{H} \otimes \mathcal{H}$ . Recall that  $\mathcal{H} \otimes \mathcal{H}$  is isometric to the space of Hilbert-Schmidt operators from  $\mathcal{H}$  to itself. Let  $V_h : \mathcal{H} \to \mathcal{H}$  be the operator corresponding to h via such isometry. We show that the SELF definition is satisfied with  $V = V_h$  and  $\psi(y) = k(y, \cdot) \in \mathcal{H}$  for any  $y \in \mathcal{M}$ . In particular, we have  $\|V\| \leq \|V\|_{\mathrm{HS}} = \|h\|_{\mathcal{H} \otimes \mathcal{H}}$ , with  $\|V\|_{\mathrm{HS}}$  denoting the Hilbert-Schmidt norm of V.

Step 2: Under Asm. 2,  $C_0^{\infty}(\mathcal{M}) \otimes C_0^{\infty}(\mathcal{M})$  "contains"  $C^{\infty}(\mathcal{Y} \times \mathcal{Y})$ . If  $\mathcal{Y}$  is compact and geodesically convex, then it is diffeomorphic to a compact set of  $\mathbb{R}^d$ . By using this fact, we prove that any function in  $C^{\infty}(\mathcal{Y} \times \mathcal{Y})$ , the space of smooth functions on  $\mathcal{Y} \times \mathcal{Y}$ , admits an extension in  $C_0^{\infty}(\mathcal{M} \times \mathcal{M})$  the space of smooth functions on  $\mathcal{M} \times \mathcal{M}$  vanishing at infinity (this is well defined since  $\mathcal{M}$  is diffeomorphic to  $\mathbb{R}^d$ , see Lemma 7 in the supplementary material), and that  $C_0^{\infty}(\mathcal{M} \times \mathcal{M}) = C_0^{\infty}(\mathcal{M}) \otimes C_0^{\infty}(\mathcal{M})$ , with  $\otimes$  the canonical topological tensor product [54].

Step 3: Under Asm. 1, there exists an RKHS on  $\mathcal{M}$  containing  $C_0^{\infty}(\mathcal{M})$ . Under Asm. 1, the Sobolev space  $\mathcal{H} = H_s^2(\mathcal{M})$  of square integrable functions with smoothness *s* is an RKHS for any s > d/2 (see [23] for a definition of Sobolev spaces on Riemannian manifolds).

The proof proceeds as follows: from Step 1, we see that to guarantee h to be SELF it is sufficient to prove the existence of an RKHS  $\mathcal{H}$  such that  $h \in \mathcal{H} \otimes \mathcal{H}$ . The rest of the proof is therefore devoted to showing that for smooth functions this is satisfied for  $\mathcal{H} = H_s^2(\mathcal{M})$ . Since h is smooth, by Step 2 we have that under Asm. 2, there exists a  $\bar{h} \in C_0^{\infty}(\mathcal{M}) \otimes C_0^{\infty}(\mathcal{M})$  whose restriction  $\bar{h}|_{\mathcal{Y} \times \mathcal{Y}}$  to  $\mathcal{Y} \times \mathcal{Y}$  corresponds to h. Now, denote by  $H_s^2(\mathcal{M})$  the Sobolev space of squared integrable functions on  $\mathcal{M}$  with smoothness index s > 0. By construction, (see [23]) for any s > 0, we have  $C_0^{\infty}(\mathcal{M})|_{\mathcal{Y}} \subseteq H_s^2(\mathcal{M})|_{\mathcal{Y}}$ , namely for any function. In particular,  $\bar{h} \in C_0^{\infty}(\mathcal{M}) \otimes C_0^{\infty}(\mathcal{M}) \subseteq H_s^2(\mathcal{M}) \otimes H_s^2(\mathcal{M})$ . Finally, Step 3 guarantees that under Asm. 1,  $\mathcal{H} = H_s^2(\mathcal{M})$  with s > d/2 is an RKHS, showing that  $h \in \mathcal{H} \otimes \mathcal{H}$  as desired.

Interestingly, Thm. 1 shows that the SELF estimator proposed in Eq. (2) can tackle *any* manifold valued learning problem in the form of Eq. (1) with smooth loss function. In the following we study the specific case of the squared geodesic distance.

**Theorem 2** ( $d^2$  is SELF). Let  $\mathcal{M}$  satisfy Asm. 1 and  $\mathcal{Y} \subseteq \mathcal{M}$  satisfy Asm. 2. Then, the squared geodesic distance  $\triangle = d^2 : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$  is smooth on  $\mathcal{Y}$ . Therefore  $\triangle$  is SELF on  $\mathcal{Y}$ .

The proof of the result above is reported in the supplementary material. The main technical aspect is to show that regularity provided by Asm. 2 guarantees the squared geodesic distance to be smooth. The fact that  $\triangle$  is SELF is then an immediate corollary of Thm. 1.

#### 3.1 Statistical Properties of Manifold Structured Prediction

In this section, we characterize the generalization properties of the manifold structured estimator Eq. (2) in light of Thm. 1 and Thm. 2.

**Theorem 3** (Universal Consistency). Let  $\mathcal{M}$  satisfy Asm. 1 and  $\mathcal{Y} \subseteq \mathcal{M}$  satisfy Asm. 2. Let  $\mathcal{X}$  be a compact set and  $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  be a bounded continuous universal kernel<sup>2</sup> For any  $n \in \mathbb{N}$ and any distribution  $\rho$  on  $\mathcal{X} \times \mathcal{Y}$  let  $\widehat{f}_n : \mathcal{X} \to \mathcal{Y}$  be the manifold structured estimator in Eq. (2) for a learning problem with smooth loss function  $\Delta$ , with  $(x_i, y_i)_{i=1}^n$  training points independently sampled from  $\rho$  and  $\lambda_n = n^{-1/4}$ . Then

$$\lim_{n \to \infty} \mathcal{E}(\hat{f}_n) = \mathcal{E}(f^*) \quad \text{with probability 1.}$$
(11)

The result above follows from Thm. 4 in [13] combined with our result in Thm. 1. It guarantees that the algorithm considered in this work finds a consistent estimator for the manifold structured problem, when the loss function is smooth (thus also in the case of the squared geodesic distance). As it is standard in statistical learning theory, we can impose regularity conditions on the learning problem, in order to derive also generalization bounds for  $\hat{f}$ . In particular, if we denote by  $\mathcal{F}$  the RKHS associated to the kernel k, we will require  $g^*$  to belong to the same space  $\mathcal{H} \otimes \mathcal{F}$  where the estimator  $\hat{g}$  introduced in Eq. (9) is learned. In the simplified case discussed in Section 2.2, with linear kernel on  $\mathcal{X} = \mathbb{R}^d$  and  $\mathcal{H} = \mathbb{R}^k$  finite dimensional, we have  $\mathcal{F} = \mathbb{R}^d$  and this assumption corresponds to require the existence of a matrix  $W_*^{\top} \in \mathbb{R}^{k \times d} = \mathcal{H} \otimes \mathcal{F}$ , such that  $g^*(x) = W_*^{\top} x$  for any  $x \in \mathcal{X}$ . In the general case, the space  $\mathcal{H} \otimes \mathcal{F}$  extends to the notion of *reproducing kernel Hilbert space for vector-valued functions* (see e.g. [31, 2]) but the same intuition applies [29, 10, 53]. **Theorem 4** (Generalization Bounds). Let  $\mathcal{M}$  satisfy Asm. 1 and  $\mathcal{Y} \subseteq \mathcal{M}$  satisfy Asm. 2. Let  $\mathcal{H} = H_s^2(\mathcal{M})$  with s > d/2 and  $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$  be a bounded continuous reproducing kernel with associated RKHS  $\mathcal{F}$ . For any  $n \in \mathbb{N}$ , let  $\hat{f}_n$  denote the manifold structured estimator in Eq. (2) for a learning problem with smooth loss  $\Delta : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$  and  $\lambda_n = n^{-1/2}$ . If the conditional mean  $g^*$  belongs to  $\mathcal{H} \otimes \mathcal{F}$ , then

$$\mathcal{E}(\widehat{f}_n) - \mathcal{E}(f^*) \le \mathsf{c}_{\wedge} \mathsf{q} \, \tau^2 \, n^{-\frac{1}{4}} \tag{12}$$

holds with probability not less than  $1 - 8e^{-\tau}$  for any  $\tau > 0$ , with  $c_{\triangle} = \| \triangle \|_{\mathcal{H} \otimes \mathcal{H}}$  and q a constant not depending on  $n, \tau$  or the loss  $\triangle$ .

The generalization bound of Thm. 4 is obtained by adapting Thm. 5 of [13] to our results in Thm. 1 as detailed in the supplementary material. To our knowledge these are the first results characterizing in such generalization properties of an estimator for manifold structured learning with generic smooth loss function. We conclude with a remark on a key quantity in the bound of Thm. 4.

**Remark 1** (The constant  $c_{\Delta}$ ). We comment on the role played in the learning rate by  $c_{\Delta}$ , the norm of the loss function  $\Delta$  seen an element of the Hilbert space  $\mathcal{H} \otimes \mathcal{H}$ . Indeed, from the discussion of Thm. 1 we have seen that any smooth function on  $\mathcal{Y}$  is SELF and belongs to the set  $\mathcal{H} \otimes \mathcal{H}$  with  $\mathcal{H} = H_s^2(\mathcal{M})$ , the Sobolev space of squared integrable functions for s > d/2. Following this interpretation, we see that the bound in Thm. 4 can improve significantly (in terms of the constants) depending on the regularity of the loss function: smoother loss functions will result in "simpler" learning problems and vice-versa. In particular, when  $\Delta$  corresponds to the squared geodesic distance, the more "regular" is the manifold  $\mathcal{M}$ , the learning problem will be. A refined quantitative characterization of  $c_{\Delta}$  in terms of the Ricci curvature and the injective radius of the manifold is left to future work.

<sup>&</sup>lt;sup>2</sup>This is standard assumption for universal consistency (see [52]). An example of continuous universal kernel on  $\mathcal{X} = \mathbb{R}^d$  is the Gaussian  $k(x, x') = \exp(-||x - x'||^2/\sigma)$ , for  $\sigma > 0$ .

Table 1: Structured loss, gradient of the structured loss and retraction for  $P_{++}^m$  and  $S_{d-1}$ .  $Z_i \in P_{++}^m$  and  $z_i \in S_{d-1}$  are the training set points.  $I \in \mathbb{R}^{d \times d}$  is the identity matrix.

	Positive definite matrix manifold $(P_{++}^m)$	Sphere $(\mathcal{S}_{d-1})$
F(y)	$\sum_{i=1}^{n} \alpha_i \ \log(Y^{-\frac{1}{2}} Z_i Y^{-\frac{1}{2}})\ _F^2$	$\sum_{i=1}^{n} \alpha_i \arccos\left(\langle z_i, y \rangle\right)^2$
$\nabla_{\mathcal{M}}F(y)$	$2\sum_{i=1}^{n} \alpha_i Y^{\frac{1}{2}} \log(Y^{\frac{1}{2}} Z_i^{-1} Y^{\frac{1}{2}}) Y^{\frac{1}{2}}$	$4\sum_{i=1}^{n} \alpha_i (yy^T - I) \frac{\arccos(\langle z_i, y \rangle)}{\sqrt{1 - \langle z_i, y \rangle}} z_i$
$R_y(v)$	$Y^{\frac{1}{2}} \exp(Y^{-\frac{1}{2}}vY^{-\frac{1}{2}})Y^{\frac{1}{2}}$	$\frac{v}{\ v\ }$

# 4 Manifold Structured Prediction Algorithm and Experiments

In this section we recall geometric optimization algorithms that can be adopted to perform the estimation of  $\hat{f}$  on a novel test point x. We then evaluate the performance of the proposed method in practice, reporting numerical results on simulated and real data.

The algorithm presented in this paper consists in two steps. In the *training* phase, the matrix  $C = (K + \lambda nI)^{-1}$  is computed, requiring a computational cost of essentially  $O(n^3)$  in time and  $O(n^2)$  in space (see Eq. (3)). In the *inference* phase, given a test point x, the coefficients in Eq. (3) are computed,  $\alpha(x) := Cv(x)$ , requiring essentially O(n) in time, and then the optimization problem in Eq. (2) is solved (see next subsection for more details). Note that it is possible to reduce the computational complexity of the training and evaluation of the coefficients, while retaining the statistical guarantees of the proposed method. Indeed the computation of the coefficients consists essentially in solving a kernel ridge regression problem [10] as analyzed in [13], for which methods based on random projection, as *Nyström* [48] or *random features* [42], reduce the complexity up to  $O(n\sqrt{n})$  in time and O(n) in space, while guaranteeing the same statistical properties [43, 45, 9, 11, 44].

#### 4.1 Optimization on Manifolds

We begin discussing the computational aspects related to evaluating the manifold structured estimator. In particular, we discuss how to address the optimization problem in Eq. (2) in specific settings. Given a test point  $x \in \mathcal{X}$ , this process consists in solving a minimization over  $\mathcal{Y}$ , namely

$$\min_{y \in \mathcal{V}} F(y) \tag{13}$$

where F(y) corresponds to the linear combination of  $\triangle(y, y_i)$  weighted by the scores  $\alpha_i(x)$  computed according to Eq. (3). If  $\mathcal{Y}$  is a linear manifold or a subset of  $\mathcal{M} = \mathbb{R}^d$ , this problem can be solved by means of gradient-based minimization algorithms, such as Gradient Descent (GD):

$$y_{t+1} = y_t - \eta_t \nabla F(y_t) \tag{14}$$

for a step size  $\eta_t \in \mathbb{R}$ . This algorithm can be extended to *Riemannian gradient descent* (RGD) [58] on manifolds, as

$$y_{t+1} = Exp_{y_t}(\eta_t \nabla_{\mathcal{M}} F(y_t)) \tag{15}$$

Where  $\nabla_{\mathcal{M}}F$  is the gradient defined with respect to the Riemannian metric (see [1]) and  $Exp_y$ :  $T_y\mathcal{M} \to \mathcal{M}$  denotes the exponential map on  $y \in \mathcal{Y}$ , mapping a vector from the tangent space  $T_y\mathcal{M}$ to the associated point on the manifold according to the Riemannian metric [28]. For completeness, the algorithm is recalled in Appendix F. For this family of gradient-based algorithms it is possible to substitute the exponential map with a retraction  $R_y: T_y\mathcal{M} \to \mathcal{M}$ , which is a first order approximation to the exponential map. Retractions are often faster to compute and still offer convergence guarantees [1]. In the following experiments we will use both retractions and exponential maps. We mention that the step size  $\eta_t$  can be found with a line search over the validation set, for more see [1]. Note that also stochastic optimization algorithms have been generalized to the Riemannian setting such as R-SGD [8]. Interestingly, methods such as R-SGD can be advantageous in our setting. Indeed, solving the inference in Eq. (2) requires solving the minimization of a sum over n elements, it might be favorable from the computational perspective to adopt a stocastic method that at each iteration minimizes the functional with respect to a single (or a mini-batch) of randomly sampled elements of the entire sum.

Table 2: Simulation experiment: average squared loss (First two columns) and  $\triangle_{PD}$  (Last two columns) error of the proposed structured prediction (SP) approach and the KRLS baseline on learning the inverse of a PD matrix for increasing matrix dimension.

	Squar	ed loss	$ riangle_{ ext{PD}}$ loss		
Dim	KRLS	SP	KRLS	SP	
5	$0.72 {\pm} 0.08$	$0.89{\pm}0.08$	$111 \pm 64$	$0.94{\pm}0.06$	
10	$0.81 {\pm} 0.03$	$0.92 {\pm} 0.05$	$44 \pm 8.3$	$1.24 {\pm} 0.06$	
15	$0.83{\pm}0.03$	$0.91 {\pm} 0.06$	$56 \pm 10$	$1.25 {\pm} 0.05$	
20	$0.85{\pm}0.02$	$0.91 {\pm} 0.03$	$59 \pm 12$	$1.33 {\pm} 0.03$	
25	$0.87{\pm}0.01$	$0.91{\pm}0.02$	$72 \pm 9$	$1.44 {\pm} 0.03$	
30	$0.88{\pm}0.01$	$0.91{\pm}0.02$	$67 \pm 7.2$	$1.55{\pm}0.03$	

Table 1 reports gradients and retraction maps for the geodesic distance of two problems of interest considered in this work: positive definite manifold and the sphere. See Sections 4.2 and 4.3 for more details on the related manifolds.

We point out that using optimization algorithms that comply with the geometry of the manifold, such as RGD, guarantees that the computed value is an element of the manifold. This is in contrast with algorithms that compute a solution in a linear space that contains  $\mathcal{M}$  and then need to project the computed solution onto  $\mathcal{M}$ . We next discuss empirical evaluations of the proposed manifold structured estimator on both synthetic and real datasets.

#### 4.2 Synthetic Experiments: Learning Positive Definite Matrices

We consider the problem of learning a function  $f : \mathbb{R}^d \to \mathcal{Y} = P_{++}^m$ , where  $P_{++}^m$  denotes the *cone of positive definite (PD)*  $m \times m$  matrices. Note that  $P_{++}^m$  is a manifold with squared geodesic distance  $\triangle_{\text{PD}}$  between any two PD matrices  $Z, Y \in P_{++}^m$  defined as

$$\Delta_{\rm PD}(Z,Y) = \|\log(Y^{-\frac{1}{2}}Z\,Y^{-\frac{1}{2}})\|_F^2 \tag{16}$$

where, for any  $M \in P_{++}^m$ , we have that  $M^{\frac{1}{2}}$  and  $\log(M)$  correspond to the matrices with same eigenvectors of M but with respectively the square root and logarithm of the eigenvalues of M. In Table 1 we show the computation of the structured loss, the gradient of the structured loss and the exponential map of the PD matrix manifold. We refer the reader to [32, 6] for a more detailed introduction on the manifold of positive definite matrices.

For the experiments reported in the following we compared the performance of the manifold structured estimator minimizing the loss  $\triangle_{PD}$  and a Kernel Regularized Least Squares classifier (KRLS) baseline (see Appendix G), both trained using the Gaussian kernel  $k(x, x') = \exp(-||x - x'||^2/2\sigma^2)$ . The matrices predicted by the KRLS estimator are projected on the PD manifold by setting to a small positive constant (1e - 12) the negative eigenvalues. For the manifold structured estimator, the optimization problem at Eq. (2) was performed with the Riemannian Gradient Descent (RGD) algorithm [1]. We refer to [58] regarding the implementation of the RGD in the case of the geodesic distance on the PD cone.

Learning the Inverse of a Positive Definite Matrix. We consider the problem of learning the function  $f: P_{++}^m \to P_{++}^m$  such that  $f(X) = X^{-1}$  for any  $X \in P_{++}^m$ . Input matrices are generated as  $X_i = U\Sigma U^\top \in P_{++}^m$  with U a random orthonormal matrix sampled from the Haar distribution [16] and  $S \in P_{++}^m$  a diagonal matrix with entries randomly sampled from the uniform distribution on [0, 10]. We generated datasets of increasing dimension m from 5 to 50, each with 1000 points for training, 100 for validation and 100 for testing. The kernel bandwidth  $\sigma$  was chosen and the regularization parameter  $\lambda$  were selected by cross-validation respectively in the ranges 0.1 to 1000 and  $10^{-6}$  to 1 (logarithmically spaced).

Table 2 reports the performance of the manifold structured estimator (SP) and the KRLS baseline with respect to both the  $\triangle_{PD}$  loss and the least squares loss (normalized with respect to the number of dimensions). Note that the KRLS estimator target is to minimize the least squares (Frobenius) loss and is not designed to capture the geometry of the PD cone. We notice that the proposed approach significantly outperforms the KRLS baseline with respect to the  $\triangle_{PD}$  loss. This is expected:  $\triangle_{PD}$ 



Figure 1: (Left) Fingerprints reconstruction: Average absolute error (in degrees) for the manifold structured estimator (SP), the manifold regression (MR) approach in [51] and the KRLS baseline. (Right) Fingerprint reconstruction of a single image where the structured predictor achieves 15.7 of average error while KRLS 25.3.

penalizes especially matrices with very different eigenvalues and our method cannot predict matrices with non-positive eigenvalues, as opposed to KRLS which computes a linear solution in  $\mathbb{R}^{d^2}$  and then projects it onto the manifold. However the two methods perform comparably with respect to the squared loss. This is consistent with the fact that our estimator is aware of the natural structure of the output space and uses it profitably during learning.

#### 4.3 Fingerprint Reconstruction

We consider the fingerprint reconstruction application in [51] in the context of manifold regression. Given a partial image of a fingerprint, the goal is to reconstruct the contour lines in output. Each fingerprint image is interpreted as a separate structured prediction problem where training input points correspond to the 2D position  $x \in \mathbb{R}^2$  of valid contour lines and the output is the local orientation of the contour line, interpreted as a point on the circumference  $S_1$ . The space  $S_1$  is a manifold with squared geodesic distance  $\Delta_{S_1}$  between two points  $z, y \in S_1$  corresponding to

$$\Delta_{\mathcal{S}_1}(z, y) = \arccos\left(\langle z, y \rangle\right)^2 \tag{17}$$

where  $\arccos$  is the inverse cosine function. In Table 1 we show the computation of the structured loss, the gradient of the structured loss and the chosen retraction for the sphere manifold. We compared the performance of the manifold structured estimator proposed in this paper with the manifold regression approach in [51] on the FVC fingerprint verification challenge dataset<sup>3</sup>. The dataset consists of 48 fingerprint pictures, each with ~ 1400 points for training, ~ 1000 points for validation and the rest (~ 25000) for test.

Fig. 1 reports the average absolute error (in degrees) between the true contour orientation and the one estimated by our structured prediction approach (SP), the manifold regression (MR) in [51] and the KRLS baseline. Our method outperforms the MR competitor by a significant margin. As expected, the KRLS baseline is not able to capture the geometry of the output space and has a significantly larger error of the two other approaches. This is also observed on the qualitative plot in Fig. 1 (Left) where the predictions of our SP approach and the KRLS baseline are compared with the ground truth on a single fingerprint. Output orientations are reported for each pixel with a color depending on their orientation (from 0 to  $\pi$ ). While the KRLS predictions are quite inconsistent, it can be noticed that our estimator is very accurate and even "smoother" than the ground truth.

### 4.4 Multilabel Classification on the Statistical Manifold

We evaluated our algorithm on multilabel prediction problems. In this context the output is an m-dimensional histogram, i.e. a discrete probability distribution over m points. We consider as manifold the space of probability distributions over m points, that is the m-dimensional simplex  $\Delta^m$  endowed with the *Fisher information metric* [3]. We will consider  $\mathcal{Y} = \Delta^m_{\epsilon}$  where we require  $y_1, \ldots, y_m \ge \epsilon$ , for  $\epsilon > 0$ . In the experiment we considered  $\epsilon = 1e - 5$ . The geodesic induced by the

<sup>&</sup>lt;sup>3</sup>http://bias.csr.unibo.it/fvc2004

Table 3: Area under the curve (AUC) on multilabel benchmark datasets [55] for KRLS and SP.

	KRLS	SP (Ours)
Emotions	0.63	0.73
CAL500	<b>0.92</b>	0.92
Scene	0.62	0.73

Fisher metric is,  $d(y, y') = \arccos \left( \sum_{i=1}^{m} \sqrt{y_i y'_i} \right)$  [38]. This geodesic comes from applying the map  $\pi: \Delta^m \to S_{m-1}, \pi(y) = (\sqrt{y_1}, \dots, \sqrt{y_{m+1}})$  to the points  $\{y_i\}_{i=1}^n \in \Delta^m$ . This results in points that belong to the intersection of the positive quadrant  $\mathbb{R}^m_{++}$  and the sphere  $S_{m-1}$ . We can therefore use the geodetic distance on the Sphere and gradient and retraction map described in Table 1. We test our approach on some of the benchmark multilabel datasets described in [55] and we compare the results with the KRLS baseline. We cross-validate  $\lambda$  and  $\sigma$  taking values, respectively, from the intervals  $[10^{-6}, 10^{-1}]$  and [0.1, 10]. We compute the area under curve (AUC) [49] metric to evaluate the quality of the predictions, results are shown in Table 3.

### 4.5 Additional Example

We conclude this section with a further relevant example of our proposed approach to the setting of applications to relational knowledge, which we plan to investigate in future work. In particular, we consider settings where the output space  $\mathcal{M}$  corresponds to the Hyperbolic space (or Poincaré disk), which has recently been adopted by the knowledge representation community to learn embeddings of relational data to encode discrete semantic/hierarchical information [37, 36]. The embedding is such that symbolic objects (e.g. words, entities, concepts) with high semantic or functional similarity are mapped into points with small hyperbolic geodesic distance. Typically, learning the embedding is a time consuming process that requires training from scratch on the whole dataset whenever a new example is provided. To address this issue, with our manifold regression approach we could learn  $f : X \to \mathcal{M}$  with  $x \in X$  the observed entity and f(x) its predicted embedding. This would allow to transfer the embedding learned using techniques such as those in [37] to new points, without retraining the entire system from scratch. Interestingly, our theory is applicable to this setting since  $\mathcal{M}$  satisfies Asm. 1.

# 5 Conclusions

In this paper we studied a structured prediction approach for manifold valued learning problems. In particular we characterized a wide class of loss functions (including the geodesic distance) for which we proved the considered algorithm to be statistically consistent, additionally providing finite sample bounds under standard regularity assumptions. Our experiments show promising results on synthetic and real data using two common manifolds: the positive definite matrices cone and the sphere. With the latter we considered applications on fingerprint reconstruction and multi-labeling. The proposed method leads to some open questions. From a statistical point of view it is of interest how invariants of the manifold explicitly affect the learning rates, see Remark 1. From a more computational perspective, even if experimentally our algorithm achieves good results we did not investigate convergence guarantees in terms of optimization.

#### Acknowledgments.

A. R. acknowledges the support of the European Research Council (grant SEQUOIA 724063). L. R. acknowledges the support of the AFOSR projects FA9550-17-1-0390 and BAA-AFRL-AFOSR-2016-0007 (European Office of Aerospace Research and Development), and the EU H2020-MSCA-RISE project NoMADS - DLV-777826. The work of L. R. and G. M. M. is supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216, and the Italian Institute of Technology. We gratefully acknowledge the support of NVIDIA Corporation for the donation of the Titan Xp GPUs and the Tesla k40 GPU used for this research. This work was supported in part by EPSRC grant EP/P009069/1.

### References

- [1] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [2] Mauricio A Alvarez, Lorenzo Rosasco, Neil D Lawrence, et al. Kernels for vector-valued functions: A review. Foundations and Trends(R) in Machine Learning, 4(3):195–266, 2012.
- [3] Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American Mathematical Soc., 2007.
- [4] Nachman Aronszajn. Theory of reproducing kernels. Transactions of the American mathematical society, 68(3):337–404, 1950.
- [5] GH Bakir, T Hofmann, B Schölkopf, AJ Smola, B Taskar, and SVN Vishwanathan. Predicting structured data. neural information processing, 2007.
- [6] Rajendra Bhatia. Positive definite matrices. Princeton university press, 2009.
- [7] Veli Bicer, Thanh Tran, and Anna Gossen. Relational kernel machines for learning from graph-structured rdf data. In *Extended Semantic Web Conference*, pages 47–62. Springer, 2011.
- [8] Silvere Bonnabel. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013.
- [9] Raffaello Camoriano, Tomás Angles, Alessandro Rudi, and Lorenzo Rosasco. Nytro: When subsampling meets early stopping. In *Artificial Intelligence and Statistics*, pages 1403–1411, 2016.
- [10] Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. Foundations of Computational Mathematics, 7(3):331–368, 2007.
- [11] Luigi Carratino, Alessandro Rudi, and Lorenzo Rosasco. Learning with sgd and random features. Advances in Neural Information Processing Systems, 2018.
- [12] Carlo Ciliberto, Francis Bach, and Alessandro Rudi. Localized structured prediction. *arXiv preprint arXiv:1806.02402*, 2018.
- [13] Carlo Ciliberto, Lorenzo Rosasco, and Alessandro Rudi. A consistent regularization approach for structured prediction. Advances in Neural Information Processing Systems 29 (NIPS), pages 4412–4420, 2016.
- [14] Carlo Ciliberto, Alessandro Rudi, Lorenzo Rosasco, and Massimiliano Pontil. Consistent multitask learning with nonlinear output relations. In Advances in Neural Information Processing Systems, pages 1986–1996, 2017.
- [15] Harold Charles Daume and Daniel Marcu. *Practical structured learning techniques for natural language processing*. Citeseer, 2006.
- [16] Joe Diestel and Angela Spalsbury. The joys of Haar measure. American Mathematical Soc., 2014.
- [17] Moussab Djerrab, Alexandre Garcia, Maxime Sangnier, and Florence d'Alché Buc. Output fisher embedding regression. *Machine Learning*, pages 1–28, 2018.
- [18] John C Duchi, Lester W Mackey, and Michael I Jordan. On the consistency of ranking algorithms. In ICML, pages 327–334, 2010.
- [19] P Thomas Fletcher. Geodesic regression and the theory of least squares on riemannian manifolds. International journal of computer vision, 105(2):171–185, 2013.
- [20] Sylvestre Gallot, Dominique Hulin, and Jacques Lafontaine. *Riemannian geometry*, volume 3. Springer, 1990.
- [21] Wolfgang Härdle and Léopold Simar. *Applied multivariate statistical analysis*, volume 22007. Springer, 2007.
- [22] Søren Hauberg, Oren Freifeld, and Michael J Black. A geometric take on metric learning. In Advances in Neural Information Processing Systems, pages 2024–2032, 2012.
- [23] Emmanuel Hebey. *Nonlinear analysis on manifolds: Sobolev spaces and inequalities*, volume 5. American Mathematical Soc., 2000.

- [24] Jacob Hinkle, Prasanna Muralidharan, P Thomas Fletcher, and Sarang Joshi. Polynomial regression on riemannian manifolds. In *European Conference on Computer Vision*, pages 1–14. Springer, 2012.
- [25] Mohammed Waleed Kadous and Claude Sammut. Classification of multivariate time series and structured data using constructive induction. *Machine learning*, 58(2):179–216, 2005.
- [26] Anna Korba, Alexandre Garcia, and Florence d'Alché Buc. A structured prediction approach for label ranking. Advances in Neural Information Processing Systems, 2018.
- [27] Quoc V Le, Tim Sears, and Alexander J Smola. Nonparametric quantile estimation. 2005.
- [28] John M Lee. Smooth manifolds. In Introduction to Smooth Manifolds, pages 1–29. Springer, 2003.
- [29] Junhong Lin, Alessandro Rudi, Lorenzo Rosasco, and Volkan Cevher. Optimal rates for spectral algorithms with least-squares regression over hilbert spaces. *Applied and Computational Harmonic Analysis*, 2018.
- [30] Giulia Luise, Alessandro Rudi, Massimiliano Pontil, and Carlo Ciliberto. Differential properties of sinkhorn approximation for learning with wasserstein distance. Advances in Neural Information Processing Systems, 2018.
- [31] Charles A Micchelli and Massimiliano Pontil. On learning vector-valued functions. *Neural computation*, 17(1):177–204, 2005.
- [32] Maher Moakher and Philipp G Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. In *Visualization and Processing of Tensor Fields*, pages 285–298. Springer, 2006.
- [33] Jeffrey S Morris. Functional regression. Annual Review of Statistics and Its Application, 2:321–359, 2015.
- [34] Janez Mrčun. On isomorphisms of algebras of smooth functions. Proceedings of the American Mathematical Society, 133(10):3109–3113, 2005.
- [35] Jet Nestruev. Smooth manifolds and observables, volume 220. Springer Science & Business Media, 2006.
- [36] Maximilian Nickel and Douwe Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. *International Conference on Machine Learning*, 2018.
- [37] Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In Advances in neural information processing systems, pages 6338–6347, 2017.
- [38] Frank Nielsen and Ke Sun. Clustering in hilbert simplex geometry. *arXiv preprint arXiv:1704.00454*, 2017.
- [39] Alex Nowak-Vila, Francis Bach, and Alessandro Rudi. Sharp analysis of learning with discrete losses. arXiv preprint arXiv:1810.06839, 2018.
- [40] Sebastian Nowozin, Christoph H Lampert, et al. Structured learning and prediction in computer vision. *Foundations and Trends* (R) *in Computer Graphics and Vision*, 6(3–4):185–365, 2011.
- [41] Benjamin Paaßen, Christina Göpfert, and Barbara Hammer. Time series prediction for graphs in kernel and dissimilarity spaces. *Neural Processing Letters*, pages 1–21, 2017.
- [42] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In Advances in neural information processing systems, pages 1177–1184, 2008.
- [43] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. In *Advances in Neural Information Processing Systems*, pages 1657–1665, 2015.
- [44] Alessandro Rudi, Luigi Carratino, and Lorenzo Rosasco. Falkon: An optimal large scale kernel method. In Advances in Neural Information Processing Systems, pages 3888–3898, 2017.
- [45] Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In Advances in Neural Information Processing Systems, pages 3215–3225, 2017.
- [46] Takashi Sakai. Riemannian geometry, volume 149. American Mathematical Soc., 1996.
- [47] John Shawe-Taylor and Nello Cristianini. *Kernel methods for pattern analysis*. Cambridge university press, 2004.
- [48] Alex J Smola and Bernhard Schölkopf. Sparse greedy matrix approximation for machine learning. 2000.

- [49] Ashwin Srinivasan. Note on the location of optimal classifiers in n-dimensional roc space. 1999.
- [50] Florian Steinke and Matthias Hein. Non-parametric regression between manifolds. In Advances in Neural Information Processing Systems, pages 1561–1568, 2009.
- [51] Florian Steinke, Matthias Hein, and Bernhard Schölkopf. Nonparametric regression between general riemannian manifolds. *SIAM Journal on Imaging Sciences*, 3(3):527–563, 2010.
- [52] Ingo Steinwart and Andreas Christmann. Support vector machines. Springer Science & Business Media, 2008.
- [53] Ingo Steinwart, Don R Hush, Clint Scovel, et al. Optimal rates for regularized least squares regression. In COLT, 2009.
- [54] François Treves. *Topological Vector Spaces, Distributions and Kernels: Pure and Applied Mathematics,* volume 25. Elsevier, 2016.
- [55] Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Mining multi-label data. In Data mining and knowledge discovery handbook, pages 667–685. Springer, 2009.
- [56] Cédric Villani. Optimal transport: old and new, volume 338. Springer Science & Business Media, 2008.
- [57] Franz-Erich Wolter. Distance function and cut loci on a complete riemannian manifold. Archiv der Mathematik, 32(1):92–96, 1979.
- [58] Hongyi Zhang and Suvrit Sra. First-order methods for geodesically convex optimization. In Conference on Learning Theory, pages 1617–1638, 2016.

# Appendix

The appendix of this work is organized in the following sections:

- A SELF property for smooth functions defined on manifolds (Thm. 1).
- B Proof of SELF property for squared geodesic distances (Thm. 2).
- C Generalization bounds for the structured estimator with squared geodesic loss (Thm. 4).
- D Basic definitions and concepts for Riemannian manifolds.
- E Riemannian gradient descent algorithm.
- F A note on KRLS for the experiments in Section 4.2 on PD matrices.

# A Proof of Thm. 1

We prove here intermediate results that will be key to prove Thm. 1. We refer to [28] for basic definitions on manifolds and to [4] for an introduction on reproducing kernel Hilbert spaces (RKHS).

Notation and Definitions. We recall here basic notations and definition that will be used in the following. Given a smooth manifold  $\mathcal{M}$ , for any open subset  $U \subseteq \mathcal{M}$  we denote by  $C^{\infty}(U)$  the set of smooth functions on U and with  $C_0^{\infty}(U)$  the set of *compactly supported* smooth functions on U, namely functions such that the closure of their support is a compact set. For a compact subset  $N \subset \mathcal{M}$  we denote by  $C_0^{\infty}(N)$  the set of all functions  $h: N \to \mathbb{R}$  that admit an extension  $\overline{h} \in C_0^{\infty}(\mathcal{M})$  such that  $\overline{h}|_N = h$  and its support is contained in N, namely it vanishes on the border of N. Finally, for any subset N of  $\mathcal{M}$  we denote  $C^{\infty}(N)$  the set of all functions that admit a smooth extension in  $C^{\infty}(\mathcal{M})$ .

In the following, a central role will be played by tensor product of topological vector spaces [54]. In particular, for a Hilbert space  $\mathcal{H}$ , we will denote  $\mathcal{H} \otimes \mathcal{H}$  the closure of the tensor product between  $\mathcal{H}$  and itself with respect to the canonical norm such that  $||h \otimes h'||_{\mathcal{H} \otimes \mathcal{H}} = ||h||_{\mathcal{H}} ||h'||_{\mathcal{H}}$  for any  $h, h' \in \mathcal{H}$ . Moreover, to given a compact set  $N \subset \mathbb{R}^d$ , we recall that  $C_0^{\infty}(N) \otimes_{\pi} C_0^{\infty}(N)$  denotes the completion of the topological tensor product between  $C_0^{\infty}(N)$  and itself with respect to the projective topology (see [54] Def. 43.2 and 43.5). In the following, for simplicity, we will denote this space with  $C_0^{\infty}(N) \otimes C_0^{\infty}(N)$  with some abuse of notation. Finally, for any subset  $\mathcal{Y} \subseteq \mathcal{M}$  and space  $\mathcal{F}$  of functions from  $\mathcal{M}$  to  $\mathbb{R}$  we denote by  $\mathcal{F}|_{\mathcal{Y}}$  the space of functions from  $\mathcal{Y}$  to  $\mathbb{R}$  that admit an extension in  $\mathcal{F}$ . In particular not that  $C^{\infty}(\mathcal{Y}) = C^{\infty}(\mathcal{M})|_{\mathcal{Y}}$ .

### A.1 Auxiliary Results

We are ready to prove the auxiliary results.

**Lemma 5.** Let  $\mathcal{M}$  be a topological space,  $Y \subseteq \mathcal{M}$  be a compact subset and  $\mathcal{H}$  a reproducing kernel Hilbert space of functions on  $\mathcal{M}$  with kernel  $K : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$  such that there exists  $\kappa > 0$  for which  $k(y, y) \leq \kappa^2$  for any  $y \in \mathcal{Y}$ . Then, for any  $\bar{h} \in \mathcal{H} \otimes \mathcal{H}$ , its restriction to  $\mathcal{Y} \times \mathcal{Y}$ ,  $h = \bar{h}|_{\mathcal{Y} \times \mathcal{Y}}$  is SELF.

*Proof.* Denote  $K_y = k(y, \cdot) \in \mathcal{H}$  for every  $y \in \mathcal{M}$ . Then the space  $\mathcal{H} \otimes \mathcal{H}$  is an RKHS with reproducing kernel  $\bar{K} : (\mathcal{M} \times \mathcal{M}) \times (\mathcal{M} \times \mathcal{M}) \to \mathbb{R}$  such that  $\bar{K}((y, z), (y', z')) = K(y, y')K(z, z')$  for any  $y, y', z, z' \in \mathcal{M}$  (see e.g. [4]). In particular  $\bar{K}_{(y,z)} = K_y \otimes K_z$ . Let now  $\bar{h} : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$  be a function in  $\mathcal{H} \otimes \mathcal{H}$ . In particular, there exist a  $V \in \mathcal{H} \otimes \mathcal{H}$  such that  $\langle V, K_y \otimes K_z \rangle_{\mathcal{H} \otimes \mathcal{H}} = \bar{h}(y, z)$  for any  $y, z \in \mathcal{Y}$  (reproducing property). Note that  $\mathcal{H} \otimes \mathcal{H}$  is isometric to the space of Hilbert-Schmidt operators from  $\mathcal{H}$  to itself, with inner product corresponding to  $\langle A, B \rangle_{\mathcal{H} \otimes \mathcal{H}} = \langle A, B \rangle_{\mathrm{HS}} = \mathrm{Tr}(A^*B)$  for any  $A, B \in \mathcal{H} \otimes \mathcal{H}$ , with  $A^*$  denoting the conjugate of  $A^* \in \mathcal{H} \otimes \mathcal{H}$ . Therefore, for any  $y, z \in \mathcal{Y}$  we have

$$\bar{h}|_{\mathcal{Y}\times\mathcal{Y}}(y,z) = \bar{h}(y,z) = \langle V, K_y \otimes K_z \rangle_{\mathcal{H}\otimes\mathcal{H}} = \operatorname{Tr}(V^*K_y \otimes K_z) = \langle K_z, V^*K_y \rangle_{\mathcal{H}}.$$
 (18)

Since  $K_y$  is bounded in  $\mathcal{H}$ , for  $y \in \mathcal{Y}$  and the operator norm of V is bounded by its Hilbert-Schmidt norm, namely  $||V|| \leq ||V||_{\text{HS}}$ , we can conclude that  $h = \bar{h}|_{\mathcal{Y} \times \mathcal{Y}}$  is indeed SELF.  $\Box$ 

**Lemma 6.** Let  $\mathcal{M}$  satisfy Asm. 1. Then there exists a reproducing kernel Hilbert space of functions  $\mathcal{H}$  on  $\mathcal{M}$ , with bounded kernel, such that  $C_0^{\infty}(\mathcal{M}) \subseteq \mathcal{H}$ .

*Proof.* Let  $H_s^2(\mathcal{M})$  denote the Sobolev space on  $\mathcal{M}$  of squared integrable functions with smoothness s > 0 (see [23] for the definition of Sobolev spaces on Riemannian manifolds). By construction (see page 47 of [23]),  $C_0^{\infty}(\mathcal{M}) \subset H_s^2(\mathcal{M})$  for any s > 0. To prove this Lemma, we will show that  $H_s^2(\mathcal{M})$  is an RKHS for any s > d/2. The proof is organized in two steps.

Step 1:  $H_s^2(\mathcal{M})$  is continuously embedded in  $C(\mathcal{M})$ . By Asm. 1, we can apply Thm. 3.4 in [23] (see also Thm. 2.7 [23] for compact manifolds), which guarantees the existence of a constant C > 0 (see last lines of the proofs for its explicit definition) such that

$$\sup_{y \in \mathcal{M}} |f(y)| \le C ||f||_{\mathcal{H}^2_s(\mathcal{M})},$$

for any  $y \in \mathcal{M}$  and  $f \in \mathcal{H}^2_s(\mathcal{M})$ .

Step 2: Constructing  $\mathcal{H}$  from  $H^2_s(\mathcal{M})$ . Prop. 2.1 of [23] proves that there exists an inner product, that we denote by  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ , whose associated norm is equivalent to  $\|\cdot\|_{H^2_s(\mathcal{M})}$  and such that the space  $\mathcal{H} = (H^2_s(\mathcal{M}), \langle \cdot, \cdot \rangle_{\mathcal{H}})$  is a Hilbert space.

Now, for any  $y \in \mathcal{M}$  denote by  $e_y : \mathcal{H} \to \mathbb{R}$ , the linear functional corresponding to the evaluation, that is  $e_y(f) = f(y)$ . Now by Step 1, we have that the linear functional  $e_y$  is uniformly bounded and so continuous, indeed,

$$|e_y(f)| = |f(y)| \le C ||f||_{\mathcal{H}}, \quad \forall f \in \mathcal{H}.$$

So by the Riesz representation theorem  $e_y \in \mathcal{H}$  and so  $\mathcal{H}$  is a reproducing kernel Hilbert space, with kernel  $k(y, y') = \langle e_y, e_{y'} \rangle_{\mathcal{H}}$ , (see [4], page 343, for more details). Note finally that the kernel is bounded since

$$\begin{aligned} \|e_y\|_{\mathcal{H}} &= \sup_{\|f\|_{\mathcal{H}} \le 1} |\langle e_y, f \rangle_{\mathcal{H}}| = \sup_{\|f\|_{\mathcal{H}} \le 1} |e_y(f)| \le C, \\ &\le \|e_y\|_{\mathcal{H}} \|e_{y'}\|_{\mathcal{H}} \le C^2. \end{aligned}$$

and therefore  $k(y, y') \leq ||e_y||_{\mathcal{H}} ||e_{y'}||_{\mathcal{H}} \leq C^2$ 

In the following, let  $A \subseteq \{f : U \to S\}$  and  $B \subseteq \{g : V \to S\}$ , with U, V, S topological spaces. We denote  $A \cong B$  if there exists an invertible map  $q : U \to V$ , such that  $B = A \circ q^{-1}$  and  $A = B \circ q$ . **Lemma 7** (see also [34, 35]). Let U be a geodesically convex open subset of a d-dimensional complete Riemannian manifold  $\mathcal{M}$  without border, then there exists a smooth map  $q : U \to \mathbb{R}^d$  with smooth inverse, such that

$$C^{\infty}(U) \cong C^{\infty}(\mathbb{R}^d), \quad and \quad C_0^{\infty}(U) \cong C_0^{\infty}(\mathbb{R}^d)$$

moreover for any compact set  $\mathcal{Y} \subset U$  there exists a compact set  $R \subset \mathbb{R}^d$  such that  $R = q(\mathcal{Y})$  and the map s, that is the restriction of q to  $\mathcal{Y} \to R$ , guarantees

$$C^{\infty}(\mathcal{Y}) \cong C^{\infty}(R), \quad and \quad C^{\infty}_{0}(\mathcal{Y}) \cong C^{\infty}_{0}(R)$$

*Proof.* By Lemma 9, there exists a point  $p \in U$  such that  $d(p, \cdot)$  admits all directional derivatives in all points  $q \in U$  (it is, in fact in  $C^{\infty}(U)$ ). We are therefore in the hypotheses of Thm. 2 in [57], from which we conclude that there exists a smooth diffeomorphism between U and  $\mathbb{R}^d$  (with smooth inverse). Denoting by q the diffeomorphism between U and  $\mathbb{R}^d$ , for any function  $f \in C^{\infty}(U)$ , we have  $f \circ q^{-1} \in C^{\infty}(\mathbb{R}^d)$ , so  $C^{\infty}(U) \circ q^{-1} \subseteq C^{\infty}(\mathbb{R}^d)$  and for any function  $g \in C^{\infty}(\mathbb{R}^d)$  we have  $g \circ q \in C^{\infty}(U)$ , so  $C^{\infty}(\mathbb{R}^d) \circ q \subseteq C^{\infty}(U)$ . Finally we recall that if  $A \subseteq B$ , then  $A \circ p \subseteq B \circ p$  for any set A, B and any map p applicable to A, B. Then

$$C^{\infty}(U) = C^{\infty}(U) \circ q^{-1} \circ q \subseteq C^{\infty}(\mathbb{R}^d) \circ q \subseteq C^{\infty}(U)$$

and so  $C^{\infty}(N) \cong C^{\infty}(\mathbb{R}^d)$ . The same reasoning holds  $C_0^{\infty}(U) \cong C_0^{\infty}(\mathbb{R}^d)$ .

Analogously, the smooth diffeomorphism q maps compact subsets of U to compact subsets of  $\mathbb{R}^d$ . Denote by  $R \subset \mathbb{R}^d$  the compact subset that is  $q(\mathcal{Y})$ , the image of  $\mathcal{Y} \subseteq U$  a compact subset of U, then s is the restriction of q to  $\mathcal{Y} \to R$ . By the same reasoning as above, we have that  $C^{\infty}(\mathcal{Y}) \cong C^{\infty}(R)$  via s. **Lemma 8.** Let U be a open geodesically convex subset of a complete Riemannian d-dimensional manifold  $\mathcal{M}$  and  $\mathcal{Y}$  a compact subset of U, then there exists a compact subset  $N \subseteq U$  such that  $\mathcal{Y}$  belongs to the interior of N and

$$C^{\infty}(\mathcal{Y} \times \mathcal{Y}) \subseteq (C_0^{\infty}(N) \otimes C_0^{\infty}(N))|_{Y \times Y}.$$

Moreover,  $C^{\infty}(\mathcal{Y}) \subseteq C_0^{\infty}(N)|_{\mathcal{Y}}$ .

*Proof.* We first consider the real case  $U = \mathcal{M} = \mathbb{R}^d$  with Euclidean metric. By Cor. 2.19 in [28], for any open subset  $V \subset \mathbb{R}^d$  we have that any  $f \in C^{\infty}(\mathcal{Y})$  admits an extension  $\tilde{f} \in C^{\infty}(\mathbb{R}^d)$  such that  $\tilde{f}|_{\mathcal{Y}} = f$  and  $\operatorname{supp} \tilde{f} \subset C_0^{\infty}(V)$ . Then, since  $\mathcal{Y}$  is bounded (compact in a complete space), there exists a bounded open set V containing  $\mathcal{Y}$ . Let  $N = \overline{V}$  the closure of V. N is a compact set as well and contains  $\mathcal{Y}$  in its interior. In particular, since for any  $f \in C^{\infty}(\mathcal{Y})$  the extension  $\tilde{f}$  has support contained in  $V \subset N$ , this shows that  $C^{\infty}(\mathcal{Y}) \subseteq C_0^{\infty}(N)$ . Analogously we have  $C^{\infty}(\mathcal{Y} \times \mathcal{Y}) \subseteq C_0^{\infty}(N \times N)$ .

Now, by Thm. 51.6(a) in [54], we have that

$$C_0^{\infty}(N) \otimes C_0^{\infty}(N) \cong C_0^{\infty}(N \times N).$$
<sup>(19)</sup>

which concludes the proof in the real setting. The proof generalizes trivially to the case where U is an open geodesically convex subset of a complete Riemannian manifold thanks to the isomorphisms between spaces of smooth functions provided by Lemma 7.

### A.2 Proof of Thm. 1

For the following results we need to introduce the concept of *cut locus*. For any  $y \in \mathcal{M}$ , denote by  $\operatorname{Cut}(y) \subseteq \mathcal{M}$  the *cut locus* of y the closure of the set of points  $z \in \mathcal{M}$  that are connected to y by more than one minimal geodesic (see [20, 46]). For any  $y \in \mathcal{Y}$  we have  $y \in \mathcal{M} \setminus \operatorname{Cut}(y)$ , see e.g. Lemma 4.4 in [46].

Finally we refine Asm. 2 to avoid pathological cases. Indeed a geodesically convex set can still have conjugate points on the boundary. To avoid this situation we restate Asm. 2 as follows

**Assumption 2'**  $\widetilde{M}$  is an open geodesically convex subset of the manifold  $\mathcal{M}$  and  $\mathcal{Y}$  is a compact subset of  $\widetilde{M}$ .

*Proof of Thm. 1.* By Asm. 2', let  $\widetilde{\mathcal{M}}$  be an open geodesically convex subset of  $\mathcal{M}$  such that  $\mathcal{Y} \subset \widetilde{\mathcal{M}} \subseteq \mathcal{M}$ . Apply Lemma 8 and let  $N \subseteq \widetilde{\mathcal{M}}$  be a compact set such that  $\mathcal{Y}$  is contained in the interior of N, namely

$$C^{\infty}(\mathcal{Y}) \subseteq C_0^{\infty}(N)|_{\mathcal{Y}} \subseteq C_0^{\infty}(\mathcal{M})|_{\mathcal{Y}} \subseteq \mathcal{H}|_{\mathcal{Y}}.$$
(20)

Then, by applying again Lemma 8 we have

$$C^{\infty}(\mathcal{Y} \times \mathcal{Y}) \subseteq (C_0^{\infty}(N) \otimes C_0^{\infty}(N))|_{\mathcal{Y} \times \mathcal{Y}} \subseteq (\mathcal{H} \otimes \mathcal{H})|_{\mathcal{Y} \times \mathcal{Y}}.$$
(21)

Therefore we conclude that for any  $h \in C^{\infty}(\mathcal{Y})$ , there exists  $\bar{h} : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$  with  $\bar{h} \in \mathcal{H} \otimes \mathcal{H}$ and  $h = \bar{h}|_{\mathcal{Y} \times \mathcal{Y}}$ . Finally we apply Lemma 5 to  $\bar{h}$ , which guarantees h to be SELF.

# B Proof of Thm. 2

We prove a preliminary result.

**Lemma 9.** Let  $\mathcal{M}$  be a Riemannian manifold and N be a geodesically convex subset of  $\mathcal{M}$ . Then,

$$d^2|_{N \times N} \in C^{\infty}(N \times N).$$

*Proof.* For any  $y \in \mathcal{M}$ , denote  $Cut(y) \subseteq \mathcal{M}$  the *cut locus* of y, that is the set of points in  $z \in \mathcal{M}$  that are connected by more than one minimal geodesic curve with y (see [20, 46]). Let  $Cut(\mathcal{M}) =$ 

 $\bigcup_{y \in \mathcal{M}} (\{y\} \times \operatorname{Cut}(y)) \subseteq \mathcal{M} \times \mathcal{M}.$  Then, then the squared geodesic distance is such that (see e.g. [56], page 336)

$$d^2 \in C^{\infty}(\mathcal{M} \times \mathcal{M} \setminus \operatorname{Cut}(\mathcal{M})).$$

Now note that by definition of geodesically convex subset  $N \subseteq \mathcal{M}$ , for any two points in N there exist one and only one minimizing geodesic curve connecting them. Therefore,  $N \times N \cap \text{Cut}(\mathcal{M}) = \emptyset$ and consequently  $N \times N \subseteq \mathcal{M} \times \mathcal{M} \setminus \text{Cut}(\mathcal{M})$ . We conclude that the restriction of  $d^2$  on  $N \times N$ is  $C^{\infty}$  as required.

*Proof of Thm.* 2. By Lemma 9, under Asm. 1 and Asm. 2, the squared geodesic distances is smooth. The desired result is then obtained by applying Thm. 1.  $\Box$ 

# C Proof of Thm. 4

*Proof.* The theorem is proved by combining Thm. 1 with Thm. 5 in [13]. To characterize the constant  $c_{\triangle}$  we need an extra step.

Under Asm. 1 and Asm. 2 and the smoothness of  $\triangle$ , we can apply Thm. 1, which characterizes  $\triangle$  as SELF. According to the proof of Thm. 1 and in particular of Lemma 5, for any  $y, z \in \mathcal{Y}$  we have

$$\Delta(y,z) = \langle \psi(y), V\psi(z) \rangle_{\mathcal{H}}$$
(22)

where  $\mathcal{H} = H_s^2(\mathcal{M})$  with s > d/2,  $\psi(y) = K_y(\cdot)$  where  $K : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$  is the reproducing kernel associated to  $\mathcal{H}$  and  $V : \mathcal{H} \to \mathcal{H}$  is the operator defined in Eq. (18). In particular, by the isometry between the tensor space  $\mathcal{H} \otimes \mathcal{H}$  and the space of Hilbert-Schmidt operators from  $\mathcal{H}$  to  $\mathcal{H}$ , we have

$$\|V\|_{\mathrm{HS}} = \|\bigtriangleup\|_{\mathcal{H}\otimes\mathcal{H}}.\tag{23}$$

To conclude, since  $\triangle$  is SELF, the following generalization bound in Thm. 5 from [13]

$$\mathcal{E}(\widehat{f}) - \mathcal{E}(f^*) \le \|V\| \operatorname{q} \tau^2 n^{-\frac{1}{4}}$$
(24)

holds with probability at least  $1 - 8e^{-\tau}$ . Here, ||V|| denotes the operator norm of ||V|| and q is a constant depending only on  $\mathcal{Y}$  and the distribution  $\rho$  (see end of proof of Lemma 18 for additional details). Finally, we recall, by the relation between the operator and Hilbert-Schmidt norm, that  $||V|| \leq ||V||_{\text{HS}} = ||\Delta||_{\mathcal{H}\otimes\mathcal{H}} = c_{\Delta}$ .

### **D** Additional Comparison

We have implemented the approach in [50, 51] and tested it on the synthetic data in Section 4.2. We note that these works were the only ones we found in the literature that address the problem of manifold-valued regression. We also note in passing that the seminal work of Steinke, Hein and Schölkopf [50, 51] did focus on proposing and solving an ERM strategy for manifold-valued learning, but did not study the theoretical properties of the resulting estimator.

The method essentially consists of a Nadaraya-Watson-like estimator combined with an iterative projection over the output manifold. In our implementation of [51] we used Gaussian kernel and first order polynomials as interpolating functions (see Sec. 4 in [51] for notation and details). Note that in general it is not always possible to uniformly sample the centers for the interpolating functions over the output manifold (as suggested in [51]). Thus, we used the training outputs as centers. To train the model, the parameters of the first order polynomials have been optimized using a quasi-Newton method on the empirical risk with geodesic distance and *Eells energy* regularizer.

$ riangle_{PD}$ err.	Krls	Ours	[29]	time (s)	train	test
n = 10	44	1.24	6.98	KRLS	0.2	0.01
n = 20	59	1.33	11.2	OURS	0.2	9.8
n = 30	67	1.55	18.1	[29]	1740	0.1



Figure 2: Pictorial representations of the exponential map.

The table above (Left) extends Table 2 adding the accuracy of [29] on the synthetic experiment. The table on the right reports the computational times of both our approach and [51] for n = 30 and PD matrices of dimension d = 30 (hence manifold dimension  $O(d^2)$ ). The long training times for [51] are due to the high-dimensional non-linear optimization required to learn the weights parametrizing the interpolating function (which are of order  $O(d^2n)$ , here  $\sim 30000$ ). Note that the complexity of inference scales proportionally (indeed linearly) with the number n of training points and corresponds to finding the barycenter of n points on the output manifold. Stochastic methods such as Riemannian SGD (R-SGD) can be adopted to obtain approximate predictions that are still within the statistical learning bound of Thm. 4, but require less computations. We will: 1) add a remark at the end of Sec 4.1 detailing the impact of adopting a stochastic algorithm in this setting; 2) Add a plot comparing time vs accuracy of Riemannian GD vs R-SGD on the synthetic experiments for different values of n and manifold dimension d. 3) Add the performance of R-SGD for the real experiments in Sec. 4.3 (Table 3 and Fig. 1).

We observe that our approach is significantly more accurate than [51]. Admittedly, it may be possible that better performance could be achieved with [51] by performing a finer cross-validation. However, in practice, the long computational times required to train this method make it prohibitive to perform extensive model-selection. For these reasons we did not run [51] on the multi-labeling tasks.

# E Differential geometry definitions

A Riemannian manifold  $(\mathcal{M}, g)$  of dimension n is a topological space  $\mathcal{M}$  such that every point  $y \in \mathcal{M}$  has a neighbourhood which is homeomorphic to an open set in Euclidean space  $\mathbb{R}^n$  and g is a collection of inner product defined in every tangent space  $T_y\mathcal{M}$  of every point  $y \in \mathcal{M}$ . Intuitively, the tangent space  $T_y\mathcal{M}$  is an approximation of a neighbourhood of  $y \in M$  that has a vector space structure. We will denote the inner product of  $u, v \in T_y\mathcal{M}$  as  $\langle u, v \rangle_y$ . Thanks the inner product structure in every tangent space of the manifold we can compute gradients of functions  $f : \mathcal{M} \to \mathbb{R}$  that we will denote with  $\nabla_{\mathcal{M}} f : \mathcal{F}_y(\mathcal{M}) \to T_y\mathcal{M}$ . Where  $F_y(\mathcal{M})$  is the set of smooth real-valued functions defined on a neighbourhood of y.

For any  $y_0, y_1 \in \mathcal{M}$  and  $v \in T_y \mathcal{M}$  there is a unique smooth *geodesic* curve  $\gamma: [0, 1] \to \mathcal{M}$  such that  $\gamma(0) = y_0, \gamma(1) = y_1$  and  $\frac{d}{dt}\gamma(0) = v$ , this curve locally minimizes the path between  $y_0$  and  $y_1$ . Given the geodesic between  $y_0$  and  $y_1$  with derivative  $\frac{d}{dt}\gamma(0) = v$ , the exponential map  $Exp_{y_0}: T_{y_0}\mathcal{M} \to \mathcal{M}$  maps vector  $v \in T_{y_0}\mathcal{M}$  to  $y_1$ . A retraction  $R_y: T_y\mathcal{M} \to \mathcal{M}$ , is a first order approximation of the exponential map. Exponential maps are retractions.

# F Riemannian Gradient Descent

In this section we report fully the algorithm Riemannian Gradient Descent.

#### Algorithm 1 Riemannian gradient descent

**Require:** number of iterations T, step size  $\eta$ , initial point  $y_0$ 1: for t = 0, ..., T - 1 do 2:  $v_t = \nabla_{\mathcal{M}} \sum_{i=1}^n \alpha_i(x) \bigtriangleup (y_t, y_i)$ 3:  $y_{t+1} \leftarrow R_{y_t}(\eta_t v_t)$ 4: end for 5: return  $y_T$ 

# G Kernel Regularized Least Squares estimator for Positive definite matrices

We consider the case where we want to use KRLS estimators to predict a positive definite matrix given a data set  $\{x_i, y_i\}_{i=1}^n$ . The KRLS estimator  $f \colon \mathbb{R}^d \to \mathbb{R}$  is a function defined as  $f(x) = \sum_{i=1}^n k(x, x_i)w_i$ , where  $k \colon \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$  is a reproducing kernel and  $w = [w_1, \ldots, w_n] \in \mathbb{R}^n$  are constant weights computed by solving the problem:

$$\min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{N} \|\hat{y}_i - \hat{K}w\|^2 + \lambda \|w\|^2$$

 $\hat{K} \in \mathbb{R}^{n \times n}$  is the kernel matrix whose elements are defined as  $(K)_{ij} = k(x_i, x_j)$ .

To predict a positive definite matrix  $y \in \mathbb{P}_d^{++}$ , a KRLS estimator is learned for every element of the flattened matrix  $vec(y) \in \mathbb{R}^{d^2}$ . Suppose  $j \in \{1, \ldots, d^2\}$  is the index of the *j*-th component of vec(y) that we want to predict, then we want to learn the estimator  $f^{(j)}(x) = \sum_{i=1}^n k(x, x_i)w_i^{(j)}$ . The corresponding problem has labels  $\hat{y}^{(j)} = [vec(y_1)_j, \ldots, vec(y_{d^2})_j]$  and we solve for  $w^{(j)} = [w_1^{(j)}, \ldots, w_n^{(j)}]$ . Indeed we compute  $d^2$  estimator to predict  $vec(f) = [f^{(1)}(x), \ldots, f^{(d^2)}]$  and then recover *y* from its vectorized form. Once the matrix is predicted we enforce it to be positive definite by performing a spectral decomposition and setting the negative eigenvalues to a small positive constant.

In general, when doing structured predictions with KRLS approach, it is necessary to project the outcome of the prediction on the desired manifold.