

TPOPT: EFFICIENT TRAINABLE TEMPLATE OPTIMIZATION ON LOW-DIMENSIONAL MANIFOLDS

Anonymous authors

Paper under double-blind review

ABSTRACT

In scientific and engineering scenarios, a recurring task is the detection of low-dimensional families of signals or patterns. A classic family of approaches, exemplified by template matching, aims to cover the search space with a dense template bank. While simple and highly interpretable, it suffers from poor computational efficiency due to unfavorable scaling in the signal space dimensionality. In this work, we study TpopT (TemPlate OPTimization) as an alternative scalable framework for detecting low-dimensional families of signals which maintains high interpretability. We provide a theoretical analysis of the convergence of Riemannian gradient descent for TpopT, and prove that it has a superior dimension scaling to covering. We also propose a practical TpopT framework for nonparametric signal sets, which incorporates techniques of embedding and kernel interpolation, and is further configurable into a trainable network architecture by unrolled optimization. The proposed trainable TpopT exhibits significantly improved efficiency-accuracy tradeoffs for gravitational wave detection, where matched filtering is currently a method of choice. We further illustrate the general applicability of this approach with experiments on handwritten digit data.

1 INTRODUCTION

Low-dimensional structure is ubiquitous in data arising from physical systems: these systems often involve relatively few intrinsic degrees of freedom, leading to low-rank Ji et al. (2010); Gibson et al. (2022), sparse Quan et al. (2015), or manifold structure Mokhtarian & Abbasi (2002); Brown et al. (2004); Lunga et al. (2013). In this paper, we study the fundamental problem of detecting and estimating signals which belong to a low-dimensional manifold, from noisy observations Wakin et al. (2005); Wakin (2007); Baraniuk & Wakin (2009).

Perhaps the most classical and intuitive approach to detecting families of signals is *matched filtering* (MF), which constructs a bank of templates, and compares them individually with the observation. Due to its simplicity and interpretability, MF remains the core method of choice in the gravitational wave detection of the scientific collaborations LIGO Abramovici et al. (1992); Abbott et al. (2015), Virgo Acernese et al. (2015) and KARGA Akutsu et al. (2021), where massive template banks are constructed to search for traces of gravitational waves produced by pairs of merging black holes in space Owen & Sathyaprakash (1999); Abbott et al. (2016; 2017); Yan et al. (2022b). Emerging advances on template placement (Roy et al., 2017; 2019) and optimization (Weerathunga & Mohanty, 2017; Pal & Nayak, 2023; Dal Canton et al., 2021) provide promising ideas of growth. The conceptual idea of large template banks for detection is also widely present in other scenarios such as neuroscience Shi et al. (2010), geophysics Caffagni et al. (2016); Rousset et al. (2017), image pose recognition Picos et al. (2016), radar signal processing Pardhu et al. (2014); Johnson (2009), and aerospace engineering Murphy et al. (2017). In the meantime, many modern learning architectures employ similar ideas of matching inputs with template banks, such as transformation-invariant neural networks which create a large number of templates by applying transformations to a smaller family of filters Sohn & Lee (2012); Kanazawa et al. (2014); Zhou et al. (2017).

One major limitation of this approach is its unfavorable scaling with respect to the signal manifold dimension. For gravitational wave detection, this leads to massive template banks in deployment, and presents a fundamental barrier to searching broader and higher dimensional signal manifolds. For

transformation-invariant neural networks, the dimension scaling limits their applications to relatively low-dimensional transformation groups such as rotations.

This paper is motivated by a simple observation: instead of using sample templates to cover the search space, we can search for a best-matching template via optimization over the search space with higher efficiency. In other words, while MF searches for the best-matching template by enumeration, a first-order optimization method can leverage the geometric properties of the signal set, and avoid the majority of unnecessary templates. We refer to this approach as template optimization (TpopT).

In many practical scenarios, we lack an analytical characterization of the signal manifold. We propose a nonparametric extension of TpopT, based on signal embedding and kernel interpolation, which retains the test-time efficiency of TpopT.¹ The components of this method can be trained on sample data, reducing the need for parameter tuning and improving the performance in Gaussian noise. Our training approach draws inspiration from unrolled optimization Chen et al. (2022), which treats the iterations of an optimization method as layers of a neural network. This approach has been widely used for estimating low-dimensional (sparse) signals Liu & Chen (2019); Xin et al. (2016) with promising results on a range of applications Monga et al. (2021); Diamond et al. (2017); Liang et al. (2019); Buchanan et al. (2022). The main contributions of this paper are as follows:

- Propose trainable TpopT as an efficient approach to detecting and estimating signals from low-dimensional families, with nonparametric extensions when an analytical data model is unavailable.
- Prove that Riemannian gradient descent for TpopT is exponentially more efficient than MF.
- Demonstrate significantly improved complexity-accuracy tradeoffs for gravitational wave detection, where MF is currently a method of choice.

2 PROBLEM FORMULATION AND METHODS

In this section, we describe the problem of detecting and recovering signals from a low-dimensional family, and provide a high-level overview of two approaches — matched filtering and template optimization (TpopT). The problem setup is simple: assume the signals of interest form a d -dimensional manifold $S \subset \mathbb{R}^D$, where $d \ll D$, and that they are normalized such that $S \subset \mathbb{S}^{D-1}$. For a given observation $\mathbf{x} \in \mathbb{R}^D$, we want to determine whether \mathbf{x} consists of a noisy copy of some signal of interest, and recover the signal if it exists. More formally, we model the observation and label as:

$$\mathbf{x} = \begin{cases} a \mathbf{s}_y + \mathbf{z} & \text{if } y = 1 \\ \mathbf{z}, & \text{if } y = 0 \end{cases} \quad (1)$$

where $a \in \mathbb{R}_+$ is the signal amplitude, $\mathbf{s}_y \in S$ is the ground truth signal, and $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Our goal is to solve this detection and estimation problem with simultaneously high statistical accuracy and computational efficiency.

Matched Filtering. A natural decision statistic for this detection problem is $\max_{\mathbf{s} \in S} \langle \mathbf{s}, \mathbf{x} \rangle$, i.e.

$$\hat{y}(\mathbf{x}) = 1 \iff \max_{\mathbf{s} \in S} \langle \mathbf{s}, \mathbf{x} \rangle \geq \tau \quad (2)$$

where τ is some threshold, and the recovered signal can be obtained as $\arg \max_{\mathbf{s} \in S} \langle \mathbf{s}, \mathbf{x} \rangle$.² *Matched filtering*, or template matching, approximates the above decision statistic with the maximum over a finite bank of templates $\mathbf{s}_1, \dots, \mathbf{s}_{n_{\text{templates}}}$:

$$\hat{y}_{\text{MF}}(\mathbf{x}) = 1 \iff \max_{i=1, \dots, n_{\text{templates}}} \langle \mathbf{s}_i, \mathbf{x} \rangle \geq \tau. \quad (3)$$

The template \mathbf{s}_i contributing to the highest correlation is thus the recovered signal. This matched filtering method is a fundamental technique in signal detection (simultaneously obtaining the estimated signals), playing an especially significant role in scientific applications Owen & Sathyaprakash (1999); Rousset et al. (2017); Shi et al. (2010).

¹In contrast to conventional manifold learning, where the goal is to learn a representation of the data manifold Bengio & Monperrus (2004); Culpepper & Olshausen (2009); Rifai et al. (2011); Park et al. (2015); Kumar et al. (2017), our goal is to learn an *optimization algorithm* on the signal manifold.

²This statistic is optimal for detecting a single signal \mathbf{s} in iid Gaussian noise; this is the classical motivation for matched filtering Helstrom (2013). For detecting a family of signals $\mathbf{s} \in S$, it is no longer statistically optimal Yan et al. (2022a). However, it remains appealing due to its simplicity.

If the template bank densely covers S , (3) will accurately approximate (2). However, dense covering is inefficient — the number n of templates required to cover S up to some target radius r grows as $n \propto 1/r^d$, making this approach impractical for all but the smallest d .³

Template Optimization. Rather than densely covering the signal space, *template optimization* (TpopT) searches for a best matching template \hat{s} , by numerically solving

$$\hat{s}(\mathbf{x}) = \arg \min_{\mathbf{s} \in S} f(\mathbf{s}) \equiv -\langle \mathbf{s}, \mathbf{x} \rangle. \quad (4)$$

The decision statistic is then $\hat{y}_{\text{TpopT}}(\mathbf{x}) = 1 \iff \langle \hat{s}(\mathbf{x}), \mathbf{x} \rangle \geq \tau$. Since the domain of optimization S is a Riemannian manifold, in principle, the optimization problem (4) can be solved by the Riemannian gradient iteration Boumal (2023)

$$\mathbf{s}^{k+1} = \exp_{\mathbf{s}^k} \left(-\alpha_k \text{grad}[f](\mathbf{s}^k) \right). \quad (5)$$

Here, k is the iteration index, $\exp_{\mathbf{s}}(\mathbf{v})$ is the exponential map at point \mathbf{s} , $\text{grad}[f](\mathbf{s})$ is the Riemannian gradient⁴ of the objective f at point \mathbf{s} , and α_k is the step size.

Alternatively, if the signal manifold S admits a global parameterization $\mathbf{s} = \mathbf{s}(\boldsymbol{\xi})$, we can optimize over the parameters $\boldsymbol{\xi}$, solving $\hat{\boldsymbol{\xi}}(\mathbf{x}) = \arg \min_{\boldsymbol{\xi}} -\langle \mathbf{s}(\boldsymbol{\xi}), \mathbf{x} \rangle$ using the (Euclidean) gradient method:

$$\boldsymbol{\xi}^{k+1} = \boldsymbol{\xi}^k + \alpha_k \cdot \left(\nabla_{\boldsymbol{\xi}} \mathbf{s}(\boldsymbol{\xi}^k) \right)^{\text{T}} \mathbf{x}, \quad (6)$$

where $\nabla_{\boldsymbol{\xi}} \mathbf{s}(\boldsymbol{\xi}^k) \in \mathbb{R}^{D \times d}$ is the Jacobian matrix of $\mathbf{s}(\boldsymbol{\xi})$ at point $\boldsymbol{\xi}^k$. Finally, the estimated signal $\hat{s}(\mathbf{x}) = \mathbf{s}(\hat{\boldsymbol{\xi}}(\mathbf{x}))$ and decision statistic \hat{y}_{TpopT} can be obtained from the estimated parameters $\hat{\boldsymbol{\xi}}$.

Of course, the optimization problem (4) is in general nonconvex, and methods (5)-(6) only converge to global optima when they are initialized sufficiently close to the solution of (4). We can guarantee global optimality by employing multiple initializations $\mathbf{s}_1^0, \dots, \mathbf{s}_{n_{\text{init}}}^0$, which cover the manifold S at some radius Δ where at least one initialization is guaranteed to produce a global optimizer.

In the next section, we will corroborate these intuitions with rigorous analysis. In subsequent sections, we will further develop more practical counterparts to (5)-(6) which (i) do not require an analytical representation of the signal manifold S [Section 4], and (ii) can be trained on sample data to improve statistical performance [Section 5].

3 THEORY: EFFICIENCY GAINS OVER MATCHED FILTERING

The efficiency advantage of optimization comes from its ability to use gradient information to rapidly converge to $\hat{s} \approx \mathbf{s}_{\hat{\eta}}$, within a basin of initializations \mathbf{s}^0 satisfying $d(\mathbf{s}^0, \mathbf{s}_{\hat{\eta}}) \leq \Delta$: the larger the basin, the fewer initializations are needed to guarantee global optimality. The basin size Δ in turn depends on the geometry of the signal set S , through its *curvature*. Figure 1 illustrates the key intuition: if the curvature is small, there exists a relatively large region in which the gradient of the objective function points towards the global optimizer \mathbf{s}^* . On the other hand, if the signal manifold is very curvy, there may only exist a relatively small region in which the gradient points in the correct direction.

We can formalize this intuition through the curvature of geodesics on the manifold S . For a smooth curve $\gamma: [0, T] \rightarrow S \subset \mathbb{R}^n$, with unit speed parameterization $\gamma(t)$, $t \in [0, T]$, the maximum curvature is

$$\kappa(\gamma) = \sup_{t \in T} \|\ddot{\gamma}(t)\|_2. \quad (7)$$

Geometrically, κ^{-1} is the minimum, over all points $\gamma(t)$, of the radius of the osculating circle whose velocity and acceleration match those of γ at t . We extend this definition to S , a Riemannian submanifold of \mathbb{R}^n , by taking κ to be the maximum curvature of any geodesic on S :

$$\kappa(S) = \sup_{\gamma \subset S : \text{unit-speed geodesic}} \kappa(\gamma). \quad (8)$$

³This inefficiency has motivated significant efforts in applied communities to optimize the placement of the templates \mathbf{s}_i , maximizing the statistical performance for a given fixed $n_{\text{templates}}$ Owen & Sathyaprakash (1999). It is also possible to learn these templates from data, leveraging connections to neural networks Yan et al. (2022a). Nevertheless, the curse of dimensionality remains in force.

⁴The Riemannian gradient is the projection of the Euclidean gradient $\nabla_{\mathbf{s}} f$ onto the tangent space $T_{\mathbf{s}} S$.




Figure 1: Relationship between curvature and convergence basins of gradient descent. Gradient descent has larger convergence basins under lower curvature (larger radius of osculating circle). Points within the convergence basin have gradient descent direction pointing “toward” s^* , while points outside the basin may have gradient descent pointing “away from” s^* .

We call this quantity the *extrinsic geodesic curvature* of S .⁵ Our main theoretical result shows that, as suggested by Figure 1 there is a $\Delta = 1/\kappa$ neighborhood within which gradient descent rapidly converges to a close approximation of s_{\natural} :

Theorem 1. *Suppose the extrinsic geodesic curvature of S is bounded by κ . Consider the Riemannian gradient method (5), with initialization satisfying $d(s^0, s_{\natural}) < 1/\kappa$, and step size $\tau = \frac{1}{64}$. Then when $\sigma \leq c/(\kappa\sqrt{d})$, with high probability, we have for all k*

$$d(s^{k+1}, s_{\natural}) \leq (1 - \epsilon) d(s^k, s_{\natural}) + C\sigma\sqrt{d}. \quad (9)$$

Moreover, when $\sigma \leq c/(\kappa\sqrt{D})$, with high probability, we have for all k

$$d(s^k, s^*) \leq C(1 - \epsilon)^k \sqrt{f(s^0) - f(s^*)}, \quad (10)$$

where s^* is the unique minimizer of f over $B(s_{\natural}, 1/\kappa)$. Here, C, c, ϵ are positive numerical constants.

Interpretation: Convergence to Optimal Statistical Precision In (9), we show that under a relatively mild condition on the noise, gradient descent exhibits linear convergence to a $\sigma\sqrt{d}$ -neighborhood of s_{\natural} . This accuracy is the best achievable up to constants: for small σ , with high probability any minimizer s^* satisfies $d(s^*, s_{\natural}) > c\sigma\sqrt{d}$, and so the accuracy guaranteed by (9) is optimal up to constants. Also noteworthy is that both the accuracy and the required bound on the noise level σ are *dictated solely by the intrinsic dimension d* . The restriction $\sigma \leq c/(\kappa\sqrt{D})$ has a natural interpretation in terms of Figure 1 — at this scale, the noise “acts locally”, ensuring that s^* is close enough to s_{\natural} so that for any initialization in $B(s_{\natural}, \Delta)$, the gradient points toward s^* . In (10) we also show that under a stronger condition on σ , gradient descent enjoys linear convergence for *all* iterations k .

Implications on Complexity. Here we compare the complexity required for MF and TpopT to achieve a target estimation accuracy $d(\hat{s}, s_{\natural}) \leq r$. The complexity of MF is simply N_r , the covering number of S with radius r . On the other hand, the complexity of TpopT is dictated by $n_{\text{init}} \times n_{\text{gradient-step}}$. We have $n_{\text{init}} = N_{1/\kappa}$ since TpopT requires initialization within radius $1/\kappa$ of s_{\natural} , and $n_{\text{gradient-step}} \propto \log 1/\kappa r$ because gradient descent enjoys a linear convergence rate. Note that the above argument applies when $C\sigma d^{1/2}/\epsilon \leq r \leq 1/\kappa$, where the upper bound on r prescribes the regime where gradient descent is in action (otherwise TpopT and MF are identical), and the lower bound on r reflects the statistical limitation due to noise. Since the covering number $N_{\text{radius}} \propto (1/\text{radius})^d$, the complexities of the two methods T_{MF} and T_{TPopT} are given by

$$T_{\text{MF}} \propto 1/r^d, \quad T_{\text{TPopT}} \propto \kappa^d \log(1/\kappa r). \quad (11)$$

Combining this with the range of r , it follows that TpopT always has superior dimensional scaling than MF whenever the allowable estimation error r is below $1/\kappa$ (and identical to MF above that). The advantage is more significant at lower noise and higher estimation accuracy.

⁵Notice that $\kappa(S)$ measures how S curves in the ambient space \mathbb{R}^n ; this is in contrast to traditional *intrinsic* curvature notions in Riemannian geometry, such as the sectional and Ricci curvatures. An extrinsic notion of curvature is relevant here because our objective function $f(s) = -\langle s, x \rangle$ is defined extrinsically. Intrinsic curvature also plays an important role in our arguments — in particular, in controlling the effect of noise.

Proof Ideas. The proof of Theorem 1 follows the intuition in Figure 1, by (i) considering a noiseless version of the problem and showing that in a $1/\kappa$ ball, the gradient points towards \mathbf{s}_b , and (ii) controlling the effect of noise, by bounding the maximum component T^{\max} of the noise \mathbf{z} along any tangent vector $\mathbf{v} \in T_s S$ at any point $\mathbf{s} \in B(\mathbf{s}_b, \Delta)$. By carefully controlling T^{\max} , we are able to achieve rates driven by intrinsic dimension, not ambient dimension. Intuitively, this is because the collection of tangent vectors, i.e., the tangent bundle, has dimension $2d$. Our proof involves a discretization argument, which uses elements of Riemannian geometry (Toponogov’s theorem on geodesic triangles, control of parallel transport via the second fundamental form Toponogov (2006); Lee (2018)). To show convergence of iterates (10), we show that in a $1/\kappa$ region, the objective f enjoys Riemannian strong convexity and Lipschitz gradients Boumal (2023). Please see the supplementary material for complete proofs.

4 NONPARAMETRIC TPOPT VIA EMBEDDING AND KERNEL INTERPOLATION

The theoretical results in Section 3 rigorously quantify the advantages of TpopT in detecting and estimating signals from low-dimensional families. A straightforward application of TpopT requires a precise analytical characterization of the signal manifold. In this section, we develop more practical, nonparametric extension of TpopT, which is applicable in scenarios in which we *only* have examples $\mathbf{s}_1, \dots, \mathbf{s}_N$ from S . This extension will maintain the test-time efficiency advantages of TpopT.

Embedding. We begin by embedding the example points $\mathbf{s}_1, \dots, \mathbf{s}_N \in \mathbb{R}^n$ into a lower-dimensional space \mathbb{R}^d , producing data points $\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_N \in \mathbb{R}^d$. The mapping φ should preserve pairwise distances and can be chosen in a variety of ways; Because the classical Multidimensional Scaling (MDS) setup on Euclidean distances is equivalent to Principal Component Analysis (PCA), we simply use PCA in our experiments. Assuming that φ is bijective over S , we can take $\mathbf{s} = \mathbf{s}(\boldsymbol{\xi})$ as an approximate parameterization of S , and develop an optimization method which, given an input \mathbf{x} , searches for a parameter $\boldsymbol{\xi} \in \mathbb{R}^d$ that minimizes $f(\mathbf{s}(\boldsymbol{\xi})) = -\langle \mathbf{s}(\boldsymbol{\xi}), \mathbf{x} \rangle$.

Kernel Interpolated Jacobian Estimates. In the nonparametric setting, we only know the values of $\mathbf{s}(\boldsymbol{\xi})$ at the finite point set $\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_N$, and we do not have any direct knowledge of the functional form of the mapping $\mathbf{s}(\cdot)$ or its derivatives. To extend TpopT to this setting, we can estimate the Jacobian $\nabla \mathbf{s}(\boldsymbol{\xi})$ at point $\boldsymbol{\xi}_i$ by solving a weighted least squares problem

$$\widehat{\nabla} \mathbf{s}(\boldsymbol{\xi}_i) = \arg \min_{\mathbf{J} \in \mathbb{R}^{d \times d}} \sum_{j=1}^N w_{j,i} \|\mathbf{s}_j - \mathbf{s}_i - \mathbf{J}(\boldsymbol{\xi}_j - \boldsymbol{\xi}_i)\|_2^2, \quad (12)$$

where the weights $w_{j,i} = \Theta(\boldsymbol{\xi}_i, \boldsymbol{\xi}_j)$ are generated by an appropriately chosen kernel Θ . The least squares problem (12) is solvable in closed form. In practice, we prefer compactly supported kernels, so that the sum in (12) involves only a small subset of the points $\boldsymbol{\xi}_j$;⁶ in experiment, we choose Θ to be a truncated radial basis function kernel $\Theta_{\lambda, \delta}(\mathbf{x}_1, \mathbf{x}_2) = \exp(-\lambda \|\mathbf{x}_1 - \mathbf{x}_2\|_2^2) \cdot \mathbb{1}_{\|\mathbf{x}_1 - \mathbf{x}_2\|_2 < \delta}$. When example points \mathbf{s}_i are sufficiently dense and the kernel Θ is sufficiently localized, $\widehat{\nabla} \mathbf{s}(\boldsymbol{\xi})$ will accurately approximate the true Jacobian $\nabla \mathbf{s}(\boldsymbol{\xi})$.

Expanding the Basin of Attraction using Smoothing. In actual applications such as computer vision and astronomy, the signal manifold S often exhibits large curvature κ , leading to a small basin of attraction. One classical heuristic for increasing the basin size is to *smooth* the objective function f . We can incorporate smoothing by taking gradient steps with a kernel smoothed Jacobian,

$$\widetilde{\nabla} \mathbf{s}(\boldsymbol{\xi}_i) = Z^{-1} \sum_j w_{j,i} \widehat{\nabla} \mathbf{s}(\boldsymbol{\xi}_j), \quad (13)$$

where $w_{j,i} = \Theta_{\lambda_s, \delta_s}(\boldsymbol{\xi}_i, \boldsymbol{\xi}_j)$ and $Z = \sum_j w_{j,i}$. The gradient iteration becomes

$$\boldsymbol{\xi}^{k+1} = \boldsymbol{\xi}^k + \alpha_k \widetilde{\nabla} \mathbf{s}(\boldsymbol{\xi}^k)^T \mathbf{x}. \quad (14)$$

When the Jacobian estimate $\widetilde{\nabla} \mathbf{s}(\boldsymbol{\xi})$ accurately approximates $\nabla \mathbf{s}(\boldsymbol{\xi})$, we have

$$\widetilde{\nabla} \mathbf{s}(\boldsymbol{\xi}_i)^T \mathbf{x} \approx Z^{-1} \sum_j w_{j,i} \nabla \mathbf{s}(\boldsymbol{\xi}_j)^T \mathbf{x} = \nabla \left[Z^{-1} \sum_j w_{j,i} f(\mathbf{s}(\boldsymbol{\xi}_j)) \right]. \quad (15)$$

⁶In our experiments on gravitational wave astronomy, we introduce an additional quantization step, computing approximate Jacobians on a regular grid $\hat{\boldsymbol{\xi}}_1, \dots, \hat{\boldsymbol{\xi}}_{N'}$ of points in the parameter space Ξ .




Figure 2: Illustration of 2-dim signal embeddings and the parameter optimization procedure for gravitational wave signals.

i.e., $\widetilde{\nabla} \mathbf{s}^\top$ is an approximate gradient for a smoothed version \widetilde{f} of the objective f . Figure 2 illustrates smoothed optimization landscapes \widetilde{f} for different levels of smoothing, i.e., different choices of λ_s . In general, the more smoothing is applied, the broader the basin of attraction. We employ a coarse-to-fine approach, which starts with a highly smoothed landscape (small λ_s) in the first iteration and decreases the level of smoothing from iteration to iteration — see Figure 2.

These observations are in line with theory: because our embedding approximately preserves Euclidean distances, $\|\xi_i - \xi_j\|_2 \approx \|s_i - s_j\|_2$, we have

$$\widetilde{f}(\mathbf{s}(\xi_i)) = Z^{-1} \sum_j \Theta(\xi_i, \xi_j) \langle s_j, \mathbf{x} \rangle \approx \langle Z^{-1} \sum_j \Theta(s_i, s_j) s_j, \mathbf{x} \rangle, \quad (16)$$

i.e., applying kernel smoothing in the parameter space is nearly equivalent to applying kernel smoothing to the signal manifold S . This smoothing operation expands the basin of attraction $\Delta = 1/\kappa$, by reducing the manifold curvature κ . Empirically, we find that with appropriate smoothing often a single initialization suffices for convergence to global optimality, suggesting this as a potential key to breaking the curse of dimensionality.

5 TRAINING NONPARAMETRIC TPOPT

In the section above, we described nonparametric TpopT for finding the matching template by the iterative gradient solver (5). Note that this framework requires pre-computing the Jacobians $\nabla \mathbf{s}(\xi)$ and determining optimization hyperparameters, including the step sizes α_k and kernel width parameters λ_k at each layer. In this section, we adapt TpopT into a trainable architecture, which essentially learns all the above quantities from data to further improve performance.

Recall the gradient descent iteration (14) in TpopT. Notice that if we define a collection of matrices $\mathbf{W}(\xi_i, k) = \alpha_k \widetilde{\nabla} \mathbf{s}(\xi_i)^\top \in \mathbb{R}^{d \times D}$ indexed by $\xi_i \in \{\xi_1, \dots, \xi_N\}$ and $k \in \{1, \dots, K\}$ where K is the total number of iterations, then the iteration can be rewritten as

$$\xi^{k+1} = C^{-1} \sum_{i=1}^N w_{k,i} (\xi_i + \mathbf{W}(\xi_i, k) \mathbf{x}), \quad (17)$$

where $w_{k,i} = \Theta_{\lambda_k, \delta_k}(\xi^k, \xi_i)$, and $C = \sum_i w_{k,i}$. Equation (17) can be interpreted as a kernel interpolated gradient step, where the \mathbf{W} matrices summarize the Jacobian and step size information. Because Θ is compactly supported, this sum involves only a small subset of the sample points ξ_i . Now, if we “unroll” the optimization by viewing each gradient descent iteration as one layer of a trainable network, we arrive at a trainable TpopT architecture, as illustrated in Figure 3. Here the trainable parameters in the network are the $\mathbf{W}(\xi_i, k)$ matrices and the kernel width parameters λ_k .

Following our heuristic that the $\mathbf{W}(\xi_i, k)$ matrices were originally the combination of Jacobian and step size, we can initialize these matrices as $\alpha_k \widetilde{\nabla} \mathbf{s}(\xi_i)^\top$. For the loss function during training, we use the square loss between the network output $\xi^K(\mathbf{x})$ and the optimal quantization point $\xi^*(\mathbf{x}) = \arg \max_{i=1, \dots, N} \langle \mathbf{s}(\xi_i), \mathbf{x} \rangle$, namely

$$L = \frac{1}{N_{\text{train}}} \sum_{j=1}^{N_{\text{train}}} \|\xi^K(\mathbf{x}_j) - \xi^*(\mathbf{x}_j)\|_2^2 \quad (18)$$

for a training set $\{\mathbf{x}_j\}_{j=1}^{N_{\text{train}}}$ with positively-labeled data only. This loss function is well-aligned with the signal estimation task, and is also applicable to detection.




Figure 3: Architecture of trainable TpopT. The model takes x as input and starts with a fixed initialization ξ^0 , and outputs ξ^K after going through K layers. The trainable parameters are the collection of $\mathbf{W}(\xi_i, k)$ matrices and kernel width parameters λ_k .

In summary, the trainable TpopT architecture consists of the following steps:

- Create embeddings $s_i \mapsto \xi_i$.
- Estimate Jacobians $\nabla s(\xi)$ at points ξ_i by weighted least squares.
- Estimate smoothed Jacobians $\widetilde{\nabla} s(\xi)$ at any ξ by kernel smoothing.
- Select a multi-level smoothing scheme.
- Train the model with unrolled optimization.

6 EXPERIMENTS

We apply the trainable TpopT to gravitational wave detection, where MF is the current method of choice, and show a significant improvement in efficiency-accuracy tradeoffs. We further demonstrate its wide applicability on low-dimensional data with experiments on handwritten digit data.

To compare the efficiency-accuracy tradeoffs of MF and TpopT models, we note that (i) for MF, the computation cost of the statistic $\max_{i=1, \dots, n} \langle s_i, x \rangle$ is dominated by the cost of n length D inner products, requiring nD multiplication operations. For TpopT, with M parallel initializations, K iterations of the gradient descent (17), m neighbors in the truncated kernel, and a final evaluation of the statistic, we require $MD(Kdm + 1)$ multiplications; other operations including the kernel interpolation and look-up of pre-computed gradients have negligible test-time cost.

6.1 GRAVITATIONAL WAVE DETECTION

We aim to detect a family of gravitational wave signals in Gaussian noise. Each gravitational wave signal is a one-dimensional chirp-like signal – see Figure 4 (left).⁷ Please refer to section F in the appendix for data generation details.

Based on their physical modeling, gravitational wave signals are equipped with a set of physical parameters, such as the masses and three-dimensional spins of the binary black holes that generate them, etc. While it is tempting to directly optimize on this native parameter space, unfortunately the optimization landscape on this space turns out to be rather unfavorable, as shown in Figure 4 (right). We see that the objective function has many spurious local optimizers and is poorly conditioned. Therefore, we still resort to signal embedding to create an alternative set of approximate “parameters” that are better suited for optimization.

For the signal embedding, we apply PCA with dimension 2 on a separate set of 30,000 noiseless waveforms drawn from the same distribution. Because the embedding dimension is relatively low, here we quantize the embedding parameter space with an evenly-spaced grid, with the range of each dimension evenly divided into 30 intervals. The value ξ^0 at the initial layer of TpopT is fixed at the center of this quantization grid. Prior to training, we first determine the optimization hyperparameters

⁷The raw data of gravitational wave detection is a noisy one-dimensional time series, where gravitational wave signals can occur at arbitrary locations. We simplify the problem by considering input segments of fixed time duration.




Figure 4: **Left:** Example of a gravitational wave signal. **Right:** Optimization landscape in the physical parameter space (mass-spin- z), shown as the heatmap of signal correlations.

(step sizes and smoothing levels) using a layer-wise greedy grid search, where we sequentially choose the step size and smoothing level at each layer as if it were the final layer. This greedy approach significantly reduces the cost of the search. From there, we use these optimization hyperparameters to initialize the trainable TpopT network, and train the parameters on the training set. We use the Adam Kingma & Ba (2014) optimizer with batch size 1000 and constant learning rate 10^{-2} . Regarding the computational cost of TpopT, we have $M = 1$, $d = 2$, $m = 4$ during training and $m = 1$ during testing. The test time complexity of K -layer TpopT is $D(2K + 1)$.

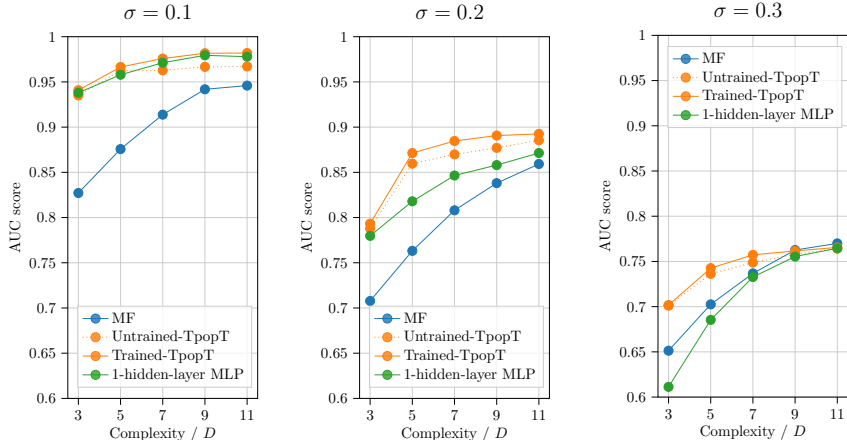


Figure 5: This figure compares the performance of four methods: (1) matched filtering (MF), (2) Template optimization (TpopT) without training, (3) TpopT with training, and (4) multi-layer perceptron (MLP) with one hidden layer. All methods are compared at three noise levels. We see that TpopT performs well in low to moderate noise, which matches theoretical results.

To evaluate the performance of matched filtering at any given complexity m , we randomly generate 1,000 independent sets of m templates drawn from the above distribution, evaluate the ROC curves of each set of templates on the validation set, and select the set with the highest area-under-curve (AUC) score. This selected template bank is then compared with TpopT on the shared test set.

Figure 6.1 shows the comparison of efficiency-accuracy tradeoffs for this task between matched filtering and TpopT after training. We see that TpopT achieves significantly higher detection accuracy compared with MF at equal complexity. At low to moderate noise levels, Trained-TpopT performs the best, followed by MLP, and matched filtering performs the worst. As noise level increases, MLP’s performance worsens most significantly, becoming the worst at $\sigma = 0.3$.

6.2 HANDWRITTEN DIGIT RECOGNITION

In this second experiment, we apply TpopT to the classic task of handwritten digit recognition using the MNIST Deng (2012) dataset, in particular detecting the digit 3 from all other digits. We apply random Euclidean transformations to all images, with translation uniformly between ± 0.1 image size on both dimensions and rotation angle uniformly between $\pm 30^\circ$.

Since the signal space here is nonparametric, we first create a 3-dimensional PCA embedding from the training set, and Figure 6 (left) shows a slice of the embedding projected onto the first two embedding dimensions. See supplementary for experiment details. Regarding the computational cost of TpopT, we have $M = 1$, $d = 3$, $m = 5$ during training and $m = 1$ during testing. Since the complexity is measured at test time, the complexity with K -layer TpopT is $D(3K + 1)$. Additional experimental details can be found in F.

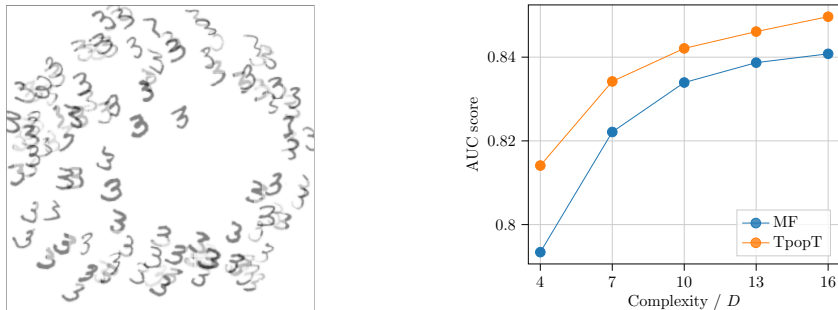


Figure 6: **Left:** A slice of the 3-d embeddings projected onto the first two dimensions. **Right:** Classification scores of MF and TpopT at different complexity levels, for handwritten digit recognition.

Matched filtering is also evaluated similarly as in the previous experiment. We first set aside a random subset of 500 images of digit 3 from the MNIST training set and construct the validation set from it. The remaining images are used to randomly generate 1,000 independent sets of transformed digits 3, and the best-performing set of templates on the validation set is selected as the MF template bank, and compared with TpopT on the shared test set. Figure 6 (right) shows the comparison of efficiency-accuracy tradeoffs between the two methods, and we see a consistently higher detection accuracy of trained TpopT over MF at equal complexities.

7 DISCUSSION AND LIMITATIONS

In this paper, we studied TpopT as an approach to efficient detection of low-dimensional signals. We provided a proof of convergence of Riemannian gradient descent on the signal manifold, and demonstrated its superior dimension scaling compared to MF. We also proposed the trainable TpopT architecture that can handle general nonparametric families of signals. Experimental results show that trained TpopT achieves significantly improved efficiency-accuracy tradeoffs than MF, especially in the gravitational wave detection task where MF is the method of choice.

The principal limitation of nonparametric TpopT is its storage complexity: it represents the manifold using a dense collection of points and Jacobians, with cost exponential in intrinsic dimension d . At the same time, we note that the same exponential storage complexity is encountered by matched filtering with a pre-designed template bank. In some sense, this exponential resource requirement reflects an intrinsic constraint of the signal detection problem, unless more structures within the signal space can be exploited. Both TpopT and its nonparametric extension achieve exponential improvements in test-time efficiency compared to MF; nevertheless, our theoretical results retain an exponential dependence on intrinsic dimension d , due to the need for multiple initializations. In experiments, the proposed smoothing allows convergence to global optimality from a single initialization. Our current theory does not fully explain this observation; this is an important direction for future work.

An advantage of MF not highlighted in this paper is its efficiency in handling noisy time series, using the fast Fourier transform. This enables MF to rapidly locate signals that occur at a-priori unknown spatial/temporal locations. Developing a convolutional version of TpopT with similar advantages is another important direction.

Finally, our gravitational wave experiments use synthetic data with known ground truth, in order to corroborate the key messages of this paper. Future experiments that explore broader and more realistic setups will be an important empirical validation of the proposed method.

REFERENCES

- B. P. Abbott et al. Advanced LIGO. *Classical and Quantum Gravity*, 32:074001, Apr 2015. doi: 10.1088/0264-9381/32/7/074001.
- Benjamin P Abbott, Richard Abbott, TD Abbott, MR Abernathy, Fausto Acernese, Kendall Ackley, Carl Adams, Thomas Adams, Paolo Addesso, RX Adhikari, et al. Observation of gravitational waves from a binary black hole merger. *Physical review letters*, 116(6):061102, 2016.
- Benjamin P Abbott, Rich Abbott, TDea Abbott, Fausto Acernese, Kendall Ackley, Carl Adams, Thomas Adams, Paolo Addesso, RX Adhikari, Vaishali B Adya, et al. Gw170817: observation of gravitational waves from a binary neutron star inspiral. *Physical review letters*, 119(16):161101, 2017.
- Alex Abramovici, William E. Althouse, Ronald W. P. Drever, Yekta Gursel, Seiji Kawamura, Frederick J. Raab, David Shoemaker, Lisa Sievers, Robert E. Spero, Kip S. Thorne, Rochus E. Vogt, Rainer Weiss, Stanley E. Whitcomb, and Michael E. Zucker. LIGO: The Laser Interferometer Gravitational-Wave Observatory. *Science*, 256(5055):325–333, April 1992. doi: 10.1126/science.256.5055.325.
- F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32:024001, Jan 2015. doi: 10.1088/0264-9381/32/2/024001.
- T Akutsu, M Ando, K Arai, Y Arai, S Araki, A Araya, N Aritomi, Y Aso, S Bae, Y Bae, et al. Overview of kagra: Detector design and construction history. *Progress of Theoretical and Experimental Physics*, 2021(5):05A101, 2021.
- Richard G Baraniuk and Michael B Wakin. Random projections of smooth manifolds. *Foundations of computational mathematics*, 9(1):51–77, 2009.
- Yoshua Bengio and Martin Monperrus. Non-local manifold tangent learning. *Advances in Neural Information Processing Systems*, 17, 2004.
- Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
- Emery N Brown, Robert E Kass, and Partha P Mitra. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature neuroscience*, 7(5):456–461, 2004.
- Sam Buchanan, Jingkai Yan, Ellie Haber, and John Wright. Resource-efficient invariant networks: Exponential gains by unrolled optimization. *arXiv preprint arXiv:2203.05006*, 2022.
- Enrico Caffagni, David W Eaton, Joshua P Jones, and Mirko van der Baan. Detection and analysis of microseismic events using a matched filtering algorithm (mfa). *Geophysical Journal International*, 206(1):644–658, 2016.
- Tianlong Chen, Xiaohan Chen, Wuyang Chen, Zhangyang Wang, Howard Heaton, Jialin Liu, and Wotao Yin. Learning to optimize: A primer and a benchmark. *The Journal of Machine Learning Research*, 23(1):8562–8620, 2022.
- Benjamin Culpepper and Bruno Olshausen. Learning transport operators for image manifolds. *Advances in neural information processing systems*, 22, 2009.
- Tito Dal Canton, Alexander H. Nitz, Bhooshan Gadre, Gareth S. Cabourn Davies, Verónica Villa-Ortega, Thomas Dent, Ian Harry, and Liting Xiao. Real-time Search for Compact Binary Mergers in Advanced LIGO and Virgo’s Third Observing Run Using PyCBC Live. , 923(2):254, December 2021. doi: 10.3847/1538-4357/ac2f9a.
- Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- Steven Diamond, Vincent Sitzmann, Felix Heide, and Gordon Wetzstein. Unrolled optimization with deep priors. *arXiv preprint arXiv:1705.08041*, 2017.

- Elizabeth A Gibson, Junhui Zhang, Jingkai Yan, Lawrence Chillrud, Jaime Benavides, Yanelli Nunez, Julie B Herbstman, Jeff Goldsmith, John Wright, and Marianthi-Anna Kioumourtzoglou. Principal component pursuit for pattern identification in environmental mixtures. *Environmental Health Perspectives*, 130(11):117008, 2022.
- Carl W Helstrom. *Statistical theory of signal detection: international series of monographs in electronics and instrumentation*, volume 9. Elsevier, 2013.
- Hui Ji, Chaoqiang Liu, Zuowei Shen, and Yuhong Xu. Robust video denoising using low rank matrix completion. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pp. 1791–1798. IEEE, 2010.
- DG Johnson. Complex scatterer reconstruction using multistatic spherical wave isar fourier template matching. In *2009 International Conference on Electromagnetics in Advanced Applications*, pp. 291–294. IEEE, 2009.
- Angjoo Kanazawa, Abhishek Sharma, and David Jacobs. Locally scale-invariant convolutional neural networks. *arXiv preprint arXiv:1412.5104*, 2014.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Abhishek Kumar, Prasanna Sattigeri, and Tom Fletcher. Semi-supervised learning with gans: Manifold invariance with improved inference. *Advances in neural information processing systems*, 30, 2017.
- John M Lee. *Introduction to Riemannian manifolds*, volume 2. Springer, 2018.
- Dong Liang, Jing Cheng, Ziwen Ke, and Leslie Ying. Deep mri reconstruction: unrolled optimization algorithms meet neural networks. *arXiv preprint arXiv:1907.11711*, 2019.
- Jialin Liu and Xiaohan Chen. Alista: Analytic weights are as good as learned weights in lista. In *International Conference on Learning Representations (ICLR)*, 2019.
- Dalton Lunga, Saurabh Prasad, Melba M Crawford, and Okan Ersoy. Manifold-learning-based feature extraction for classification of hyperspectral data: A review of advances in manifold learning. *IEEE Signal Processing Magazine*, 31(1):55–66, 2013.
- Farzin Mokhtarian and Sadegh Abbasi. Shape similarity retrieval under affine transforms. *Pattern Recognition*, 35(1):31–41, 2002.
- Vishal Monga, Yuelong Li, and Yonina C Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.
- Timothy S Murphy, Marcus J Holzinger, and Brien Flewelling. Space object detection in images using matched filter bank and bayesian update. *Journal of Guidance, Control, and Dynamics*, 40(3):497–509, 2017.
- Alex Nitz, Ian Harry, Duncan Brown, Christopher M. Biwer, Josh Willis, Tito Dal Canton, Collin Capano, Thomas Dent, Larne Pekowsky, Soumi De, Miriam Cabero, Gareth S Cabourn Davies, Andrew R. Williamson, Duncan Macleod, Bernd Machenschalk, Francesco Pannarale, Prayush Kumar, Steven Reyes, dfinstad, Sumit Kumar, Shichao Wu, Márton Tápai, Leo Singer, veronica villa, Sebastian Khan, Stephen Fairhurst, Koustav Chandra, Alex Nielsen, Shashwat Singh, and Thomas Massinger. gwastro/pycbc: v2.1.2 release of pycbc, May 2023. URL <https://doi.org/10.5281/zenodo.7885796>.
- Benjamin J Owen and Bangalore Suryanarayana Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Physical Review D*, 60(2):022002, 1999.
- Souradeep Pal and K Rajesh Nayak. Swarm-intelligent search for gravitational waves from eccentric binary mergers. *arXiv e-prints*, art. arXiv:2307.03736, July 2023. doi: 10.48550/arXiv.2307.03736.

- Thottempudi Pardhu, A Kavya Sree, and K Tanuja. Design of matched filter for radar applications. *Electrical and Electronics Engineering: An International Journal (ELELIJ) Vol, 3*, 2014.
- Mijung Park, Wittawat Jitkrittum, Ahmad Qamar, Zoltán Szabó, Lars Buesing, and Maneesh Sahani. Bayesian manifold learning: the locally linear latent variable model (ll-lvm). *Advances in neural information processing systems*, 28, 2015.
- Kenia Picos, Victor H Diaz-Ramirez, Vitaly Kober, Antonio S Montemayor, and Juan J Pantrigo. Accurate three-dimensional pose recognition from monocular images using template matched filtering. *Optical Engineering*, 55(6):063102–063102, 2016.
- Xiaojun Quan, Chunyu Kit, Yong Ge, and Sinno Jialin Pan. Short and sparse text topic modeling via self-aggregation. In *24th International Joint Conference on Artificial Intelligence, IJCAI 2015*, pp. 2270–2276. AAAI Press/International Joint Conferences on Artificial Intelligence, 2015.
- Salah Rifai, Pascal Vincent, Xavier Muller, Xavier Glorot, and Yoshua Bengio. Contractive auto-encoders: Explicit invariance during feature extraction. In *Proceedings of the 28th international conference on international conference on machine learning*, pp. 833–840, 2011.
- Baptiste Rousset, Michel Campillo, Cecile Lasserre, William Benjamin Frank, Nathalie Cotte, Andrea Walpersdorf, Anne Socquet, and Vladimir Kostoglodov. A geodetic matched filter search for slow slip with application to the mexico subduction zone. *Journal of Geophysical Research: Solid Earth*, 122(12):10–498, 2017.
- Soumen Roy, Anand S. Sengupta, and Nilay Thakor. Hybrid geometric-random template-placement algorithm for gravitational wave searches from compact binary coalescences. , 95(10):104045, May 2017. doi: 10.1103/PhysRevD.95.104045.
- Soumen Roy, Anand S. Sengupta, and Parameswaran Ajith. Effectual template banks for upcoming compact binary searches in Advanced-LIGO and Virgo data. , 99(2):024048, January 2019. doi: 10.1103/PhysRevD.99.024048.
- Yulin Shi, Zoran Nenadic, and Xiangmin Xu. Novel use of matched filtering for synaptic event detection and extraction. *PLoS One*, 5(11):e15517, 2010.
- Kihyuk Sohn and Honglak Lee. Learning invariant representations with local transformations. *arXiv preprint arXiv:1206.6418*, 2012.
- Victor A Toponogov. *Differential geometry of curves and surfaces*. Springer, 2006.
- Ramon Van Handel. *Probability in High Dimension*, volume 1. APC550 Lecture Notes Princeton University, 2016.
- Michael B Wakin. *The geometry of low-dimensional signal models*. PhD thesis, Rice University, 2007.
- Michael B Wakin, David L Donoho, Hyeokho Choi, and Richard G Baraniuk. The multiscale structure of non-differentiable image manifolds. In *Wavelets XI*, volume 5914, pp. 413–429. SPIE, 2005.
- Thilina S. Weerathunga and Soumya D. Mohanty. Performance of particle swarm optimization on the fully-coherent all-sky search for gravitational waves from compact binary coalescences. , 95(12):124030, June 2017. doi: 10.1103/PhysRevD.95.124030.
- Bo Xin, Yizhou Wang, Wen Gao, David Wipf, and Baoyuan Wang. Maximal sparsity with deep networks? *Advances in Neural Information Processing Systems*, 29, 2016.
- Jingkai Yan, Mariam Avagyan, Robert E Colgan, Doğa Veske, Imre Bartos, John Wright, Zsuzsa Márka, and Szabolcs Márka. Generalized approach to matched filtering using neural networks. *Physical Review D*, 105(4):043006, 2022a.
- Jingkai Yan, Robert Colgan, John Wright, Zsuzsa Márka, Imre Bartos, and Szabolcs Márka. Boosting the efficiency of parametric detection with hierarchical neural networks. *Physical Review D*, 106(6):063008, 2022b.
- Yanzhao Zhou, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Oriented response networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 519–528, 2017.

A APPENDIX

B OVERVIEW

In the appendices, we will prove Theorem 1 from the main paper.

For the rest of the supplementary materials, Section C proves result (10) under the stricter constraint on the noise level σ , Section D bounds the effect of noise on the tangent bundle, and Section E uses this bound to prove result (9) under the looser constraint on the noise level.

C PROOF OF RESULT (10)

In this section, we state and prove one of the two parts of our main claims about gradient descent:

Theorem 2. *Let S be a complete manifold. Suppose the extrinsic geodesic curvature of S is bounded by κ . Consider the Riemannian gradient method, with initialization satisfying $d(\mathbf{s}^0, \mathbf{s}_\natural) < \Delta = 1/\kappa$, and step size $\tau = \frac{1}{64}$. Then when $\sigma \leq 1/(60\kappa\sqrt{D})$, with probability at least $1 - e^{-D/2}$, we have for all k*

$$d(\mathbf{s}^k, \mathbf{s}^*) \leq (1 - \epsilon)^k d(\mathbf{s}^0, \mathbf{s}^*), \quad (19)$$

where \mathbf{s}^* is the unique minimizer of f over $B(\mathbf{s}_\natural, 1/\kappa)$. Here, c, ϵ are positive numerical constants.

Proof. Since the closed neighborhood $B(\mathbf{s}_\natural, 1/\kappa)$ is a compact set and f is continuous, there must exist a minimizer of f on $B(\mathbf{s}_\natural, 1/\kappa)$, which we denote as \mathbf{s}^* . We will show that with high probability \mathbf{s}^* does not lie on the boundary $\partial B(\mathbf{s}_\natural, 1/\kappa)$. It suffices to show that $\forall \mathbf{s} \in \partial B(\mathbf{s}_\natural, 1/\kappa)$:

$$\left\langle -\text{grad}[f](\mathbf{s}), \frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2} \right\rangle > 0, \quad (20)$$

namely that the gradient descent direction points inward the neighborhood for all points on the boundary. Here $\log_{\mathbf{s}} : S \rightarrow T_{\mathbf{s}}S$ denotes the logarithmic map at point $\mathbf{s} \in S$. To show this, we have

$$\begin{aligned} \left\langle -\text{grad}[f](\mathbf{s}), \frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2} \right\rangle &= \left\langle P_{T_{\mathbf{s}}S}[\mathbf{s}_\natural + \mathbf{z}], \frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2} \right\rangle \\ &= \left\langle \mathbf{s}_\natural + \mathbf{z}, \frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2} \right\rangle \\ &\geq \left\langle \mathbf{s}_\natural, \frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2} \right\rangle - \|\mathbf{z}\|_2, \end{aligned} \quad (21)$$

where the operator $P_{T_{\mathbf{s}}S}[\cdot]$ denotes projection onto the tangent space at \mathbf{s} , and we used the fact that $\log_{\mathbf{s}} \mathbf{s}_\natural \in T_{\mathbf{s}}S$. Let γ be a unit-speed geodesic of S with $\gamma(0) = \mathbf{s}_\natural$ and $\gamma(\Delta) = \mathbf{s}$, the existence of which is ensured by the completeness of S . Hence $\dot{\gamma}(\Delta) = -\frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2}$. Using Lemma 3, it follows that

$$\left\langle \mathbf{s}_\natural, \frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2} \right\rangle = \langle \gamma(0), -\dot{\gamma}(\Delta) \rangle \geq \Delta - \frac{1}{6}\kappa^2\Delta^3 = \frac{5}{6\kappa}. \quad (22)$$

Throughout this proof, we will use the result from measure concentration that $\|\mathbf{z}\|_2 \leq 2\sigma\sqrt{D}$ with probability at least $1 - e^{-D/2}$. Hence with high probability we have $\left\langle -\text{grad}[f](\mathbf{s}), \frac{\log_{\mathbf{s}} \mathbf{s}_\natural}{\|\log_{\mathbf{s}} \mathbf{s}_\natural\|_2} \right\rangle \geq \frac{5}{6\kappa} - \frac{1}{30\kappa} > 0$. Therefore, with high probability \mathbf{s}^* lies in the interior of $B(\mathbf{s}_\natural, 1/\kappa)$, and hence the gradient vanishes at \mathbf{s}^* , i.e. $\text{grad}[f](\mathbf{s}^*) = \mathbf{0}$.

Suppose we are currently at the k -th iteration with iterate \mathbf{s}^k . Define $\mathbf{s}_t = \exp_{\mathbf{s}^k}(-t \text{grad}[f](\mathbf{s}^k))$ with variable $t \in [0, \tau]$, and the next iterate can be represented as $\mathbf{s}^{k+1} = \mathbf{s}_\tau$. The global definition

of the exponential map is ensured by the completeness of S . We have that

$$\begin{aligned}
d(\mathbf{s}^{k+1}, \mathbf{s}^*) - d(\mathbf{s}^k, \mathbf{s}^*) &= \int_0^\tau \frac{d}{dr} d(\mathbf{s}_r, \mathbf{s}^*) \Big|_t dt \\
&= \int_0^\tau \left\langle \frac{d}{dr} \mathbf{s}_r \Big|_t, \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle dt \\
&= \int_0^\tau \left\langle \Pi_{\mathbf{s}_t, \mathbf{s}^k} \{-\text{grad}[f](\mathbf{s}^k)\}, \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle dt \\
&= \int_0^\tau \left\langle -\text{grad}[f](\mathbf{s}_t), \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle dt \\
&\quad + \int_0^\tau \left\langle \text{grad}[f](\mathbf{s}_t) + \Pi_{\mathbf{s}_t, \mathbf{s}^k} \{-\text{grad}[f](\mathbf{s}^k)\}, \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle dt \\
&= \int_0^\tau \left\langle P_{T_{\mathbf{s}_t} S}[\mathbf{s}^*], \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle dt \\
&\quad + \int_0^\tau \left\langle P_{T_{\mathbf{s}_t} S}[\mathbf{x} - \mathbf{s}^*], \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle dt \\
&\quad + \int_0^\tau \left\langle \text{grad}[f](\mathbf{s}_t) + \Pi_{\mathbf{s}_t, \mathbf{s}^k} \{-\text{grad}[f](\mathbf{s}^k)\}, \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle dt.
\end{aligned} \tag{23}$$

The third equation holds because the velocity at the new point \mathbf{s}_t is the same velocity vector at \mathbf{s}^k but transported along the curve since there is no acceleration along the curve. The last equation follows from the fact that $\text{grad}[f](\mathbf{s}_t) = -P_{T_{\mathbf{s}_t} S}[\mathbf{x}]$. In the following, we will bound the three terms in (23) separately.

For convenience, write

$$d(t) = d(\mathbf{s}_t, \mathbf{s}^*) \tag{24}$$

so that $d(0) = d(\mathbf{s}^k, \mathbf{s}^*)$ and $d(\tau) = d(\mathbf{s}^{k+1}, \mathbf{s}^*)$.

For the integrand of the first term in (23), let γ be a unit-speed geodesic between \mathbf{s}_t and \mathbf{s}^* , where $\gamma(0) = \mathbf{s}^*$ and $\gamma(d(t)) = \mathbf{s}_t$. We have

$$\begin{aligned}
\left\langle P_{T_{\mathbf{s}_t} S}[\mathbf{s}^*], \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle &= \left\langle \mathbf{s}^*, \frac{-\log_{\mathbf{s}_t}(\mathbf{s}^*)}{\|\log_{\mathbf{s}_t}(\mathbf{s}^*)\|_2} \right\rangle \\
&= \langle \gamma(0), \dot{\gamma}(d(t)) \rangle \\
&\leq -d(t) + \frac{1}{6} \kappa^2 d^3(t),
\end{aligned} \tag{25}$$

where we used the fact that $\dot{\gamma}(d(t)) = \frac{-\log_{\mathbf{s}_t}(\mathbf{s}^*)}{\|\log_{\mathbf{s}_t}(\mathbf{s}^*)\|_2} \in T_{\mathbf{s}_t} S$, and the inequality is given by Lemma 3.

For the integrand of the second term in (23), we have

$$\begin{aligned}
\left\langle P_{T_{\mathbf{s}_t} S}[\mathbf{x} - \mathbf{s}^*], \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle &\leq \|P_{T_{\mathbf{s}_t} S}[\mathbf{x} - \mathbf{s}^*]\|_2 \\
&= \|P_{T_{\mathbf{s}_t} S} P_{(T_{\mathbf{s}^*} S)^\perp}[\mathbf{x} - \mathbf{s}^*]\|_2 \\
&\leq \|P_{T_{\mathbf{s}_t} S} P_{(T_{\mathbf{s}^*} S)^\perp}\| \|\mathbf{x} - \mathbf{s}^*\|_2 \\
&\leq \|P_{(T_{\mathbf{s}^*} S)^\perp} P_{T_{\mathbf{s}_t} S}\| \|\mathbf{z}\|_2,
\end{aligned} \tag{26}$$

where we used the optimality of \mathbf{s}^* and the symmetry of projection operators. The operator norm $\|P_{(T_{\mathbf{s}^*} S)^\perp} P_{T_{\mathbf{s}_t} S}\|$ can be rewritten as

$$\|P_{(T_{\mathbf{s}^*} S)^\perp} P_{T_{\mathbf{s}_t} S}\| = \sup_{\mathbf{v} \in T_{\mathbf{s}_t} S, \|\mathbf{v}\|_2=1} d(\mathbf{v}, T_{\mathbf{s}^*} S). \tag{27}$$

For any unit vector $\mathbf{v} \in T_{\mathbf{s}_t} S$, we will construct a vector in $T_{\mathbf{s}^*} S$ and use its distance from \mathbf{v} to upper bound $d(\mathbf{v}, T_{\mathbf{s}^*} S)$. We again use the unit-speed geodesic γ joining \mathbf{s}^* and \mathbf{s}_t , where

$\gamma(0) = \mathbf{s}^*$ and $\gamma(d(t)) = \mathbf{s}_t$. Let $\mathbf{v}_r = \mathcal{P}_{r,d(t)}\mathbf{v}$ for $r \in [0, d(t)]$, where $\mathcal{P}_{r,d(t)}$ denotes the parallel transport backward along γ . The derivative of \mathbf{v}_r can be expressed by the second fundamental form $\frac{d}{dr}\mathbf{v}_r = \mathbb{I}(\dot{\gamma}(r), \mathbf{v}_r)$, which can further be bounded by Lemma 4 to get $\|\frac{d}{dr}\mathbf{v}_r\|_2 \leq 3\kappa$. Hence $\|\mathbf{v}_{d(t)} - \mathbf{v}_0\|_2 \leq 3\kappa d(t)$. Since $\mathbf{v}_{d(t)} = \mathbf{v}$ and $\mathbf{v}_0 \in T_{\mathbf{s}^*}S$, it follows that $d(\mathbf{v}, T_{\mathbf{s}^*}S) \leq 3\kappa d(t)$ for any unit vector $\mathbf{v} \in T_{\mathbf{s}_t}S$. Hence

$$\|P_{(T_{\mathbf{s}^*}S)^\perp} P_{T_{\mathbf{s}_t}S}\| \leq 3\kappa \cdot d(t). \quad (28)$$

Since $\|\mathbf{z}\|_2 \leq 2\sigma\sqrt{D}$ with high probability, plugging these into (26), we have with high probability

$$\left\langle P_{T_{\mathbf{s}_t}S}[\mathbf{x} - \mathbf{s}^*], \frac{-\log_{\mathbf{s}_t} \mathbf{s}^*}{\|\log_{\mathbf{s}_t} \mathbf{s}^*\|_2} \right\rangle \leq 6\sigma\kappa\sqrt{D} \cdot d(t). \quad (29)$$

The integrand of the third term in (23) can be bounded using the Riemannian Hessian. We have

$$\text{grad}[f](\mathbf{s}_t) = \Pi_{\mathbf{s}_t, \mathbf{s}^k} \{\text{grad}[f](\mathbf{s}^k)\} + \int_{r=0}^t \Pi_{\mathbf{s}_t, \mathbf{s}_r} \text{Hess}[f](\mathbf{s}_r) \Pi_{\mathbf{s}_r, \mathbf{s}^k} \{-\text{grad}[f](\mathbf{s}^k)\} dr. \quad (30)$$

Using the L -Lipschitz gradient property of the function f from Lemma 5, we have

$$\begin{aligned} \left\langle \text{grad}[f](\mathbf{s}_t) + \Pi_{\mathbf{s}_t, \mathbf{s}^k} \{-\text{grad}[f](\mathbf{s}^k)\}, \frac{-\log_{\mathbf{s}_r} \mathbf{s}^*}{\|\log_{\mathbf{s}_r} \mathbf{s}^*\|_2} \right\rangle &\leq \|\text{grad}[f](\mathbf{s}_t) - \Pi_{\mathbf{s}_t, \mathbf{s}^k} \{\text{grad}[f](\mathbf{s}^k)\}\|_2 \\ &\leq t \max_{\bar{\mathbf{s}}} \|\text{Hess}[f](\bar{\mathbf{s}})\| \|\text{grad}[f](\mathbf{s}^k)\|_2 \\ &\leq tL^2 d(\mathbf{s}^k, \mathbf{s}^*). \end{aligned} \quad (31)$$

Hence

$$\int_0^\tau \left\langle \text{grad}[f](\mathbf{s}_t) + \Pi_{\mathbf{s}_t, \mathbf{s}^k} \{-\text{grad}[f](\mathbf{s}^k)\}, \frac{-\log_{\mathbf{s}_r} \mathbf{s}^*}{\|\log_{\mathbf{s}_r} \mathbf{s}^*\|_2} \right\rangle dt \leq \frac{1}{2} \tau^2 L^2 d(\mathbf{s}^k, \mathbf{s}^*). \quad (32)$$

Gathering the separate bounds of the three terms in (23), we have

$$\begin{aligned} d(\tau) &\leq d(0) + \int_0^\tau \left(-d(t) + \frac{1}{6} \kappa^2 d^3(t) + 6\sigma\kappa\sqrt{D}d(t) \right) dt + \frac{1}{2} L^2 \tau^2 d(0) \\ &= (1 + \frac{1}{2} L^2 \tau^2) d(0) + \int_0^\tau \left(-(1 - c_1)d(t) + \frac{1}{6} \kappa^2 d^3(t) \right) dt, \end{aligned} \quad (33)$$

where $c_1 = 6\sigma\kappa\sqrt{D} \leq \frac{1}{10}$. By triangle inequality we have

$$d(\mathbf{s}^k, \mathbf{s}^*) - d(\mathbf{s}^k, \mathbf{s}_t) \leq d(\mathbf{s}_t, \mathbf{s}^*) \leq d(\mathbf{s}^k, \mathbf{s}^*) + d(\mathbf{s}^k, \mathbf{s}_t). \quad (34)$$

Since $\mathbf{s}_t = \exp_{\mathbf{s}^k}(-t \text{grad}[f](\mathbf{s}^k))$, we have

$$d(\mathbf{s}^k, \mathbf{s}_t) \leq t \|\text{grad}[f](\mathbf{s}^k)\|_2 \leq tL \cdot d(\mathbf{s}^k, \mathbf{s}^*), \quad (35)$$

and thus

$$(1 - tL)d(0) \leq d(t) \leq (1 + tL)d(0). \quad (36)$$

Hence for the integrand in (33), we have

$$\begin{aligned} -(1 - c_1)d(t) + \frac{1}{6} \kappa^2 d^3(t) &\leq -(1 - c_1)d(t) + \frac{1}{6} \kappa^2 (1 + tL)^2 d^2(0)d(t) \\ &\leq -(1 - c_1)d(t) + \frac{1}{6} \kappa^2 (1 + \tau L)^2 (2\Delta)^2 d(t) \\ &= (-1 + c_2)d(t) \\ &\leq (-1 + c_2)(1 - tL)d(0) \end{aligned} \quad (37)$$

where $c_2 = c_1 + \frac{2}{3}(1 + \tau L)^2$.

Plugging this back, we get

$$\begin{aligned} d(\tau) &\leq (1 + \frac{1}{2} L^2 \tau^2) d(0) + (-1 + c_2) d(0) \int_0^\tau (1 - tL) dt \\ &= \left(1 - (1 - c_2)\tau + \left(\frac{1}{2} L^2 + \frac{1}{2} L(1 - c_2) \right) \tau^2 \right) d(0). \end{aligned} \quad (38)$$

Substituting in $L = \frac{121}{30}$ from Lemma 5 and $\tau = \frac{1}{64}$, we get

$$d(\mathbf{s}^{k+1}, \mathbf{s}^*) \leq (1 - \epsilon)d(\mathbf{s}^k, \mathbf{s}^*) \quad (39)$$

where $\epsilon \approx 2.3 \times 10^{-4}$, which proves result (10). Note that this also implies the uniqueness of the minimizer \mathbf{s}^* . \square

C.1 SUPPORTING LEMMAS

Lemma 3. *Let γ be a regular unit-speed curve on the manifold $S \subset \mathbb{S}^{d-1}$ with extrinsic curvature κ . Then,*

$$\langle \dot{\gamma}(t), \gamma(0) \rangle \leq -t + \frac{\kappa^2 t^3}{6} \quad (40)$$

Proof. Since $\gamma \subset \mathbb{S}^{d-1}$, by differentiating both sides of $\|\gamma(t)\|_2^2 = 1$ we get $\langle \dot{\gamma}(t), \gamma(t) \rangle = 0$. Further, since γ is unit-speed, we have $\|\dot{\gamma}(t)\|_2^2 = 1$ and by differentiating it $\langle \ddot{\gamma}(t), \dot{\gamma}(t) \rangle = 0$. Therefore,

$$\begin{aligned} \langle \dot{\gamma}(t), \gamma(0) \rangle &= \left\langle \dot{\gamma}(t), \gamma(t) - \int_0^t \dot{\gamma}(t_1) dt_1 \right\rangle \\ &= - \left\langle \dot{\gamma}(t), \int_0^t \dot{\gamma}(t_1) dt_1 \right\rangle \\ &= - \int_{t_1=0}^t \left\langle \dot{\gamma}(t), \dot{\gamma}(t) - \int_{t_2=t_1}^t \ddot{\gamma}(t_2) dt_2 \right\rangle dt_1 \\ &= -t + \int_{t_1=0}^t \int_{t_2=t_1}^t \langle \dot{\gamma}(t), \ddot{\gamma}(t_2) \rangle dt_2 dt_1 \\ &= -t + \int_{t_1=0}^t \int_{t_2=t_1}^t \left\langle \dot{\gamma}(t_2) + \int_{t_3=t_2}^t \ddot{\gamma}(t_3) dt_3, \ddot{\gamma}(t_2) \right\rangle dt_2 dt_1 \\ &= -t + \int_{t_1=0}^t \int_{t_2=t_1}^t \int_{t_3=t_2}^t \langle \ddot{\gamma}(t_3), \ddot{\gamma}(t_2) \rangle dt_3 dt_2 dt_1 \\ &\leq -t + \kappa^2 \int_{t_1=0}^t \int_{t_2=t_1}^t \int_{t_3=t_2}^t dt_3 dt_2 dt_1 \\ &= -t + \frac{\kappa^2 t^3}{6}. \end{aligned} \quad (41)$$

\square

Lemma 4. *Let $\mathbb{I}(\mathbf{u}, \mathbf{v})$ denote the second fundamental form at some point $\mathbf{s} \in S$, and let κ denote the extrinsic (\mathbb{R}^D) geodesic curvature of S . Then*

$$\sup_{\|\mathbf{u}\|_2=1, \|\mathbf{v}\|_2=1} \|\mathbb{I}(\mathbf{u}, \mathbf{v})\|_2 \leq 3\kappa. \quad (42)$$

Proof. Set

$$\kappa^{\mathbb{I}} = \max_{\|\mathbf{u}\|_2=1, \|\mathbf{v}\|_2=1} \|\mathbb{I}(\mathbf{u}, \mathbf{v})\|_2^2. \quad (43)$$

Choose unit vectors \mathbf{u}, \mathbf{v} which realize this maximum value (these must exist, by continuity of \mathbb{I} and compactness of the constraint set). Because \mathbb{I} is bilinear, $\|\mathbb{I}(\mathbf{u}, \mathbf{v})\|_2^2 = \|\mathbb{I}(\mathbf{u}, -\mathbf{v})\|_2^2$, and without loss of generality, we can assume $\langle \mathbf{u}, \mathbf{v} \rangle \leq 0$.

Since $\mathbb{I}(\mathbf{u}, \mathbf{v})$ is a symmetric bilinear form, each coordinate of the vector $\mathbb{I}(\mathbf{u}, \mathbf{v})$ has the form $\mathbb{I}_i(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T \Phi_i \mathbf{v}$ for some symmetric $d \times d$ matrix Φ_i . Now,

$$\mathbf{u}^T \Phi_i \mathbf{v} = \frac{1}{2}(\mathbf{u} + \mathbf{v})^T \Phi_i (\mathbf{u} + \mathbf{v}) - \frac{1}{2}\mathbf{u}^T \Phi_i \mathbf{u} - \frac{1}{2}\mathbf{v}^T \Phi_i \mathbf{v}, \quad (44)$$

so

$$|\frac{1}{2}\mathbf{u}^T\Phi_i\mathbf{u}| + |\frac{1}{2}\mathbf{v}^T\Phi_i\mathbf{v}| + |\frac{1}{2}(\mathbf{u} + \mathbf{v})^T\Phi_i(\mathbf{u} + \mathbf{v})| \geq |\mathbf{u}^T\Phi_i\mathbf{v}| \quad (45)$$

and

$$3|\frac{1}{2}\mathbf{u}^T\Phi_i\mathbf{u}|^2 + 3|\frac{1}{2}\mathbf{v}^T\Phi_i\mathbf{v}|^2 + 3|\frac{1}{2}(\mathbf{u} + \mathbf{v})^T\Phi_i(\mathbf{u} + \mathbf{v})|^2 \geq |\mathbf{u}^T\Phi_i\mathbf{v}|^2 \quad (46)$$

where we have used the inequality $(a + b + c)^2 \leq 3a^2 + 3b^2 + 3c^2$ which follows from convexity of the square. Summing over i , we obtain that

$$\frac{3}{4}\|\mathbb{I}(\mathbf{u}, \mathbf{u})\|_2^2 + \frac{3}{4}\|\mathbb{I}(\mathbf{v}, \mathbf{v})\|_2^2 + \frac{3}{4}\|\mathbb{I}(\mathbf{u} + \mathbf{v}, \mathbf{u} + \mathbf{v})\|_2^2 \geq \|\mathbb{I}(\mathbf{u}, \mathbf{v})\|_2^2 \quad (47)$$

this implies that

$$\frac{9}{4}\max\left\{\|\mathbb{I}(\mathbf{u}, \mathbf{u})\|_2^2, \|\mathbb{I}(\mathbf{v}, \mathbf{v})\|_2^2, \|\mathbb{I}(\mathbf{u} + \mathbf{v}, \mathbf{u} + \mathbf{v})\|_2^2\right\} \geq \|\mathbb{I}(\mathbf{u}, \mathbf{v})\|_2^2 \quad (48)$$

Because \mathbf{u}, \mathbf{v} are unit vectors with $\langle \mathbf{u}, \mathbf{v} \rangle \leq 0$, we have $\|\mathbf{u} + \mathbf{v}\|_2 \leq \sqrt{2}$, and so

$$4\kappa^2 \geq \max\left\{\|\mathbb{I}(\mathbf{u}, \mathbf{u})\|_2^2, \|\mathbb{I}(\mathbf{v}, \mathbf{v})\|_2^2, \|\mathbb{I}(\mathbf{u} + \mathbf{v}, \mathbf{u} + \mathbf{v})\|_2^2\right\}, \quad (49)$$

whence

$$9\kappa^2 \geq \|\mathbb{I}(\mathbf{u}, \mathbf{v})\|_2^2 = (\kappa^{\mathbb{I}})^2, \quad (50)$$

which is the claimed inequality. \square

Lemma 5. Assume $\sigma \leq 1/(60\kappa\sqrt{D})$. The objective function $f(\mathbf{s}) = -\langle \mathbf{s}, \mathbf{x} \rangle$ has L -Lipschitz gradient in a $1/\kappa$ -neighborhood of \mathbf{s}_{\natural} with probability at least $1 - e^{-D/2}$, where $L = \frac{121}{30}$.

Proof. On a Riemannian manifold S , the conditions for L -Lipschitz gradient in a subset can be expressed as $\frac{d^2}{dt^2}(f \circ \gamma)(t) \leq L$ for all unit-speed geodesics $\gamma(t)$ in the subset Boumal (2023).

Let $\Delta = 1/\kappa$, and let $\gamma(t)$ be a unit-speed geodesic of S in the neighborhood $B(\mathbf{s}_{\natural}, \Delta)$, $t \in [0, T]$. The neighborhood constraint implies that $T = d(\gamma(0), \gamma(T)) \leq d(\gamma(0), \mathbf{s}_{\natural}) + d(\gamma(T), \mathbf{s}_{\natural}) \leq 2\Delta$.

To bound the second derivative $\frac{d^2}{dt^2}(f \circ \gamma)(t)$, we have

$$\begin{aligned} \frac{d^2}{dt^2}(f \circ \gamma)(t) &= -\langle \ddot{\gamma}(t), \mathbf{x} \rangle \\ &= -\langle \ddot{\gamma}(t), \gamma(0) \rangle - \langle \ddot{\gamma}(t), \mathbf{s}_{\natural} - \gamma(0) \rangle - \langle \ddot{\gamma}(t), \mathbf{z} \rangle. \end{aligned} \quad (51)$$

The first term can be bounded as

$$\begin{aligned} -\langle \ddot{\gamma}(t), \gamma(0) \rangle &= -\left\langle \ddot{\gamma}(t), \gamma(t) - \int_{t_1=0}^t \dot{\gamma}(t_1) dt_1 \right\rangle \\ &= -\langle \ddot{\gamma}(t), \gamma(t) \rangle + \int_{t_1=0}^t \langle \ddot{\gamma}(t), \dot{\gamma}(t_1) \rangle dt_1 \\ &= 1 + \int_{t_1=0}^t \left\langle \ddot{\gamma}(t), \dot{\gamma}(t) - \int_{t_2=t_1}^t \ddot{\gamma}(t_2) dt_2 \right\rangle dt_1 \\ &= 1 - \int_{t_1=0}^t \int_{t_2=t_1}^t \langle \ddot{\gamma}(t), \ddot{\gamma}(t_2) \rangle dt_2 dt_1 \\ &\leq 1 + \kappa^2 \int_{t_1=0}^t \int_{t_2=t_1}^t dt_2 dt_1 \\ &\leq 1 + \frac{1}{2}\kappa^2 T^2 \\ &\leq 1 + 2\kappa^2 \Delta^2, \end{aligned} \quad (52)$$

where we used $\langle \ddot{\gamma}(t), \gamma(t) \rangle = -1$ (by differentiating both sides of $\langle \dot{\gamma}(t), \gamma(t) \rangle = 0$) and $\langle \ddot{\gamma}(t), \dot{\gamma}(t) \rangle = 0$.

Hence

$$\frac{d^2}{dt^2}(f \circ \gamma)(t) \leq 1 + 2\kappa^2 \Delta^2 + \kappa \Delta + \kappa \|\mathbf{z}\|. \quad (53)$$

Since $\|\mathbf{z}\|_2 \leq 2\sigma\sqrt{D}$ with probability at least $1 - e^{-D/2}$, combining this with $\Delta = 1/\kappa$ and $\sigma \leq \frac{1}{60\kappa\sqrt{D}}$, we get with high probability $\frac{d^2}{dt^2}(f \circ \gamma)(t) \leq \frac{121}{30}$. \square

D CHAINING BOUNDS FOR THE TANGENT BUNDLE PROCESS

In this section, we prove the following lemma, which bounds a crucial Gaussian process that arises in the analysis of gradient descent.

Main Bound for Tangent Bundle Process

Theorem 6. *Suppose that $\Delta \leq 1/\kappa$, and set*

$$T^{\max} = \sup \left\{ \langle \mathbf{v}, \mathbf{z} \rangle \mid d_S(\mathbf{s}, \mathbf{s}_{\natural}) \leq \Delta, \mathbf{v} \in T_{\mathbf{s}}S, \|\mathbf{v}\|_2 = 1 \right\}. \quad (54)$$

Then with probability at least $1 - 1.6e^{-\frac{x^2}{2\sigma^2}}$, we have

$$T^{\max} \leq 12\sigma(\kappa\sqrt{2\pi(d+1)} + \sqrt{\log 12\kappa}) + 30x. \quad (55)$$

We prove Theorem 6 below. We directly follow the proof of Theorem 5.29 from Van Handel (2016), establishing a chaining argument while accounting for slight discrepancies and establishing exact constants. The main geometric content of this argument is in Lemma 9, which bounds the size of ε -nets for the tangent bundle.

Proof. Set

$$\mathcal{V} = \left\{ \mathbf{v} \mid \mathbf{v} \in T_{\mathbf{s}}S, \|\mathbf{v}\|_2 = 1, d_S(\mathbf{s}, \mathbf{s}_{\natural}) \leq \Delta \right\}, \quad (56)$$

We first prove that $\mathcal{T} = \{\langle \mathbf{v}, \mathbf{z} \rangle\}_{\mathbf{v} \in \mathcal{V}}$ defines a separable, sub-gaussian process. Take any $\mathbf{v}, \mathbf{v}' \in \mathcal{V}$.

Then

$$\langle \mathbf{v}, \mathbf{z} \rangle - \langle \mathbf{v}', \mathbf{z} \rangle = \langle \mathbf{v} - \mathbf{v}', \mathbf{z} \rangle \sim \mathcal{N}(0, \sigma^2 d(\mathbf{v}, \mathbf{v}')^2), \quad (57)$$

immediately satisfying sub-gaussianity. By Lemma 9, there exists an ε -net $\mathcal{N}(\mathcal{V}, d, \varepsilon)$ for \mathcal{V} of size at most $N = (12\kappa/\varepsilon)^{2d+1}$. To see separability, let $\mathcal{N}_k = \mathcal{N}(\mathcal{V}, d, 2^{-k})$ be the epsilon net corresponding to $\varepsilon = \frac{1}{2^k}$. We can construct a countable dense subset of \mathcal{V} by letting

$$\mathcal{N}_{\infty} = \bigcup_{k=1}^{\infty} \mathcal{N}(\mathcal{V}, d, 2^{-k}). \quad (58)$$

Therefore, the existence of a countable dense subset implies separability of \mathcal{V} immediately implying separability of \mathcal{T} . Using these facts, we first prove the result in the finite case $|\mathcal{V}| < \infty$, after which we use separability to extend to the infinite case.

Let $|\mathcal{V}| < \infty$ and k_0 be the largest integer such that $2^{-k_0} \geq \text{diam}(\mathcal{V})$. Define $\mathcal{N}_{k_0} = \mathcal{N}(\mathcal{V}, d, 2^{-k_0})$ to be a 2^{-k_0} net of \mathcal{V} with respect to the metric d . Then for all $\mathbf{v} \in \mathcal{V}$, there exists $\pi_{k_0}(\mathbf{v}) \in \mathcal{N}_{k_0}$ such that $d(\mathbf{v}, \pi_{k_0}(\mathbf{v})) < 2^{-k_0}$.

For $k > k_0$, let $\mathcal{N}_k = \mathcal{N}(\mathcal{V}, d, 2^{-k})$ be a 2^{-k} net of \mathcal{V} . Subsequently for all $\mathbf{v} \in \mathcal{V}$, there exists $\pi_k(\mathbf{v}) \in \mathcal{N}_k$ such that $d(\mathbf{v}, \pi_k(\mathbf{v})) < 2^{-k}$.

Now fix any $\mathbf{v}_0 \in \mathcal{V}$. For any $\mathbf{v} \in \mathcal{V}$, sufficiently large n yields $\pi_n(\mathbf{v}) = \mathbf{v}$. Thus,

$$\langle \mathbf{v}, \mathbf{z} \rangle - \langle \mathbf{v}_0, \mathbf{z} \rangle = \sum_{k>k_0} \{ \langle \pi_k(\mathbf{v}), \mathbf{z} \rangle - \langle \pi_{k-1}(\mathbf{v}), \mathbf{z} \rangle \} \quad (59)$$

by the telescoping property, implying

$$\sup_{\mathbf{v} \in \mathcal{T}} \{ \langle \mathbf{v}, \mathbf{z} \rangle - \langle \mathbf{v}_0, \mathbf{z} \rangle \} \leq \sum_{k>k_0} \sup_{\mathbf{v} \in \mathcal{V}} \{ \langle \pi_k(\mathbf{v}), \mathbf{z} \rangle - \langle \pi_{k-1}(\mathbf{v}), \mathbf{z} \rangle \}. \quad (60)$$

Using the fact that \mathcal{T} is a sub-gaussian process and Lemma 5.2 of Van Handel (2016), we can bound each individual sum as

$$\mathbb{P}(\sup_{\mathbf{v} \in \mathcal{V}} \{\langle \pi_k(\mathbf{v}), \mathbf{z} \rangle - \langle \pi_{k-1}(\mathbf{v}), \mathbf{z} \rangle\} \geq 6 \times 2^{-k} \sigma \sqrt{\log |\mathcal{N}_k|} + 3 \times 2^{-k} x_k) \leq e^{-\frac{x_k^2}{2\sigma^2}}. \quad (61)$$

By ensuring that all of the sums are simultaneously controlled, we can arrive at the desired bound. We first derive the complement (i.e. there exists one sum which exceeds the desired value)

$$\mathbb{P}(A^c) := \mathbb{P}(\exists k > k_0 \text{ s.t. } \sup_{\mathbf{v} \in \mathcal{V}} \{\langle \pi_k(\mathbf{v}), \mathbf{z} \rangle - \langle \pi_{k-1}(\mathbf{v}), \mathbf{z} \rangle\} \geq 6 \times 2^{-k} \sigma \sqrt{\log |\mathcal{N}_k|} + 3 \times 2^{-k} x_k) \quad (62)$$

$$\leq \sum_{k > k_0} \mathbb{P}(\sup_{\mathbf{v} \in \mathcal{V}} \{\langle \pi_k(\mathbf{v}), \mathbf{z} \rangle - \langle \pi_{k-1}(\mathbf{v}), \mathbf{z} \rangle\} \geq 6 \times 2^{-k} \sigma \sqrt{\log |\mathcal{N}_k|} + 3 \times 2^{-k} x_k) \quad (63)$$

$$\leq \sum_{k > k_0} e^{-\frac{x_k^2}{2\sigma^2}} \quad (64)$$

$$\leq e^{-\frac{x^2}{2\sigma^2}} \sum_{k > 0} e^{-k/2} \leq 1.6e^{-\frac{x^2}{2\sigma^2}} \quad (65)$$

Now, using corollary 5.25 of Van Handel (2016) and $|\mathcal{N}| \leq (\frac{12\kappa}{\epsilon})^{2d+1}$, we have

$$\sup_{\mathbf{v} \in \mathcal{V}} \{\langle \mathbf{v}, \mathbf{z} \rangle - \langle \mathbf{v}_0, \mathbf{z} \rangle\} \leq \sum_{k > k_0} \sup_{\mathbf{v} \in \mathcal{V}} \{\langle \pi_k(\mathbf{v}), \mathbf{z} \rangle - \langle \pi_{k-1}(\mathbf{v}), \mathbf{z} \rangle\} \quad (66)$$

$$\leq 6 \sum_{k > k_0} 2^{-k} \sigma \sqrt{\log |\mathcal{N}_k|} + 3 \times 2^{-k_0} \sum_{k > 0} 2^{-k} \sqrt{k} + 3 \times 2^{-k_0} \sum_{k > 0} 2^{-k} x \quad (67)$$

$$\leq 12 \int_0^\infty \sigma \sqrt{\log \mathcal{N}(\mathcal{V}, d, \epsilon)} d\epsilon + 15 \text{diam}(\mathcal{V}) x \quad (68)$$

$$\leq 12\sigma \sqrt{2d+1} \int_0^\infty \sqrt{\log(\frac{12\kappa}{\epsilon})} d\epsilon + 15 \text{diam}(\mathcal{V}) x \quad (69)$$

$$= 12\sigma \sqrt{2d+1} (\kappa \sqrt{\pi} \text{erf}(\log 12\kappa) + \sqrt{\log 12\kappa}) + 15 \text{diam}(\mathcal{V}) x \quad (70)$$

$$\leq 12\sigma (\kappa \sqrt{2\pi(d+1)} + \sqrt{\log 12\kappa}) + 15 \text{diam}(\mathcal{V}) x, \quad (71)$$

where we have used $2^{-k_0} \leq 2 \text{diam}(\mathcal{V})$, $\sum_{k > 0} 2^{-k} \sqrt{k} \leq 1.35$ and $\sum_{k > 0} 2^{-k} \leq 1$ in (67), and $\text{erf} z = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2/2} dt \leq 1$ in (71).

Thus, if A occurs the above equation holds, implying

$$\mathbb{P}[\sup_{\mathbf{v} \in \mathcal{V}} \{\langle \mathbf{v}, \mathbf{z} \rangle - \langle \mathbf{v}_0, \mathbf{z} \rangle\} \geq 12\sigma (\kappa \sqrt{2\pi(d+1)} + \sqrt{\log 12\kappa}) + 15 \text{diam}(\mathcal{V}) x] \leq \mathbb{P}(A^c) \leq 1.6e^{-\frac{x^2}{2\sigma^2}} \quad (72)$$

Since \mathcal{T} is a separable process, Theorem 5.24 of Van Handel (2016) directly extends the result to infinite/uncountable \mathcal{T} . Letting $\langle \mathbf{v}_0, \mathbf{z} \rangle = 0$ and noting $\text{diam}(\mathcal{V}) = \sup_{\mathbf{v}, \mathbf{v}' \in \mathcal{V}} \|\mathbf{v} - \mathbf{v}'\|_2 \leq 2$ yields the claim. \square

NETS FOR $B(\mathbf{s}_q, \Delta)$

Lemma 7. *Suppose that $\Delta < 1/\kappa$. For any $\varepsilon \in (0, \dots]$, there exists an ε -net \widehat{S} for $B(\mathbf{s}_q, \Delta)$ of size $\#\widehat{S} < (12/\varepsilon)^{d+1}$.*

At a high level, the proof of this lemma proceeds as follows: we form an ε_0 net N_0 for $T_{\mathbf{s}_q} S$, and then set $\widehat{S} = \{\exp_{\mathbf{s}_q}(\mathbf{v}) \mid \mathbf{v} \in N_0\}$. We will argue that \widehat{S} is a $C\varepsilon_0$ -net for $B(\mathbf{s}_q, \Delta)$, by arguing that

at length scales $\Delta < 1/\kappa$, the distortion induced by the exponential map is bounded. Crucial to this argument is the following lemma on geodesic triangles:

Lemma 8. *Consider $\mathbf{v}, \mathbf{v}' \in T_{\mathbf{s}_i} S$, with $\|\mathbf{v}\|_2 = \|\mathbf{v}'\|_2 < \Delta$. Then if $\angle(\mathbf{v}, \mathbf{v}') < \frac{1}{\sqrt{3}}$,*

$$d_S\left(\exp_{\mathbf{s}_i}(\mathbf{v}), \exp_{\mathbf{s}_i}(\mathbf{v}')\right) \leq \sqrt{6} \Delta \angle(\mathbf{v}, \mathbf{v}'). \quad (73)$$

This lemma says that the third side of the triangle with vertices $\mathbf{s}_i, \exp_{\mathbf{s}_i}(\mathbf{v}), \exp_{\mathbf{s}_i}(\mathbf{v}')$ is at most a constant longer than the third side of an analogous triangle in Euclidean space. The proof of this is a direct application of Toponogov’s theorem, a fundamental result in Riemannian geometry which allows one to compare triangles in an arbitrary Riemannian manifold whose sectional curvature is lower bounded to triangles in a constant curvature model space, where one can apply concrete trigonometric reasoning.

Proof of Lemma 8. By Lemma 10, the sectional curvatures κ_s of S are uniformly bounded in terms of the extrinsic geodesic curvature κ :

$$\kappa_s \geq -\kappa^2. \quad (74)$$

By Toponogov’s theorem Toponogov (2006), the length $d_S\left(\exp_{\mathbf{s}_i}(\mathbf{v}), \exp_{\mathbf{s}_i}(\mathbf{v}')\right)$, of the third side of the geodesic triangle $\mathbf{s}_i, \exp_{\mathbf{s}_i}(\mathbf{v}), \exp_{\mathbf{s}_i}(\mathbf{v}')$ is bounded by the length of the third side of a geodesic triangles with two sides of length $r = \|\mathbf{v}\| = \|\mathbf{v}'\|$ and angle $\theta = \angle(\mathbf{v}, \mathbf{v}')$ in the constant curvature model space $M_{-\kappa^2}$. We can rescale, so that this third length is bounded by L/κ , where L is the length of the third side of a geodesic triangle with two sides of length $r\kappa$ and an angle of $\angle(\mathbf{v}, \mathbf{v}')$, in the hyperbolic space M_{-1} . Using hyperbolic trigonometry (cf Fact 11 and the identity $\cosh^2 t - \sinh^2 t = 1$), we have

$$\cosh L = 1 + \sinh^2(\kappa r) \times (1 - \cos \theta). \quad (75)$$

By convexity of \sinh over $[0, \infty)$, for $t \in [0, 1]$, we have $\sinh(t) \leq t \sinh(1)$, and $\sinh^2(t) \leq t^2 \sinh^2(1) < \frac{3}{2}t^2$; since $\kappa r < 1$, $\sinh^2(\kappa r) < \frac{3}{2}\kappa^2 r^2$. Since $\cos(t) \geq 1 - t^2$ for all t , we have

$$\cosh L \leq 1 + \frac{3}{2}\kappa^2 r^2 \theta^2. \quad (76)$$

Using $\kappa r < 1$, for $\theta < \frac{1}{\sqrt{3}}$ we have $\cosh(L) \leq \frac{3}{2} < \cosh(1)$. Noting that for $t \in [0, 1]$,

$$\cosh(t) \geq g(t) = 1 + \frac{1}{4}t^2, \quad (77)$$

on $s \in [0, \cosh(1)]$, we have $\cosh^{-1}(s) \leq g^{-1}(s) = 2\sqrt{s-1}$, giving

$$L \leq \sqrt{6} \cdot \kappa r \theta. \quad (78)$$

Dividing by κ gives the claimed bound. \square

Proof of Lemma 7. Form an (angular) ε_0 -net N_0 for $\{\mathbf{v} \in T_{\mathbf{s}_i} S \mid \|\mathbf{v}\|_2 = 1\}$ satisfying

$$\forall \mathbf{v} \in T_{\mathbf{s}_i} S, \exists \hat{\mathbf{v}} \in N_0 \text{ with } \angle(\mathbf{v}, \hat{\mathbf{v}}) \leq \varepsilon, \quad (79)$$

and an ε_0 -net

$$N_r = \{0, \varepsilon_0, 2\varepsilon_0, \dots, \lfloor \Delta/\varepsilon_0 \rfloor\} \quad (80)$$

for the interval $[0, \Delta]$. We can take $\#N_0 \leq (3/\varepsilon_0)^d$ and $\#N_r \leq \Delta/\varepsilon_0 \leq 1/\varepsilon_0$. Combine these two to form a net N for $\{\mathbf{v} \in T_{\mathbf{s}_0} S \mid \|\mathbf{v}\|_2 \leq \Delta\}$ by setting

$$N = \bigcup_{r \in N_r} rN_0. \quad (81)$$

Note that $\#N \leq (3/\varepsilon_0)^{d+1}$. Let $\hat{S} = \{\exp_{\mathbf{s}_i}(\mathbf{v}) \mid \mathbf{v} \in N\}$. Consider an arbitrary element \mathbf{s} of $B(\mathbf{s}_i, \Delta)$. There exists $\mathbf{v} \in T_{\mathbf{s}_i} S$ such that $\exp_{\mathbf{s}_i}(\mathbf{v}) = \mathbf{s}$. Set

$$\bar{\mathbf{v}} = \varepsilon_0 \left[\frac{\|\mathbf{v}\|_2}{\varepsilon_0} \right] \mathbf{v}. \quad (82)$$

There exists $\widehat{\mathbf{v}} \in N$ with $\|\widehat{\mathbf{v}}\|_2 = \|\bar{\mathbf{v}}\|_2$ and $\angle(\widehat{\mathbf{v}}, \bar{\mathbf{v}}) \leq \varepsilon_0$. Note that

$$\widehat{\mathbf{s}} = \exp_{\mathbf{s}_\natural}(\widehat{\mathbf{v}}) \in \widehat{S}. \quad (83)$$

By Lemma 8, we have

$$\begin{aligned} d_S(\mathbf{s}, \widehat{\mathbf{s}}) &\leq d_S(\mathbf{s}, \exp_{\mathbf{s}_\natural}(\bar{\mathbf{v}})) + d_S(\exp_{\mathbf{s}_\natural}(\bar{\mathbf{v}}), \widehat{\mathbf{s}}) \\ &\leq \varepsilon_0 + 3\Delta\varepsilon_0 \\ &< 4\varepsilon_0. \end{aligned} \quad (84)$$

Setting $\varepsilon_0 = \varepsilon/4$, we obtain that \widehat{S} is an ε -net for $B(\mathbf{s}_\natural, \Delta)$. \square

NETS FOR THE TANGENT BUNDLE

Lemma 9. *Set*

$$T = \left\{ \mathbf{v} \mid \mathbf{v} \in T_{\mathbf{s}}S, \|\mathbf{v}\|_2 = 1, d_S(\mathbf{s}, \mathbf{s}_\natural) \leq \Delta \right\}, \quad (85)$$

Then there exists an ε -net \widehat{T} for T of size

$$\#\widehat{T} \leq \left(\frac{12\kappa}{\varepsilon} \right)^{2d+1}. \quad (86)$$

Proof. Let \widehat{S} be the ε_0 -net for $B(\mathbf{s}_\natural, \Delta)$. By Lemma 7, there exists such a net of size at most $(12/\varepsilon_0)^{d+1}$. For each $\widehat{\mathbf{s}} \in \widehat{S}$, form an ε_1 -net $N_{\widehat{\mathbf{s}}}$ for

$$\left\{ \mathbf{v} \in T_{\widehat{\mathbf{s}}}S \mid \|\mathbf{v}\|_2 = 1 \right\}. \quad (87)$$

We set

$$\widehat{T} = \bigcup_{\widehat{\mathbf{s}} \in \widehat{S}} N_{\widehat{\mathbf{s}}}. \quad (88)$$

By ? Lemma 5.2, we can take $\#N_{\widehat{\mathbf{s}}} \leq (3/\varepsilon_1)^d$, and so

$$\#\widehat{T} \leq \left(\frac{3}{\varepsilon_1} \right)^d \left(\frac{12}{\varepsilon_0} \right)^{d+1}. \quad (89)$$

Consider an arbitrary element $\mathbf{v} \in T$. The vector \mathbf{v} belongs to the tangent space $T_{\mathbf{s}}S$ for some \mathbf{s} . By construction, there exists $\widehat{\mathbf{s}} \in \widehat{S}$ with $d_S(\mathbf{s}, \widehat{\mathbf{s}}) \leq \varepsilon$. Consider a minimal geodesic γ joining \mathbf{s} and $\widehat{\mathbf{s}}$. We generate $\bar{\mathbf{v}} \in T_{\widehat{\mathbf{s}}}S$ by parallel transporting \mathbf{v} along γ . Let $\mathcal{P}_{t,0}$ denote this parallel transport. By ? Lemma 8.5, the vector field $\mathbf{v}_t = \mathcal{P}_{t,0}\mathbf{v}$ satisfies

$$\frac{d}{dt}\mathbf{v}_t = \mathbb{I}(\dot{\gamma}(s), \mathbf{v}_t), \quad (90)$$

where $\mathbb{I}(\cdot, \cdot)$ is the second fundamental form. So,

$$\mathcal{P}_{t,0}\mathbf{v} = \mathbf{v} + \int_0^t \mathbb{I}(\dot{\gamma}(s), \mathbf{v}_s) ds. \quad (91)$$

By Lemma 4, for every s

$$\left\| \mathbb{I}(\dot{\gamma}(s), \mathbf{v}_s) \right\| \leq 3\kappa \quad (92)$$

and

$$\|\bar{\mathbf{v}} - \mathbf{v}\| \leq 3\varepsilon_0\kappa. \quad (93)$$

By construction, there is an element $\widehat{\mathbf{v}}$ of $N_{\widehat{\mathbf{s}}}$ with

$$\|\widehat{\mathbf{v}} - \bar{\mathbf{v}}\| \leq \varepsilon_1, \quad (94)$$

and so \widehat{T} is an $\varepsilon_1 + 3\kappa\varepsilon_0$ -net for T . Setting $\varepsilon_1 = \varepsilon/4$ and $\varepsilon_0 = \varepsilon/4\kappa$ completes the proof. \square

SUPPORTING RESULTS ON GEOMETRY

Lemma 10. For a Riemannian submanifold S of \mathbb{R}^D , the sectional curvatures $\kappa_s(\mathbf{v}, \mathbf{v}')$ are bounded by the extrinsic geodesic curvature κ , as

$$\kappa_s(\mathbf{v}, \mathbf{v}') \geq -\kappa^2. \quad (95)$$

Proof. Using the Gauss formula (Theorem 8.4 of ?), the Riemann curvature tensor R_S of S is related to the Riemann curvature tensor $R_{\mathbb{R}^D}$ of the ambient space via

$$\begin{aligned} \langle R_S(\mathbf{u}, \mathbf{v})\mathbf{v}, \mathbf{u} \rangle &= \langle R_{\mathbb{R}^D}(\mathbf{u}, \mathbf{v})\mathbf{v}, \mathbf{u} \rangle + \langle \mathbb{I}(\mathbf{u}, \mathbf{v}), \mathbb{I}(\mathbf{u}, \mathbf{v}) \rangle - \langle \mathbb{I}(\mathbf{u}, \mathbf{u}), \mathbb{I}(\mathbf{v}, \mathbf{v}) \rangle \\ &= \langle \mathbb{I}(\mathbf{u}, \mathbf{v}), \mathbb{I}(\mathbf{u}, \mathbf{v}) \rangle - \langle \mathbb{I}(\mathbf{u}, \mathbf{u}), \mathbb{I}(\mathbf{v}, \mathbf{v}) \rangle \\ &\geq -\langle \mathbb{I}(\mathbf{u}, \mathbf{u}), \mathbb{I}(\mathbf{v}, \mathbf{v}) \rangle, \end{aligned} \quad (96)$$

where we have used that $R_{\mathbb{R}^D} = 0$ and $\langle \mathbb{I}(\mathbf{u}, \mathbf{v}), \mathbb{I}(\mathbf{u}, \mathbf{v}) \rangle \geq 0$. Take any $\mathbf{v}, \mathbf{v}' \in T_s S$. The sectional curvature $\kappa_s(\mathbf{v}, \mathbf{v}')$ satisfies

$$\kappa_s(\mathbf{v}, \mathbf{v}') = \kappa_s(\mathbf{u}, \mathbf{u}') = \langle R_S(\mathbf{u}, \mathbf{u}')\mathbf{u}', \mathbf{u} \rangle, \quad (97)$$

for any orthonormal basis \mathbf{u}, \mathbf{u}' for $\text{span}(\mathbf{v}, \mathbf{v}')$. So

$$\kappa_s(\mathbf{v}, \mathbf{v}') = \langle R_S(\mathbf{u}, \mathbf{u}')\mathbf{u}', \mathbf{u} \rangle \geq -\langle \mathbb{I}(\mathbf{u}, \mathbf{u}), \mathbb{I}(\mathbf{u}', \mathbf{u}') \rangle \geq -\kappa^2, \quad (98)$$

as claimed. \square

Fact 11. For a hyperbolic triangle with side lengths a, b, c and corresponding (opposite) angles A, B, C , we have

$$\cosh c = \cosh a \cosh b - \sinh a \sinh b \cos C. \quad (99)$$

E PROOF OF RESULT (9)

In this section, we state and prove the other part of our main claims about gradient descent:

Theorem 12. Suppose that $\mathbf{x} = \mathbf{s}_{\natural} + \mathbf{z}$, with $T^{\max}(\mathbf{z}) < 1/\kappa$. Consider the constant-stepping Riemannian gradient method, with initial point \mathbf{s}^0 satisfying $d(\mathbf{s}^0, \mathbf{s}_{\natural}) < 1/\kappa$, and step size $\tau = \frac{1}{64}$.

$$d(\mathbf{s}^{k+1}, \mathbf{s}_{\natural}) \leq (1 - \varepsilon) \cdot d(\mathbf{s}^k, \mathbf{s}_{\natural}) + CT^{\max}. \quad (100)$$

Here, C and ε are positive numerical constants.

Together with Theorem 6, this result shows that gradient descent rapidly converges to a neighborhood of the truth of radius $C\sigma\sqrt{d}$.

Proof. Let

$$\bar{\mathbf{s}}_t = \exp\left(-t \cdot \text{grad}[f](\mathbf{s}^k)\right) \quad (101)$$

be a geodesic joining \mathbf{s}^k and \mathbf{s}^{k+1} , with $\bar{\mathbf{s}}_0 = \mathbf{s}^k$ and $\bar{\mathbf{s}}_\tau = \mathbf{s}^{k+1}$. Let f_{\natural} denote a noise-free version of the objective function, i.e.,

$$f_{\natural}(\mathbf{s}) = -\langle \mathbf{s}, \mathbf{s}_{\natural} \rangle, \quad (102)$$

and notice that for all \mathbf{s} ,

$$\text{grad}[f_{\natural}](\mathbf{s}) = \text{grad}[f](\mathbf{s}) + P_{T_s S} \mathbf{z}. \quad (103)$$

Furthermore, following calculations in Lemma 5, on $B(\mathbf{s}_{\natural}, 1/\kappa)$, the Riemannian hessian of f_{\natural} is bounded as

$$\|\text{Hess}[f_{\natural}](\mathbf{s})\| \leq 4. \quad (104)$$

Using the relationship

$$\text{grad}[f_{\natural}](\bar{\mathbf{s}}_t) = \mathcal{P}_{\bar{\mathbf{s}}_t, \bar{\mathbf{s}}_0} \text{grad}[f_{\natural}](\bar{\mathbf{s}}_0) + \int_{r=0}^t \mathcal{P}_{\bar{\mathbf{s}}_t, \bar{\mathbf{s}}_r} \text{Hess}[f_{\natural}](\bar{\mathbf{s}}_r) \mathcal{P}_{\bar{\mathbf{s}}_r, \bar{\mathbf{s}}_0} \text{grad}[f](\bar{\mathbf{s}}_0) dr, \quad (105)$$

where $\mathcal{P}_{\bar{\mathbf{s}}_t, \bar{\mathbf{s}}_0}$ to denote parallel transport along the curve $\bar{\mathbf{s}}_t$, we obtain that

$$\left\| \text{grad}[f_{\natural}](\bar{\mathbf{s}}_t) - \mathcal{P}_{\bar{\mathbf{s}}_t, \bar{\mathbf{s}}_0} \text{grad}[f_{\natural}](\bar{\mathbf{s}}_0) \right\| \leq 4t \|\text{grad}[f](\bar{\mathbf{s}}_0)\|_2. \quad (106)$$

Along the curve $\bar{\mathbf{s}}_t$, the distance to \mathbf{s}_{\natural} evolves as

$$\begin{aligned} \frac{d}{dt} d(\bar{\mathbf{s}}_t, \mathbf{s}_{\natural}) &= - \left\langle \mathcal{P}_{t,0} \text{grad}[f](\bar{\mathbf{s}}_0), \frac{-\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}}{\|\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}\|_2} \right\rangle \\ &= \left\langle -\text{grad}[f](\bar{\mathbf{s}}_t), \frac{-\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}}{\|\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}\|_2} \right\rangle + \left\langle \text{grad}[f](\bar{\mathbf{s}}_t) - \mathcal{P}_{t,0} \text{grad}[f](\bar{\mathbf{s}}_0), \frac{-\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}}{\|\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}\|_2} \right\rangle \\ &\leq \left\langle -\text{grad}[f](\bar{\mathbf{s}}_t), \frac{-\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}}{\|\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}\|_2} \right\rangle + \left\langle \text{grad}[f_{\natural}](\bar{\mathbf{s}}_t) - \mathcal{P}_{t,0} \text{grad}[f_{\natural}](\bar{\mathbf{s}}_0), \frac{-\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}}{\|\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}\|_2} \right\rangle + 2T^{\max} \\ &\leq \left\langle -\text{grad}[f](\bar{\mathbf{s}}_t), \frac{-\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}}{\|\log_{\bar{\mathbf{s}}_t} \mathbf{s}_{\natural}\|_2} \right\rangle + 4t \|\text{grad}[f](\bar{\mathbf{s}}_0)\| + 2T^{\max} \\ &\leq -\frac{1}{2} d(\bar{\mathbf{s}}_t, \mathbf{s}_{\natural}) + 4td(\bar{\mathbf{s}}_0, \mathbf{s}_{\natural}) + (3+4t)T^{\max} \\ &\leq -\frac{1}{2} d(\bar{\mathbf{s}}_t, \mathbf{s}_{\natural}) + \frac{1}{16} d(\bar{\mathbf{s}}_0, \mathbf{s}_{\natural}) + 4T^{\max} \end{aligned} \quad (107)$$

where we have used Lemma 13. Setting $X_t = d(\bar{\mathbf{s}}_t, \mathbf{s}_{\natural})$, we have

$$\dot{X}_t \leq -\frac{1}{4} X_t \quad (108)$$

whenever $X_t \geq \frac{1}{4} X_0 + 16T^{\max}$. Hence,

$$X_{\tau} \leq \max \left\{ e^{-\frac{\tau}{4}} X_0, \frac{1}{4} X_0 + 16T^{\max} \right\}, \quad (109)$$

and so

$$d(\mathbf{s}^{k+1}, \mathbf{s}_{\natural}) \leq \exp\left(-\frac{1}{256}\right) \cdot d(\mathbf{s}^k, \mathbf{s}_{\natural}) + 16T^{\max}, \quad (110)$$

as claimed. \square

E.1 SUPPORTING LEMMAS

Lemma 13. *Suppose that $\Delta < 1/\kappa$. For all $\mathbf{s} \in B(\mathbf{s}_{\natural}, \Delta)$, we have*

$$\left\langle -\text{grad}[f](\mathbf{s}), \frac{-\log_{\mathbf{s}} \mathbf{s}_{\natural}}{\|\log_{\mathbf{s}} \mathbf{s}_{\natural}\|_2} \right\rangle \leq -\frac{1}{2} d(\mathbf{s}, \mathbf{s}_{\natural}) + T^{\max}. \quad (111)$$

Proof. Notice that

$$\begin{aligned} \left\langle -\text{grad}[f](\mathbf{s}), \frac{-\log_{\mathbf{s}} \mathbf{s}_{\natural}}{\|\log_{\mathbf{s}} \mathbf{s}_{\natural}\|_2} \right\rangle &= \left\langle P_{T_s S}(\mathbf{s}_{\natural} + \mathbf{z}), \frac{-\log_{\mathbf{s}} \mathbf{s}_{\natural}}{\|\log_{\mathbf{s}} \mathbf{s}_{\natural}\|_2} \right\rangle \\ &\leq \left\langle P_{T_s S} \mathbf{s}_{\natural}, \frac{-\log_{\mathbf{s}} \mathbf{s}_{\natural}}{\|\log_{\mathbf{s}} \mathbf{s}_{\natural}\|_2} \right\rangle + T^{\max}. \end{aligned} \quad (112)$$

Consider a unit speed geodesic γ joining \mathbf{s}_{\natural} and \mathbf{s} , with $\gamma(0) = \mathbf{s}_{\natural}$ and $\gamma(t) = \mathbf{s}$. Then

$$\frac{-\log_{\mathbf{s}} \mathbf{s}_{\natural}}{\|\log_{\mathbf{s}} \mathbf{s}_{\natural}\|_2} = \dot{\gamma}(t), \quad (113)$$

and

$$\begin{aligned}
\left\langle P_{T_s S} \mathbf{s}_{\natural}, \frac{-\log_{\mathbf{s}} \mathbf{s}_{\natural}}{\|\log_{\mathbf{s}} \mathbf{s}_{\natural}\|_2} \right\rangle &= \langle \gamma(0), \dot{\gamma}(t) \rangle \\
&= \underbrace{\langle \gamma(t), \dot{\gamma}(t) \rangle}_{\text{this term} = 0} - \int_0^t \langle \dot{\gamma}(s), \dot{\gamma}(t) \rangle ds \\
&= -t \|\dot{\gamma}(t)\|_2^2 - \int_0^t \int_t^s \langle \ddot{\gamma}(r), \dot{\gamma}(t) \rangle dr ds \\
&\leq -d(\mathbf{s}, \mathbf{s}_{\natural}) + \frac{1}{2} \kappa d^2(\mathbf{s}, \mathbf{s}_{\natural}).
\end{aligned} \tag{114}$$

In particular, this term is bounded by $-\frac{1}{2}d(\mathbf{s}, \mathbf{s}_{\natural})$ when $\Delta < 1/\kappa$. \square

Lemma 14. For $\mathbf{s} \in B(\mathbf{s}_{\natural}, \Delta)$, we have

$$\left\| \text{grad}[f](\mathbf{s}) \right\| \leq d(\mathbf{s}, \mathbf{s}_{\natural}) + T^{\max} \tag{115}$$

Proof. Notice that

$$\begin{aligned}
\left\| \text{grad}[f](\mathbf{s}) \right\| &= \left\| P_{T_s S}(\mathbf{s}_{\natural} + \mathbf{z}) \right\| \\
&\leq \left\| P_{T_s S} \mathbf{s}_{\natural} \right\| + T^{\max} \\
&\leq \left\| P_{T_s \mathbb{S}^{D-1}} \mathbf{s}_{\natural} \right\| + T^{\max} \\
&= \sin \angle(\mathbf{s}, \mathbf{s}_{\natural}) + T^{\max} \\
&\leq d_{\mathbb{S}^{D-1}}(\mathbf{s}, \mathbf{s}_{\natural}) + T^{\max}, \\
&\leq d_S(\mathbf{s}, \mathbf{s}_{\natural}) + T^{\max},
\end{aligned} \tag{116}$$

as claimed. \square

F ADDITIONAL EXPERIMENTAL DETAILS

F.1 GRAVITATIONAL WAVE GENERATION

Below we introduce some details on Gravitational Wave data generation. Synthetic gravitational waveforms are generated with the PyCBC package Nitz et al. (2023) with masses uniformly drawn from $[20, 50]$ (times solar mass M_{\odot}) and 3-dimensional spins drawn from a uniform distribution over the unit ball, at sampling rate 2048Hz. Each waveform is padded or truncated to 1 second long such that the peak is aligned at the 0.9 second location, and then normalized to have unit ℓ^2 norm. Noise is simulated as iid Gaussian with standard deviation $\sigma = 0.1$. The signal amplitude is constant $a = 1$. The training set contains 100,000 noisy waveforms, the test set contains 10,000 noisy waveforms and pure noise each, and a separate validation set constructed iid as the test set is used to select optimal template banks for MF.

F.2 HANDWRITTEN DIGIT RECOGNITION EXPERIMENT SETUP

The MNIST training set contains 6,131 images of the digit 3. In particular, we create a training set containing 10,000 images of randomly transformed digit 3 from the MNIST training set, and a test set containing 10,000 images each of randomly transformed digit 3 and other digits from the MNIST test set. We select a random subset of 1,000 embedded points as the quantization $\hat{\Xi}$ of the parameter space, and construct a k -d tree from it to perform efficient nearest neighbor search for kernel interpolation. Parameters of the trainable TpopT are initialized using heuristics based on the Jacobians, step sizes and smoothing levels from the unrolled optimization, similar to the previous experiment. ξ^0 is initialized at the center of the embedding space. We use the Adam optimizer with batch size 100 and constant learning rate 10^{-3} .