
LEARNING GLOBALLY SMOOTH FUNCTIONS ON MANIFOLDS

Anonymous authors

Paper under double-blind review

ABSTRACT

Smoothness and low dimensional structures play central roles in improving generalization and stability in learning and statistics. The combination of these properties has led to many advances in semi-supervised learning, generative modeling, and control of dynamical systems. However, learning smooth functions is generally challenging, except in simple cases such as learning linear or kernel models. Typical methods are either too conservative, relying on crude upper bounds such as spectral normalization, too lax, penalizing smoothness on average, or too computationally intensive, requiring the solution of large-scale semi-definite programs. These issues are only exacerbated when trying to simultaneously exploit low dimensionality using, e.g., manifolds. This work proposes to overcome these obstacles by combining techniques from semi-infinite constrained learning and manifold regularization. To do so, it shows that, under typical conditions, the problem of learning a Lipschitz continuous function on a manifold is equivalent to a dynamically weighted manifold regularization problem. This observation leads to a practical algorithm based on a weighted Laplacian penalty whose weights are adapted using stochastic gradient techniques. We prove that, under mild conditions, this method estimates the Lipschitz constant of the solution, learning a globally smooth solution as a byproduct. Numerical examples illustrate the advantages of using this method to impose global smoothness on manifolds as opposed to imposing smoothness on average.

1 INTRODUCTION

Learning smooth functions has been shown to be advantageous in general and is of particular interest in physical systems. This is because of the general observation that close input features tend to be associated with close outputs and of the particular fact that in physical systems Lipschitz continuity of input-output maps translates to stability and safety (Oberman and Calder, 2018; Finlay et al., 2018a; Couellan, 2021; Finlay et al., 2018b; Pauli et al., 2021; Krishnan et al., 2020; Shi et al., 2019; Lindemann et al., 2021; Arghal et al., 2021).

To learn smooth functions one can require the parameterization to be smooth. Such is the idea, e.g., of spectral normalization of weights in neural networks (Miyato et al., 2018; Zhao and Liu, 2020). Smooth parameterizations have the advantage of being globally smooth, but they may be restrictive because they impose smoothness for inputs that are not necessarily realized in the data. This drawback motivates the use of Lipschitz *penalties* in risk minimization (Oberman and Calder, 2018; Finlay et al., 2018a; Couellan, 2021; Pauli et al., 2021; Bungert et al., 2021), which offers the opposite tradeoff. Since penalties encourage but do not enforce small Lipschitz constants, we may learn functions that are smooth on average, but with no global guarantees of smoothness at every point in the support of the data. Formulations that guarantee *global* smoothness can be obtained if the risk minimization problem is modified by the addition of a Lipschitz constant *constraint* (Krishnan et al., 2020; Shi et al., 2019; Lindemann et al., 2021; Arghal et al., 2021). This yields formulations that guarantee Lipschitz smoothness in all possible inputs without the drawback of enforcing smoothness outside of the input data distribution. Several empirical studies (Krishnan et al., 2020; Shi et al., 2019; Lindemann et al., 2021; Arghal et al., 2021) have demonstrated the advantage of imposing global smoothness constraints only on observed inputs.

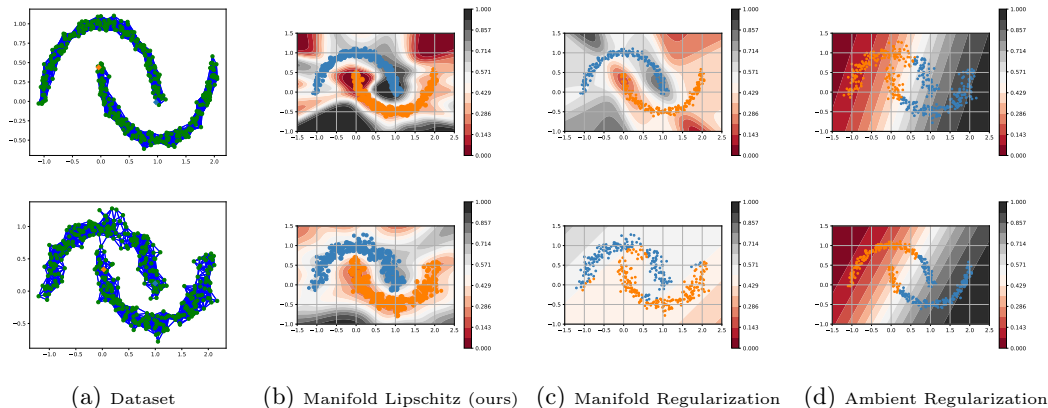


Figure 1: Two moons dataset. The setting consists of a two dimensional classification problem of two classes with 1 labeled, and 200 unlabeled samples per class. The objective is to correctly classify the 200 unlabeled samples. We consider two cases (top) the estimated manifold has two connected components, and (bottom) the manifold is weakly connected (cf. Figure 1). We plot the output of a one layer neural network trained using Manifold Regularization, Manifold/Ambient Lipschitz. Ambient Regularization fails to classify the unlabeled samples, given that ignores the distribution of samples given by the Manifold. The case in which the manifold has two connected component (cf. Figure 1a), our method works as good as Manifold Regularization, due to the fact that the Lipschitz constant will be made small in both components separately. However, when the manifold is weakly connected, Manifold Regularization fails to recognize the transition between the components, as it will penalize large gradients across the manifold, converging to a plane that connects the two samples. Our Manifold Lipschitz method, as it requires the Lipschitz constant to be small, forces a sharp transition along the point with maximal separation.

In this paper we exploit the fact that data can be often modeled as points in a low-dimensional manifold. We therefore consider manifold Lipschitz constants in which function smoothness is assessed with respect to distances measured over the data manifold (Definition 1). Although this looks like a minor difference, controlling Lipschitz constants over data manifolds is quite different from controlling Lipschitz constants in the ambient space. In Figure 1 we look at a classification problem with classes arranged in two separate half moons. Constraining Lipschitz constants in the ambient space effectively assumes the underlying data is uniformly distributed in space [cf. Figure 1-(d)]. Constraining Lipschitz constants in the data manifold, however, properly accounts for the data distribution [cf. Figure 1-(a)].

This example also illustrates how constraining manifold Lipschitz constants is related to manifold regularization (Belkin et al., 2005; Niyogi, 2013; Li et al., 2022). The difference is that manifold regularization penalizes the average norm of the manifold gradient. This distinction is significant because regularizing is more brittle than imposing constraints. In the example in Figure 1, manifold regularization fails to separate the dataset when the moons are close [cf. Figure 1-(c), bottom]. Classification with a manifold Lipschitz constant constraint is more robust to this change in the data distribution [cf. Figure 1-(a), bottom].

Global constraints in the manifold gradient yield a statistical constrained learning problem with an infinite and dense number of constraints. This is a challenging problem to approximate and solve. Here, we approach the solution of this problem in the Lagrangian dual domain and establish connections with manifold regularization that allow for the use of point cloud Laplacians. Our specific contributions are the following:

- (C1) We introduce a constrained statistical risk minimization problem in which we learn a function that: (i) attains a target loss and (ii) attains the smallest possible manifold Lipschitz constant among these functions that satisfy the target loss (Section 2).

-
- (C2) We introduce the Lagrangian dual problem and show that its empirical version is a statistically consistent approximation of the primal. This results does *not* require the learning parametrization to be convex (Section 3.1).
 - (C3) We generalize results from the manifold regularization literature to show that under regularity conditions, the evaluation of manifold Lipschitz constants can be recast in a more amenable form utilizing a weighted point cloud Laplacian (Proposition 3 in Section 3.2).
 - (C4) We present a dual ascent algorithm to find optimal multipliers. The function that attains the target loss and minimizes the manifold Lipschitz constant follows as a byproduct (Section 3.3).
 - (C5) We illustrate the merits of learning with global manifold smoothness guarantees with respect to ambient space and standard manifold regularization through two examples: (i) learning robot navigation policies in a space with obstacles and (ii) learning model mismatches in differential drive steering over non-ideal surfaces (Section 4).

RELATED WORK

This paper is at the intersection of learning with Lipschitz constant constraints (Oberman and Calder, 2018; Finlay et al., 2018a; Couellan, 2021; Pauli et al., 2021; Bungert et al., 2021; Miyato et al., 2018; Zhao and Liu, 2020; Krishnan et al., 2020; Shi et al., 2019; Lindemann et al., 2021; Arghal et al., 2021) and manifold regularization (Belkin et al., 2005; Niyogi, 2013; Li et al., 2022; Hein et al., 2005; Belkin and Niyogi, 2005). Relative to the literature on learning with Lipschitz constraints we offer the ability to leverage data manifolds. Since data manifolds are often characterized with unlabeled data (Kejani et al., 2020; Belkin and Niyogi, 2004; Jiang et al., 2019; Kipf and Welling, 2016; Yang et al., 2016; Zhu, 2005; Lecouat et al., 2018; Ouali et al., 2020; Cabannes et al., 2021), similar to these works, we utilize the point cloud Laplacian technique to compute the integral of the norm of the gradient. Relative to the literature on manifold regularization we offer global smoothness assurances instead of an average penalty of large manifold gradients. In (Krishnan et al., 2020), the space is discretized maintaining convexity, and at the expense of deriving a sensitive solution, their approach leads to a large optimization problem. Similar to us, (Krishnan et al., 2020) poses the problem of minimizing a Lipschitz constant, however, they utilize a softer surrogate (i.e. p -norm loss) which is a smoother version of the Lipschitz constant. Their approach therefore, tradeoffs numerical stability (small p) with accurate Lipschitz constant estimation ($p = \infty$). We do not work with surrogates, and we seek to minimize the maximum norm of the gradient, utilizing an epigraph technique.

2 GLOBAL CONSTRAINING OF MANIFOLD LIPSCHITZ CONSTANTS

We consider data pairs (x, y) in which the input features $x \in \mathcal{M} \subset \mathbb{R}^D$ lie in a compact oriented Riemannian manifold \mathcal{M} and the output features are real valued $y \in \mathbb{R}$. We study the regression problem of finding a function $f_\theta : \mathcal{M} \rightarrow \mathbb{R}$, parameterized by $\theta \in \Theta \subset \mathbb{R}^Q$ that minimizes the expectation of a nonnegative loss $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$, where $\ell(f_\theta(x), y)$ represents the loss of predicting output $f_\theta(x)$ when the world realizes the pair (x, y) . Data pairs (x, y) are drawn according to an unknown probability distribution $p(x, y)$ on $\mathcal{M} \times \mathbb{R}$ which we can factor as $p(x, y) = p(x)p(y|x)$.

We are interested in learning smooth functions, i.e. functions with *controlled* variability over the manifold \mathcal{M} . We therefore let $\nabla_{\mathcal{M}} f_\theta(x)$ represent the manifold gradient of f_θ and introduce the following definition.

Definition 1 (Manifold Lipschitz Constant). *Given a Riemannian manifold \mathcal{M} , the function $f_\theta : \mathcal{M} \rightarrow \mathbb{R}$ is said to be L -Lipschitz continuous if there exists a strictly positive constant $L > 0$ such that for all pairs of points $x_1, x_2 \in \mathcal{M}$,*

$$|f_\theta(x_1) - f_\theta(x_2)| \leq L d_{\mathcal{M}}(x_1, x_2), \quad (1)$$

where $d_{\mathcal{M}}(x_1, x_2)$ denotes the distance between x_1 and x_2 in the manifold \mathcal{M} . If the function f_{θ} is differentiable, (1) is equivalent to requiring the gradient norm to be bounded by L ,

$$\|\nabla_{\mathcal{M}} f_{\theta}(x)\| = \lim_{\delta \rightarrow 0} \sup_{\substack{x' \in \mathcal{M} : x' \neq x, \\ d_{\mathcal{M}}(x, x') \leq \delta}} \frac{|f_{\theta}(x) - f_{\theta}(x')|}{d_{\mathcal{M}}(x, x')} \leq L, \quad \text{for all } x \in \mathcal{M}. \quad (2)$$

With definition 1 in place and restricting attention to differentiable functions f_{θ} our stated goal of learning functions f_{θ} with controlled variability over the manifold \mathcal{M} can then be written as

$$\begin{aligned} P^* = & \min_{\theta \in \Theta, \rho \geq 0} \quad \rho, \\ \text{subject to} \quad & \mathbb{E}_{p(x, y)}[\ell(f_{\theta}(x), y)] \leq \epsilon, \\ & \|\nabla_{\mathcal{M}} f_{\theta}(z)\|^2 \leq \rho, \quad p(z)\text{-a.e.}, \quad z \in \mathcal{M}. \end{aligned} \quad (3)$$

In this formulation the statistical loss $\mathbb{E}_{p(x, y)}[\ell(f_{\theta}(x), y)]$ is required to be below a target level ϵ . Of all functions f_{θ} that can satisfy this loss requirement, Problem (3) defines as optimal the one whose Lipschitz constant $L = \sqrt{\rho}$ is the smallest.

The goal of this paper is to develop methodologies to solve (3) when the data distribution and the manifold are unknown. To characterize the distribution we are given sample pairs (x_i, y_i) drawn from the joint distribution $p(x, y)$. To characterize the manifold we are given samples z_i drawn from the marginal distribution $p(x)$. This includes the samples x_i from the (labeled) data pairs (x_i, y_i) and may also include (unlabeled) samples z_i .

It is interesting to observe that (3) shows that the problems of manifold regularization (Belkin et al., 2005; Niyogi, 2013; Kejani et al., 2020; Li et al., 2022) and Lipschitz constant control (Oberman and Calder, 2018; Finlay et al., 2018a; Couellan, 2021; Finlay et al., 2018b; Pauli et al., 2021; Krishnan et al., 2020; Shi et al., 2019; Lindemann et al., 2021; Arghal et al., 2021) are related. This connection is important to understand the merit of (3). To explain this better observe that there are three motivations for the problem formulation in (3): (i) It is often the case that if samples x_1 and x_2 are close, then the conditional distributions $p(y | x_1)$ and $p(y | x_2)$ are close as well. A function f_{θ} with small Lipschitz constant leverages this property. (ii) The Lipschitz constant of f_{θ} is *guaranteed* to be smaller than $L = \sqrt{\rho}$. This provides advantages in, e.g., physical systems where Lipschitz constant guarantees translates to stability and safety assurances. (iii) It leverages the intrinsic low-dimensional structure of the manifold \mathcal{M} embedded in the ambient space. In particular, this permits taking advantage of unlabeled data.

Motivations (i) and (ii) are tropes of the Lipschitz regularization literature; e.g., Oberman and Calder (2018); Finlay et al. (2018a); Couellan (2021); Finlay et al. (2018b); Pauli et al. (2021); Krishnan et al. (2020); Shi et al. (2019); Lindemann et al. (2021); Arghal et al. (2021). Indeed, the problem formulation in (3) is inspired in similar problem formulations in which the Lipschitz constant is regularized in the ambient space,

$$\begin{aligned} & \minimize_{\theta \in \Theta, \rho \geq 0} \quad L, \\ \text{subject to} \quad & \mathbb{E}_{p(x, y)}[\ell(f_{\theta}(x), y)] \leq \epsilon, \\ & |f_{\theta}(y) - f_{\theta}(z)| \leq L \|y - z\|, \quad (y, z) \sim p(x). \end{aligned} \quad (4)$$

A difference between (3) and (4) is that in the latter we use a Lipschitz condition that does not require differentiability. A more important difference is that in (4) the Lipschitz constant is regularized in the ambient space. The distance between features y and z in (4) is the Euclidean distance $\|y - z\|$. This is disparate from the manifold metric $d_{\mathcal{M}}(y, z)$ that is implicit in the manifold gradient constraint in (3). Thus, the formulation in (3) improves upon (4) because of it leverages the structure of the manifold \mathcal{M} [cf. Motivation (iii)].

Motivations (i) and (iii) are themes of the manifold regularization literature (Belkin et al., 2005; Niyogi, 2013; Kejani et al., 2020; Li et al., 2022). And, indeed, it is ready to conclude by invoking Green's first identity (see Section 3.2) that the formulation in (3) is also inspired in the manifold regularization problem,

$$\minimize_{\theta \in \Theta} \quad \mathbb{E}_{p(x, y)}[\ell(f_{\theta}(x), y)] + \gamma \int_{\mathcal{M}} \|\nabla_{\mathcal{M}} f_{\theta}(z)\|^2 p(z) dV(z). \quad (5)$$

The difference between (3) and (5) is that in the latter the manifold Lipschitz constant is added as a regularization penalty. This is disparate from the imposition of a manifold Lipschitz constraint in (3). The regularization in (5) *favors* solutions with small Lipschitz constant by penalizing large Lipschitz constants, while the constraint in (3) *guarantees* that the Lipschitz constant is bounded by ρ . This is the distinction between *regularizing* a Lipschitz constant versus *constraining* a Lipschitz constant. The constraint in (5) is also imposed at *all* points in the manifold, whereas the regularization in (5) is an average over the manifold. Taking an average allows for large Lipschitz constants at some specific points if this is canceled out by small Lipschitz constants in other points of the manifold. Both of these observations imply that (3) improves upon (5) because it offers global smoothness guarantees that are important in, e.g., physical systems [cf. Motivation (iii)].

Remark 1 (Alternative Manifold Lipschitz Control Formulations). There are three arbitrary choices in (3). (a) We choose to constrain the average statistical loss $\mathbb{E}_{p(x,y)}[\ell(f_\theta(x), y)] \leq \epsilon$; (b) we choose to constrain the pointwise Lipschitz constant $\|\nabla_{\mathcal{M}} f_\theta(z)\|^2 \leq \rho$; and (c) we choose as our objective to require a target loss ϵ and minimize the Lipschitz constant $L = \sqrt{\rho}$. We can alternatively choose to constrain the pointwise loss $\ell(f_\theta(x), y) \leq \epsilon$, to constrain the average Lipschitz constant $\int_{\mathcal{M}} \|\nabla_{\mathcal{M}} f_\theta(z)\|^2 p(z) dV(z) \leq \rho$ [cf. (5)], or to require a target smoothness $L = \sqrt{\rho}$ and minimize the loss ϵ . All of the possible eight combinations of choices are of interest. We formulate (3) because it is the most natural intersection between the regularization of Lipschitz constants in ambient spaces [cf. (4)] and manifold regularization [cf. (5)]. The techniques we develop in this paper can be adapted to any of the other seven alternative formulations.

3 LEARNING WITH GLOBAL LIPSCHITZ CONSTRAINTS

Problem (3) is a constrained learning problem that we will solve in the dual domain (Chamon and Ribeiro, 2020). To that end, observe that (3) has statistical and pointwise constraints. The loss constraint $\mathbb{E}_{p(x,y)}[\ell(f_\theta(x), y)] \leq \epsilon$ is said to be statistical because it restricts the expected loss over the data distribution. The Lipschitz constant constraints $\|\nabla_{\mathcal{M}} f_\theta(z)\|^2 \leq \rho$ are said to be pointwise because they are imposed for all individual points in the manifold except for a set of zero measure. Consider then a Lagrange multiplier μ associated with the statistical constraint $\mathbb{E}_{p(x,y)}[\ell(f_\theta(x), y)] \leq \epsilon$ and a Lagrange multiplier distribution $\lambda(z)$ associated with the set of pointwise constraints $\|\nabla_{\mathcal{M}} f_\theta(z)\|^2 \leq \rho$. The dual problem associated with (3) can then be written as

$$\begin{aligned} D^* = \max_{\mu, \lambda \geq 0} \min_{\theta} L(\theta, \mu, \lambda) &:= \mu \left(\mathbb{E}[\ell(f_\theta(x), y)] - \epsilon \right) + \int_{\mathcal{M}} \lambda(z) \|\nabla_{\mathcal{M}} f_\theta(z)\|^2 p(z) dV(z), \\ \text{subject to } \int_{\mathcal{M}} \lambda(z) p(z) dV(z) &= 1. \end{aligned} \quad (6)$$

We point out that in (6) we remove ρ from the Lagrangian by incorporating the dual variable constraint $\int_{\mathcal{M}} \lambda(x) p(x) dV(x) = 1$; see Appendix A for details.

We henceforth use the dual problem (6) in lieu of (3). Since we are interested in situations in which we do not have access to the data distribution $p(x, y)$ we further consider empirical versions of (6). Consider then N i.i.d. samples (x_n, y_n) drawn from $p(x, y)$ and define the empirical dual problem as,

$$\begin{aligned} \hat{D}^* = \max_{\hat{\mu}, \hat{\lambda} \geq 0} \min_{\theta} \hat{L}(\theta, \hat{\mu}, \hat{\lambda}) &:= \hat{\mu} \left(\frac{1}{N} \sum_{n=1}^N \ell(f_\theta(x_n), y_n) - \epsilon \right) + \frac{1}{N} \sum_{n=1}^N \hat{\lambda}(x_n) \|\nabla_{\mathcal{M}} f_\theta(x_n)\|^2, \\ \text{subject to } \frac{1}{N} \sum_{n=1}^N \hat{\lambda}(x_n) &= 1, \end{aligned} \quad (7)$$

where, to simplify notation, we assume no unlabeled samples are available. If unlabeled samples are given the modification is straight-forward; see Appendix B.

The remainder of this section provides three technical contributions:

- (C2) To justify solutions of (7) we must show statistical consistency with respect to the primal problem (3). This is challenging because of two reasons: (i) Since we did not assume the use of a convex parameterization, (3) is not a convex problem on θ . Thus, the primal and dual problems are not necessarily equivalent. (ii) Since we are maximizing over the dual variables μ and $\lambda(z)$, we do not know if the empirical dual formulation in (7) is close to the statistical dual formulation in (6). We will overcome these two challenges and show that the empirical dual problem (7) is a consistent approximation of the statistical primal problem (Proposition 1 in Section 3.1).
- (C3) Solving (7) requires evaluating the gradient norm sum $(1/N) \sum_{n=1}^N \hat{\lambda}(x_n) \|\nabla_{\mathcal{M}} f_{\theta}(x_n)\|^2$. We will generalize results from the manifold regularization literature to show that under regularity conditions on λ the gradient norm integral can be computed in a more amenable form utilizing a weighted point cloud Laplacian (Proposition 3 in Section 3.2).
- (C4) We introduce a primal-dual algorithm to solve (7) (Section 3.3).

3.1 STATISTICAL CONSISTENCY OF THE EMPIRICAL DUAL PROBLEM

We show in this section that (7) is close to (3). Doing so requires the following assumptions:

Assumption 1. *The loss ℓ is M -Lipschitz continuous, B -bounded, and convex.*

Assumption 2. *Let $\mathcal{H} = \{f_{\theta} \mid \theta \in \Theta \subset \mathbb{R}^Q\}$ with compact Θ be the hypothesis class, and let $\bar{\mathcal{H}} = \overline{\text{conv}}(\mathcal{H})$ be the closure of its convex hull. For each $\nu > 0$ and $\varphi \in \bar{\mathcal{H}}$ there exists $\theta \in \Theta$ such that simultaneously $\sup_{z \in \mathcal{M}} |\varphi(z) - f_{\theta}(z)| \leq \nu$ and $\sup_{z \in \mathcal{M}} \|\nabla_{\mathcal{M}} \varphi(z) - \nabla_{\mathcal{M}} f_{\theta}(z)\| \leq \nu$.*

Assumption 3. *The hypothesis class \mathcal{H} is G -Lipschitz in its gradients, i.e. $\|\nabla_{\mathcal{M}} f_{\theta_1}(z) - \nabla_{\mathcal{M}} f_{\theta_2}(z)\| \leq G |f_{\theta_1}(z) - f_{\theta_2}(z)|$.*

Assumption 4. *There exists $\zeta(N, \delta) \geq 0$, and $\hat{\zeta}(N, \delta) \geq 0$ monotonically decreasing with N , such that*

$$\left| \mathbb{E}[\ell(f_{\theta}(x), y)] - \frac{1}{N} \sum_{n=1}^N \ell(f_{\theta}(x_n), y_n) \right| \leq \zeta(N, \delta), \left| \mathbb{E}[f_{\theta}(x)] - \frac{1}{N} \sum_{n=1}^N f_{\theta}(x_n) \right| \leq \hat{\zeta}(N, \delta), \quad (8)$$

for all $\theta \in \Theta$, with probability $1 - \delta$ over independent draws $(x_n, y_n) \sim p$.

Assumption 5. *There exists a feasible solution $\tilde{\theta} \in \Theta$ such that $\mathbb{E}[\ell(f_{\tilde{\theta}}(x), y)] < \epsilon - M\nu$.*

Assumption 1 holds for most losses utilized in practice. Assumption 2 is a requirement on the richness of the parametrization. In the particular case of neural networks the covering constant ν is upper bounded by the universal approximation bound of the neural network. Assumption 3 also holds for neural networks if we use smooth nonlinearities – e.g., hyperbolic tangents. The *uniform convergence* property (8) is customary in learning theory to prove PAC learnability and is implied by bounds on complexity measures such as VC dimension or Rademacher complexity (Mohri et al., 2018; Vapnik, 1999; Shalev-Shwartz and Ben-David, 2014). The following theorem provides the desired bound.

Proposition 1. *Let $\hat{\mu}^*, \hat{\lambda}^*$ be solutions of the empirical dual problem (7). Under assumptions 1–5, there exists $\hat{\theta}^* \in \arg\min_{\theta} \hat{L}(\theta, \hat{\mu}^*, \hat{\lambda}^*)$ such that, with probability $1 - 5\delta$,*

$$|P^* - \hat{D}^*| \leq \mathcal{O}(\nu) + (1 + \Delta)\zeta(N, \delta) + \mathcal{O}(\hat{\zeta}(N, \delta)), \quad (9)$$

where $\Delta = \max(\hat{\mu}^*, \mu^*) \leq C$ for a constant $C < \infty$, where $\mu^*, \hat{\mu}^*$ are solutions of (6), and (7) respectively.

Proposition 1 shows that the empirical dual problem 7 is statistically consistent. That is to say, for any realization of N according to p , the difference between the empirical dual problem (7) and the statistical smooth learning problem (3) decreases as N increases. This difference is bounded in terms of the richness of the parametrization (ν), the difficulty of the fit requirement (as expressed by the optimal dual variables $\mu^*, \hat{\mu}^*$), and the number of samples (N). The guarantee has a form typical for constrained learning problems (Chamon and Ribeiro, 2020). Proposition 1 states that we are able to predict what is the minimum norm of the gradient that a function class can have while satisfying an expected ϵ loss. This

is important, because we do not require access to the distribution of the samples p , only a set of N samples sampled according to this distribution. On the other hand, Proposition 1 does not state that by solving the dual problem 7 we will obtain a solution of the primal problem 3. The following proposition provides a bound on the near feasibility of the solution of 7 with respect to the solution of 3.

Proposition 2. *Let $\hat{\mu}^*, \hat{\lambda}^*$ be solutions of the empirical dual problem (7). Under assumptions 1–5, there exists $\hat{\theta}^* \in \operatorname{argmin}_{\theta} \hat{L}(\theta, \hat{\mu}^*, \hat{\lambda}^*)$ such that, with probability $1 - 5\delta$,*

$$\max_{z \in \mathcal{M}} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\|^2 \leq P^* + \hat{\mu} \left| \frac{1}{N} \sum_{n=1}^N \ell(f_{\hat{\theta}^*}(x_n), y_n) - \epsilon \right| + \mathcal{O}(\nu) + \mathcal{O}(\hat{\zeta}(N, \delta)), \text{ and} \quad (10)$$

$$\mathbb{E}[\ell(f_{\hat{\theta}^*}(x), y)] \leq \epsilon + \zeta(N, \delta). \quad (11)$$

Proposition 2 provides near optimality, and near feasibility conditions for solutions $\hat{\theta}^*$ obtained through the empirical dual problem (7). The difference between the maximum gradient of the obtained solution θ and the optimal value P^* is bounded by the number of samples N , as well as the empirical constraint violation. Notice that even though the optimal dual variable $\hat{\mu}$ is not known, the constraint violation can be evaluated in practice, as it only requires to evaluate the obtained function over the N given samples.

Remark 2 (Interpolators). *In practice, the number of parameters in a parametric function (e.g. Neural Network) tends to exceed the dimension of the input, which allows functions to interpolate the data, i.e. to attain zero loss on the dataset. Proposition 2, presents a connection to interpolating functions. By setting $\epsilon = 0$, if the function achieves zero error over the empirical distribution i.e. $\ell(f_{\hat{\theta}^*}(x_n), y_n) = 0$, for all $n \in [N]$, then the dependency on μ^* disappears. This implies that over interpolating classifiers, the one with the minimum Lipschitz constant over the samples, is probably the one that attains the minimum Lipschitz constant over the whole manifold.*

3.2 FROM MANIFOLD GRADIENT TO DISCRETE LAPLACIAN

We derive an alternative way of computing the integral of the norm of the gradient utilizing samples. To do so, we define the normalized point cloud Laplacian according a probability distribution λ .

Definition 2 (Point-cloud Laplacian). *Consider a set of points $x_1, \dots, x_N \in \mathcal{M}$, sampled according to probability $\lambda : \mathcal{M} \rightarrow \mathbb{R}$. The normalized graph Laplacian of f_{θ} at $z \in \mathcal{M}$ is defined as*

$$\mathbf{L}_{\lambda, N}^t f_{\theta}(z) = \frac{1}{N} \sum_{n=1}^N W(z, x_n) (f_{\theta}(z) - f_{\theta}(x_n)), \quad (12)$$

for

$$W(z, x_n) = \frac{1}{t} \frac{G_t(z, x_n)}{\sqrt{\hat{w}(z) \hat{W}(x_n)}}, \quad \text{with } G_t(z, x_n) = \frac{1}{(4\pi t)^{d/2}} e^{-\frac{\|z - x_n\|^2}{4t}}, \quad (13)$$

$$\hat{w}(z) = \frac{1}{N} \sum_{n=1}^N G_t(z, x_n), \quad \text{and } \hat{W}(x_n) = \frac{1}{N-1} \sum_{m \neq n} G_t(x_m, x_n). \quad (14)$$

As long as the function considered is smooth enough, the following convergence result holds:

Proposition 3 (Point Cloud Estimate). *Let Λ be the set of probability distributions defined a compact d -dimensional differentiable manifold \mathcal{M} isometrically embedded in \mathbb{R}^D such that $\Lambda = \{\lambda : 0 < a \leq \lambda(z) \leq b < \infty, |\frac{\partial \lambda}{\partial x}| \leq c < \infty \text{ and } |\frac{\partial^2 \lambda}{\partial x^2}| \leq d < \infty \text{ for all } z \in \mathcal{M}\}$, and let f_{θ} , with $\theta \in \Theta$, be a family of functions with uniformly bounded derivatives up to order 3 vanishing at the boundary $\partial \mathcal{M}$. For any $\epsilon > 0, \delta > 0$, for all $N > N_0$,*

$$P \left[\sup_{\lambda \in \Lambda, \theta \in \Theta} \left| \int \|\nabla_{\mathcal{M}} f_{\theta}(z)\|^2 \lambda(z) dV(z) - \frac{1}{N} \sum_{i=1}^N f_{\theta}(z_i) \mathbf{L}_{\lambda, N}^t f_{\theta}(z_i) \lambda(z_i) \right| > \epsilon \right] \leq \delta \quad (15)$$

where the point cloud laplacian $\mathbf{L}_{\lambda, N}^t f_{\theta}$ is as defined in 2, with $t = N^{-\frac{1}{d+2+\alpha}}$ for any $\alpha > 0$.

The proof of Proposition 3 relies on two steps. First, we relate the integral of the norm of the gradient over the manifold with the integral of the continuous Laplace-Beltrami operator by virtue of Green’s identity. Second, we approximate the value of the Laplace-Beltrami operator by the point-cloud Laplacian. Proposition 3 connects the integral of the norm of the gradient of function f_θ with a point cloud Laplacian operator, this result connects the dual problem 7, with the primal 3 while allowing for a more amenable way of computing the integral.

Remark 3 (Laplacian Regularization and Manifold Lipschitz). *The dual problem (6) is closely related to the manifold regularization problem (5). In particular, the two become equivalent by substituting $\rho = \mu^{-1}$, and utilizing a uniform distribution for λ . The key difference between the problems is given by the dual variable λ , which can be thought as a probability distribution over the manifold that penalizes regions of the manifold where the norm of the gradient of f_θ is larger. The standard procedure in Laplacian regularization is to calculate the graph-Laplacian of the set of points \mathbf{L} utilizing the heat kernel, and compute the integral by $\mathbf{f}^T \mathbf{L} \mathbf{f}$, where $\mathbf{f} = [f_\theta(x_1), \dots, f_\theta(x_n)]^T$. In the case of Manifold Lipschitz, the same product can be computed, but utilizing the re-weighted point cloud laplacian.*

3.3 DUAL ASCENT ALGORITHM

We outline an iterative and empirical primal-dual algorithm to solve the dual problem. Upon the initialization of θ_0 , and λ_0, μ_0 , we set to minimize the dual function as follows,

$$\theta_{k+1} = \theta_k - \eta_\theta \nabla_\theta L(\theta_k, \mu_k, \lambda_k), \quad (16)$$

where η_θ is a positive step-size. Note that to solve problem (16), we can either utilize the gradient version the the lagrangian (cf. equation (7)), or the point-cloud Laplacian (cf. equation (15)). Consequent to updating θ , we update the dual variables as follows,

$$\mu_{k+1} = \left[\mu_k + \eta_\mu \left(\frac{1}{N} \sum_{n=1}^N \ell(f_\theta(x_n), y_n) - \epsilon \right) \right]_+, \quad (17)$$

$$\tilde{\lambda}_{k+1}(x_n) = \lambda_k(x_n) + \eta_\lambda \|\nabla_{\mathcal{M}} f_\theta(x_n)\|^2 \text{ for all data points } x_n \quad (18)$$

$$\lambda_{k+1} = \underset{\lambda}{\operatorname{argmin}} \|\tilde{\lambda}_{k+1} - \lambda_{k+1}\|, \text{ such that } \sum_{n=1}^N \lambda_{k+1}(x_n) = N. \quad (19)$$

where η_μ, η_λ are a positive step-sizes. Note that we require a convex projection over $\tilde{\lambda}$ to satisfy the normalizing constraint. In step (18) we require to estimate the norm of the gradient at data point x_n , which we estimate for the neighboring data points. Intuitively, the primal dual procedure increases the value of $\lambda(x_n)$ at points in which the norm of the gradient is larger. The role of μ is to enforce the loss ℓ to be smaller than ϵ , a larger value of μ would increase the relative importance over the norm of the integral. See Appendix F.

4 NUMERICAL RESULTS

4.1 NAVIGATION CONTROLS PROBLEM

In this section, we consider the problem of continuous navigation of an agent. The agent’s objective is to reach a goal while avoiding obstacles. The state space of the agent is given by its position, and navigates by taking actions on the velocity. We construct a square grid of points in the free space i.e. outside of the obstacles, and find the shortest path between two starting positions and the goal along the grid. For those two grid trajectories, we compute the optimal actions to be taken at each point in order to follow the trajectory. The learner is equipped with both labeled trajectories, as well as the unlabeled point grid. To leverage the manifold

Method	Trajectories
ERM	85
Ambient Reg.	66
Manifold Reg.	77
Manifold Lipschitz	94

Table 1: Number of successful trajectories from 100 random starting points.

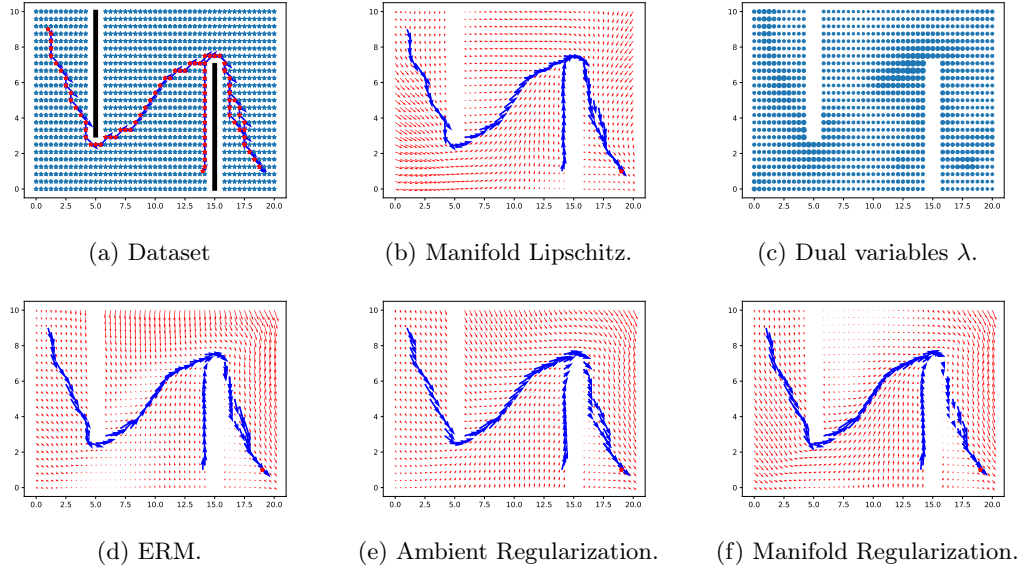


Figure 2: Figure 2a shows the training dataset, blue stars depict unlabeled point, and blue arrow the optimal action at the red star. Figure 2b shows the learned function using Manifold Learning, and 2c its associated dual variables associated. Figures 2d, 2e, 2f show the functions learned using ERM, ambient regularization, and Manifold regularization respectively.

structure of the data, we consider the grid of points, and we construct the point cloud Laplacian considering adjacent points in the grid. As measure of merit, we take 100 random points and we compute the trajectories. A trajectory is successful if it achieves the goal without colliding. The results are shown in Table 1, and the learned functions in Figure 2.

4.2 GROUND ROBOT ERROR PREDICTION

In this real world experiment, we seek to learn the error that a model would make in predicting the dynamics of a ground robot Koppel et al. (2016). We posses trajectories, in the form of time series of the control signals of an *iRobot* Packbot making turns on both pavement, and grass. For each trajectory, the mismatch between the real and the model-predicted states is quantified. Given a dataset of trajectories and errors made by the model, the objective is to learn the error that a model will make. Details of the experiment can be found in G.3.

Method	Grass	Pavement
ERM	0.42	0.0120
Ambient Reg.	0.31	0.0065
Manifold Reg.	0.38	0.0045
Manifold Lipschitz	0.25	0.0032

Table 2: Error Prediction accuracy for the Ground Robot Experiment.

5 CONCLUSION

In this work, we presented a constrain learning method to obtain smooth functions over manifold data. We have shown that under mild conditions, the problem of finding smooth functions over a manifold can be reformulated as a weighted point cloud Laplacian penalty over varying probability distributions whose dynamics are govern by the constraint violations. Three numerical examples validate the empirical advantages of obtaining functions that vary smoothly over the data.

REFERENCES

- Adam M Oberman and Jeff Calder. Lipschitz regularized deep neural networks converge and generalize. *arXiv preprint arXiv:1808.09540*, 2018.
- Chris Finlay, Adam M. Oberman, and Bilal Abbasi. Improved robustness to adversarial examples using lipschitz regularization of the loss. *CoRR*, abs/1810.00953, 2018a. URL <http://arxiv.org/abs/1810.00953>.
- Nicolas Couellan. The coupling effect of lipschitz regularization in neural networks. *SN Computer Science*, 2(2):1–9, 2021.
- Chris Finlay, Jeff Calder, Bilal Abbasi, and Adam Oberman. Lipschitz regularized deep neural networks generalize and are adversarially robust. *arXiv preprint arXiv:1808.09540*, 2018b.
- Patricia Pauli, Anne Koch, Julian Berberich, Paul Kohler, and Frank Allgöwer. Training robust neural networks using lipschitz bounds. *IEEE Control Systems Letters*, 6:121–126, 2021.
- Vishaal Krishnan, Al Makdah, Abed AlRahman, and Fabio Pasqualetti. Lipschitz bounds and provably robust training by laplacian smoothing. *Advances in Neural Information Processing Systems*, 33:10924–10935, 2020.
- Guanya Shi, Xichen Shi, Michael O’Connell, Rose Yu, Kamyar Azizzadenesheli, Animashree Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural lander: Stable drone landing control using learned dynamics. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9784–9790. IEEE, 2019.
- Lars Lindemann, Haimin Hu, Alexander Robey, Hanwen Zhang, Dimos Dimarogonas, Stephen Tu, and Nikolai Matni. Learning hybrid control barrier functions from data. In *Conference on Robot Learning*, pages 1351–1370. PMLR, 2021.
- Raghu Arghal, Eric Lei, and Shirin Saeedi Bidokhti. Robust graph neural networks via probabilistic lipschitz constraints. *arXiv preprint arXiv:2112.07575*, 2021.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Liquan Zhao and Yan Liu. Spectral normalization for domain adaptation. *Information*, 11(2):68, 2020.
- Leon Bungert, René Raab, Tim Roith, Leo Schwinn, and Daniel Tenbrinck. Clip: Cheap lipschitz training of neural networks. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 307–319. Springer, 2021.
- Misha Belkin, Partha Niyogi, and Vikas Sindhwani. On manifold regularization. In *International Workshop on Artificial Intelligence and Statistics*, pages 17–24. PMLR, 2005.
- Partha Niyogi. Manifold regularization and semi-supervised learning: Some theoretical analyses. *Journal of Machine Learning Research*, 14(5), 2013.
- Zengyi Li, Yubei Chen, Yann LeCun, and Friedrich T Sommer. Neural manifold clustering and embedding. *arXiv preprint arXiv:2201.10000*, 2022.
- Matthias Hein, Jean-Yves Audibert, and Ulrike von Luxburg. From graphs to manifolds—weak and strong pointwise consistency of graph laplacians. In *International Conference on Computational Learning Theory*, pages 470–485. Springer, 2005.
- Mikhail Belkin and Partha Niyogi. Towards a theoretical foundation for laplacian-based manifold methods. In *International Conference on Computational Learning Theory*, pages 486–500. Springer, 2005.
- M Tavassoli Kejani, Fadi Dornaika, and H Talebi. Graph convolution networks with manifold regularization for semi-supervised learning. *Neural Networks*, 127:160–167, 2020.

-
- Mikhail Belkin and Partha Niyogi. Semi-supervised learning on riemannian manifolds. *Machine learning*, 56(1):209–239, 2004.
- Bo Jiang, Ziyang Zhang, Doudou Lin, Jin Tang, and Bin Luo. Semi-supervised learning with graph learning-convolutional networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11313–11320, 2019.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, pages 40–48. PMLR, 2016.
- Xiaojin Jerry Zhu. Semi-supervised learning literature survey. *University of Wisconsin-Madison Department of Computer Sciences*, 2005.
- Bruno Lecouat, Chuan-Sheng Foo, Houssam Zenati, and Vijay R Chandrasekhar. Semi-supervised learning with gans: Revisiting manifold regularization. *arXiv preprint arXiv:1805.08957*, 2018.
- Yassine Ouali, Céline Hudelot, and Myriam Tami. An overview of deep semi-supervised learning. *arXiv preprint arXiv:2006.05278*, 2020.
- Vivien Cabannes, Loucas Pillaud-Vivien, Francis Bach, and Alessandro Rudi. Overcoming the curse of dimensionality with laplacian regularization in semi-supervised learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- Luiz Chamon and Alejandro Ribeiro. Probably approximately correct constrained learning. *Advances in Neural Information Processing Systems*, 33:16722–16735, 2020.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 1999.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Alec Koppel, Jonathan Fink, Garrett Warnell, Ethan Stump, and Alejandro Ribeiro. Online learning for characterizing unknown environments in ground robotic vehicle models. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 626–633, 2016. doi: 10.1109/IROS.2016.7759118.
- Luiz FO Chamon, Santiago Paternain, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained learning with non-convex losses. *arXiv preprint arXiv:2103.05134*, 2021.
- Andrzej Ruszczyński. *Nonlinear optimization*. Princeton university press, 2011.
- Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- Manfredo Perdigao Do Carmo and J Flaherty Francis. *Riemannian geometry*, volume 6. Springer, 1992.
- Steven Rosenberg. *The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds*. London Mathematical Society Student Texts. Cambridge University Press, 1997. doi: 10.1017/CBO9780511623783.
- Weiran Wang and Miguel A Carreira-Perpinán. Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. *arXiv preprint arXiv:1309.1541*, 2013.

-
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- Valentin Khruikov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic image embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Tianbao Yang. Advancing non-convex and constrained learning: Challenges and opportunities. *AI Matters*, 5(3):29–39, 2019.
- Alexander Robey, Luiz Chamon, George J Pappas, Hamed Hassani, and Alejandro Ribeiro. Adversarial robustness with semi-infinite constrained learning. *Advances in Neural Information Processing Systems*, 34:6198–6215, 2021.
- Juan Cervino, Luana Ruiz, and Alejandro Ribeiro. Training stable graph neural networks through constrained learning. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4223–4227. IEEE, 2022.
- Zebang Shen, Juan Cervino, Hamed Hassani, and Alejandro Ribeiro. An agnostic approach to federated learning with class imbalance. In *International Conference on Learning Representations*, 2021.
- Juan Elenter, Navid NaderiAlizadeh, and Alejandro Ribeiro. A lagrangian duality approach to active learning. *arXiv preprint arXiv:2202.04108*, 2022.
- Santiago Paternain, Luiz Chamon, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained reinforcement learning has zero duality gap. *Advances in Neural Information Processing Systems*, 32, 2019.
- Santiago Paternain, Miguel Calvo-Fullana, Luiz FO Chamon, and Alejandro Ribeiro. Safe policies for reinforcement learning via primal-dual methods. *IEEE Transactions on Automatic Control*, 2022.
- Agustin Castellano, Hancheng Min, Juan Bazerque, and Enrique Mallada. Reinforcement learning with almost sure constraints. *arXiv preprint arXiv:2112.05198*, 2021.
- Qinbo Bai, Amrit Singh Bedi, Mridul Agarwal, Alec Koppel, and Vaneet Aggarwal. Achieving zero constraint violation for constrained reinforcement learning via primal-dual approach. *arXiv preprint arXiv:2109.06332*, 2021.
- Mohammadhosein Hasanbeig, Alessandro Abate, and Daniel Kroening. Logically-constrained reinforcement learning. *arXiv preprint arXiv:1801.08099*, 2018.
- Mark Eisen, Clark Zhang, Luiz FO Chamon, Daniel D Lee, and Alejandro Ribeiro. Learning optimal resource allocations in wireless systems. *IEEE Transactions on Signal Processing*, 67(10):2775–2790, 2019.
- Shufei Zhang, Kaizhu Huang, Jianke Zhu, and Yang Liu. Manifold adversarial training for supervised and semi-supervised learning. *Neural Networks*, 140:282–293, 2021. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2021.03.031>. URL <https://www.sciencedirect.com/science/article/pii/S0893608021001192>.
- Marc Khoury and Dylan Hadfield-Menell. On the geometry of adversarial examples. *arXiv preprint arXiv:1811.00525*, 2018.
- Xingjun Ma, Bo Li, Yisen Wang, Sarah M Erfani, Sudanthi Wijewickrema, Grant Schoenbeck, Dawn Song, Michael E Houle, and James Bailey. Characterizing adversarial subspaces using local intrinsic dimensionality. *arXiv preprint arXiv:1801.02613*, 2018.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Jonathan Uesato, and Pascal Frossard. Robustness via curvature regularization, and vice versa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9078–9086, 2019.

-
- Charles Jin and Martin Rinard. Manifold regularization for locally stable deep neural networks. *arXiv preprint arXiv:2003.04286*, 2020.
- Carlos Lassance, Vincent Gripon, and Antonio Ortega. Laplacian networks: Bounding indicator function smoothness for neural networks robustness. *APSIPA Transactions on Signal and Information Processing*, 10, 2021.
- David Stutz, Matthias Hein, and Bernt Schiele. Disentangling adversarial robustness and generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6976–6987, 2019.
- Mihaela Rosca, Theophane Weber, Arthur Gretton, and Shakir Mohamed. A case for new neural network smoothness constraints. In Jessica Zosa Forde, Francisco Ruiz, Melanie F. Pradier, and Aaron Schein, editors, *Proceedings on "I Can't Believe It's Not Better!" at NeurIPS Workshops*, volume 137 of *Proceedings of Machine Learning Research*, pages 21–32. PMLR, 12 Dec 2020. URL <https://proceedings.mlr.press/v137/rosca20a.html>.
- Sébastien Bubeck and Mark Sellke. A universal law of robustness via isoperimetry. *Advances in Neural Information Processing Systems*, 34, 2021.
- Matthias Hein, Jean-Yves Audibert, and Ulrike von Luxburg. Graph laplacians and their convergence on random neighborhood graphs. *Journal of Machine Learning Research*, 8(6), 2007.
- David B Dunson, Hau-Tieng Wu, and Nan Wu. Spectral convergence of graph laplacian and heat kernel reconstruction in l_∞ from random samples. *Applied and Computational Harmonic Analysis*, 55:282–336, 2021.
- Hau-Tieng Wu and Nan Wu. Think globally, fit locally under the manifold setup: Asymptotic analysis of locally linear embedding. *The Annals of Statistics*, 46(6B):3805–3837, 2018.
- Mahyar Fazlyab, Alexander Robey, Hamed Hassani, Manfred Morari, and George Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Tom Z Jiahao, Kong Yao Chee, and M Ani Hsieh. Online dynamics learning for predictive control with an application to aerial robots. *arXiv preprint arXiv:2207.09344*, 2022.
- Kong Yao Chee, Tom Z Jiahao, and M Ani Hsieh. Knode-mpc: A knowledge-based data-driven predictive control framework for aerial robots. *IEEE Robotics and Automation Letters*, 7(2):2819–2826, 2022.
- Tom Z Jiahao, M Ani Hsieh, and Eric Forgoston. Knowledge-based learning of nonlinear dynamics and chaos. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(11): 111101, 2021.

A DUAL PROBLEM FORMULATION

In Section 3 we introduce the optimization program in (6) as the dual of (3). Strictly speaking (6) is equivalent to the actual dual problem of (3). This follows from a ready reformulation of the dual problem as we show in the following proposition.

Proposition 4. *The optimization program in (6) is equivalent to the Lagrangian dual of (3).*

Proof. The result is true because the dual problem is linear in ρ . To see this, recall that μ is the dual variable associated with the statistical constraint $\mathbb{E}_{p(x,y)}[\ell(f_\theta(x), y)] \leq \epsilon$ and that $\lambda(z)$ is the Lagrange multiplier distribution associated with the set of pointwise constraints $\|\nabla_{\mathcal{M}} f_\theta(z)\|^2 \leq \rho$. The Lagrangian $\tilde{\mathcal{L}}(\rho, \theta, \mu, \lambda)$ of (3) is therefore written as

$$\tilde{\mathcal{L}}(\theta, \rho, \mu, \lambda) = \rho + \mu \left(\mathbb{E}[\ell(f_\theta(x), y)] - \epsilon \right) + \int_{\mathcal{M}} \lambda(z) \left(\|\nabla f_\theta(z)\|^2 - \rho \right) p(z) dV(z) \quad (20)$$

Reorder terms in (20) to group the two summands that involve ρ to write

$$\begin{aligned} \tilde{\mathcal{L}}(\theta, \rho, \mu, \lambda) = & \rho \left(1 - \int \lambda(z) p(z) dV(z) dz \right) + \mu \left(\mathbb{E}_{p(x,y)}[\ell(f_\theta(x), y)] - \epsilon \right) \\ & + \int \lambda(z) \|\nabla_{\mathcal{M}} f_\theta(z)\|^2 p(z) dV(z). \end{aligned} \quad (21)$$

An important observation to make is that the Lagrangian decomposes in a term that involves only ρ and a term that involves only f_θ . Define then

$$\begin{aligned} \tilde{\mathcal{L}}_1(\rho, \lambda) &:= \rho \left(1 - \int \lambda(z) p(z) dV(z) dz \right), \\ \tilde{\mathcal{L}}_2(\theta, \mu, \lambda) &:= \mu \left(\mathbb{E}_{p(x,y)}[\ell(f_\theta(x), y)] - \epsilon \right) + \int \lambda(z) \|\nabla_{\mathcal{M}} f_\theta(z)\|^2 p(z) dV(z), \end{aligned} \quad (22)$$

so that we can write the Lagrangian in (21) as $\tilde{\mathcal{L}}(\theta, \rho, \mu, \lambda) = \tilde{\mathcal{L}}_1(\rho, \lambda) + \tilde{\mathcal{L}}_2(\theta, \mu, \lambda)$.

The dual problem can now be written as the maximization over multipliers of the minimum of the Lagrangian over primal variables

$$\begin{aligned} \tilde{D}^* &= \max_{\mu, \lambda} \min_{\theta, \rho} \tilde{\mathcal{L}}(\rho, \theta, \mu, \lambda) \\ &= \max_{\mu, \lambda} \left[\min_{\rho} \tilde{\mathcal{L}}_1(\rho, \mu, \lambda) + \min_{\theta} \tilde{\mathcal{L}}_2(\theta, \mu, \lambda) \right]. \end{aligned} \quad (23)$$

where we utilized the decomposition of the Lagrangian to write the second equality.

The important observation to make is that the minimization over ρ of $\tilde{\mathcal{L}}_1(\rho, \mu, \lambda)$ has an elementary solution. Indeed, as per its definition we have

$$\min_{\rho} \tilde{\mathcal{L}}_1(\rho, \mu, \lambda) = \min_{\rho} \rho \left(1 - \int \lambda(z) p(z) dV(z) dz \right). \quad (24)$$

This minimization yields $-\infty$ when $\int \lambda(z) p(z) dV(z) dz \neq 1$ and 0 when $\int \lambda(z) p(z) dV(z) dz = 1$. Since in the dual problem we are interested in the maximum over all dual variables, we know that: (i) The maximum will be attained for a dual distribution that satisfies $\int \lambda(z) p(z) dV(z) dz = 1$. (ii) When dual variables satisfy this property we know that $\min_{\rho} \rho (1 - \int \lambda(z) p(z) dV(z) dz) = 0$. It then follows that the dual problem in (23) is equivalent to

$$\begin{aligned} \tilde{D}^* &= \max_{\mu, \lambda} \min_{\theta} \tilde{\mathcal{L}}_2(\theta, \mu, \lambda), \\ &\text{subject to } \int_{\mathcal{M}} \lambda(z) p(z) dV(z) = 1. \end{aligned} \quad (25)$$

This is the problem in (6) given the definition of $\tilde{\mathcal{L}}_2(\theta, \mu, \lambda)$ in (22) which is the same as the definition of $L(\theta, \mu, \lambda)$ in (6). ■

□

Notice that the constraint $\int_{\mathcal{M}} \lambda(z)p(z)dV(z) = 1$ implies that $\lambda(z)$ is a probability distribution over the manifold \mathcal{M} . This is an important observation for the connections we establish to manifold regularization in Section 3.2. It also implies that even though the dual problem 6 is not an unconstrained problem, the constraint is easy to enforce as an orthogonal projection on the space of probability distributions over the manifold \mathcal{M} . This is a simple normalization.

B EMPIRICAL DUAL PROBLEM WITH UNLABELED SAMPLES

In Section 3 we introduce the empirical dual program in (7) in the case in which unlabeled samples are unavailable. The ability to leverage unlabeled samples is an important feature of this work and ready to incorporate in (7). Indeed, if in addition to N i.i.d. labeled samples (x_n, y_n) with $N \in [1, N]$ drawn from $p(x, y)$ we are also given \tilde{N} i.i.d. unlabeled samples x_n with $n \in [N + 1, N + \tilde{N}]$ drawn from the input distribution $p(x)$ we redefine (7) as

$$\begin{aligned} \hat{D}^* = \max_{\hat{\mu}, \hat{\lambda} \geq 0} \min_{\theta} \hat{L}(\theta, \hat{\mu}, \hat{\lambda}) &:= \hat{\mu} \left(\frac{1}{N} \sum_{n=1}^N \ell(f_{\theta}(x_n), y_n) - \epsilon \right) + \frac{1}{N + \tilde{N}} \sum_{n=1}^{N + \tilde{N}} \hat{\lambda}(x_n) \|\nabla_{\mathcal{M}} f_{\theta}(x_n)\|^2, \\ \text{subject to } \frac{1}{N + \tilde{N}} \sum_{n=1}^{N + \tilde{N}} \hat{\lambda}(x_n) &= 1, \end{aligned} \quad (26)$$

Results in Section 3 hold with proper modifications.

C PROOF OF PROPOSITION 1

The proof in this appendix is a generalization of the proof in Chamon et al. (2021). In order to show Proposition 1, we need to introduce an auxiliary problem formulation over a functional domain of functions. In this case, we take the optimization problem (3) over the convex hull of the domain of parametric functions $\phi \in \tilde{\mathcal{H}}$,

$$\begin{aligned} \tilde{P}^* = \min_{\phi \in \tilde{\mathcal{H}}, \rho \geq 0} \quad & \rho, \\ \text{subject to} \quad & \mathbb{E}_{p(x, y)} [\ell(\phi(x), y)] \leq \epsilon, \\ & \|\nabla_{\mathcal{M}} \phi(z)\|^2 \leq \rho, \quad p(z)\text{-a.e.}, \quad z \in \mathcal{M}. \end{aligned} \quad (27)$$

We can now show that problem 27 is strongly dual as follows,

Lemma 1. *Under the assumptions of Proposition 1, the functional smooth learning problem (27) has zero duality gap, i.e. $\tilde{P}^* = \tilde{D}^*$.*

Proof Lemma 1: To begin with, the non-parametric space $\tilde{\mathcal{H}}$ is convex as it is the convex hull of the space of parametric functions \mathcal{H} , therefore the domain of the optimization problem is a convex set. Note that the objective is linear. Given that the gradient of ϕ is a linear function of ϕ , and that by taking the norm we preserve convexity, the constraint $\|\nabla_{\mathcal{M}} \phi(z)\|^2 \leq \rho$ is convex. By assumption 1, the loss ℓ is convex. Therefore, problem (3) is a semi-infinite convex problem. Moreover, as $\tilde{\theta}$ belongs to the relative interior of the feasible set, and $\mathcal{H} \subset \tilde{\mathcal{H}}$, it suffices to take $\phi^{\dagger}(\cdot) = f(\tilde{\theta}, \cdot)$, and $\rho^{\dagger} > R^2$, and by Slater's condition as $\phi^{\dagger}, \rho^{\dagger}$ belongs to the interior of the feasible domain, problem (3) has zero duality gap, i.e. it is strongly dual.

Now we need to define the *supergradient* of function $d(\mu, \lambda)$.

Definition 3 (Supergradient and Superdifferential). *We say that $c \in \mathbb{R}$ is a supergradient of d at μ if,*

$$d(\mu') \leq d(\mu) + c(\mu' - \mu) \text{ for all } \mu' \in \mathbb{R}. \quad (28)$$

The set of all supergradients of d at μ is called the superdifferential, and we denote it $\partial d(\mu)$.

Lemma 2 (Danskin's Theorem). *Consider the function,*

$$F(x) = \sup_{y \in Y} f(x, y). \quad (29)$$

where $f : \mathbb{R}^n \times Y \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$, if the following conditions are satisfied

1. Function $f(\cdot, y)$ is convex for all $y \in Y$.
2. Function $f(x, \cdot)$ is upper semicontinuous for all x in a certain neighborhood of a point x_0 .
3. The set $Y \subset \mathbb{R}^m$ is compact

Then,

$$\partial F(x_0) = \text{conv} \left(\bigcup_{y \in \hat{Y}(x_0)} \partial_x f(x_0, y) \right) \quad (30)$$

where $\partial_x f(x_0, y)$ denotes the subdifferential of the function $f(\cdot, y)$ at x_0

Proof of Lemma 2: The proof can be found in (Ruszczynski, 2011, Theorem 2.87).

We now need to define the primal problem 3 associated with Lagrangian $L(\theta, \mu, \lambda)$ as follows,

$$P^* = \min_{\theta \in \Theta} \max_{\mu, \lambda \geq 0} L(\theta, \mu, \lambda). \quad (31)$$

note that 3 and 31 are equivalent problems, given that in the optimal solution, $\int_{\mathcal{M}} \lambda(z) p(z) dV(z)$, to avoid an infinite result.

Lemma 3. *Let μ^*, λ^* be a solution of (31). Under the conditions of Proposition 1, there exists a feasible $\theta^\dagger \in \text{argmin}_\theta \mathcal{L}(\theta, \mu^*, \lambda^*)$, and the value D^* is bounded by,*

$$P^* - \nu(\mu_\nu^* M - 2P_{\phi_\nu}^* - \nu) \leq D^* \leq P^* \quad (32)$$

where $P_{\phi_\nu}^*$ and μ_ν^* are the optimal value, and optimal dual variable of the functional version of problem 3 with constraint $\epsilon - \nu M$.

Proof: First we need to show that there is a feasible θ^\dagger for problem (3). The constraint $\|\nabla_{\mathcal{M}} f_{\theta^\dagger}(x)\|^2 \leq \rho$ can be trivially satisfied by taking $\rho = R^2$, as $\|\nabla_{\mathcal{M}} f_\theta(x)\| \leq R$ by Assumption 2. However, we need to verify that there exists a θ^\dagger that is feasible. To do so, We begin by considering the set of Lagrange minimizers as follows,

$$\Theta^\dagger(\mu^*, \lambda^*) = \text{argmin}_\theta L(\theta, \mu^*, \lambda^*) \quad (33)$$

we can also define the constraint slack associated with parameters θ as follows,

$$c(\theta) = [\mathbb{E}[\ell(f_\theta(x), y)] - \epsilon]_+. \quad (34)$$

Therefore, to show that there exists a feasible solution, it is analogous to show that there is an element θ^\dagger of $\Theta^\dagger(\mu^*, \lambda^*)$ whose slack is equal to zero, i.e. $c(\theta^\dagger) = 0$. To do so, we leverage Lemma 2, given that condition (1) is satisfied by convexity of ℓ , and linearity of integral, condition (2) is satisfied by linearity of integral, condition (3) is satisfied by the compactness of Θ , and condition (4) is satisfied by the smoothness of Θ .

By contradiction, we can say that if no element of $\Theta(\mu^*, \lambda^*)$ is feasible, then $\mathbb{E}[\ell(f_\theta(x), y)] - \epsilon > 0$, for all $\theta \in \Theta(\mu^*, \lambda^*)$. From Lemma 2, we get that $\mathbf{0} \notin \partial \Theta^\dagger(\mu^*, \lambda^*)$, which contradicts the optimality of μ^*, λ^* . Given that the dual variable μ^* is finite, considering the existence of a feasible solution by Assumption 2. Hence, there must be one element of $\Theta(\mu^*, \lambda^*)$ that is feasible.

Now we need to show that the inequality 32 holds. The upper bound is trivially verified by weak duality Boyd et al. (2004), i.e.,

$$\check{D}^* \leq \check{P}^* \quad (35)$$

To show the lower bound, we consider the functional version of problem 3 with constraint $\epsilon - M\nu$ as follows,

$$\begin{aligned} \tilde{P}_\nu^* = & \min_{\phi \in \mathcal{H}, \rho \geq 0} \rho, \\ \text{subject to} \quad & \mathbb{E}_{p(x,y)}[\ell(\phi(x), y)] \leq \epsilon - M\nu, \\ & \|\nabla_{\mathcal{M}}\phi(z)\|^2 \leq \rho, \quad p(z)\text{-a.e.}, \quad z \in \mathcal{M}. \end{aligned} \quad (36)$$

Now note that problem 36 is strongly dual by Lemma 1, i.e.,

$$\tilde{P}_\nu^* = \min_{\phi} \max_{\mu, \lambda \geq 0} \tilde{L}_\nu(\phi, \mu, \lambda) = \max_{\mu, \lambda \geq 0} \min_{\phi} \tilde{L}_\nu(\phi, \mu, \lambda) = \tilde{D}_\nu^* \quad (37)$$

with the Lagrangian defined as,

$$\begin{aligned} \tilde{L}_\nu(\phi, \mu, \lambda) = & \mu \left(\mathbb{E}[\ell(\phi(x), y)] - (\epsilon - M\nu) \right) + \int_{\mathcal{M}} \lambda(x) \|\nabla_{\mathcal{M}}\phi(x)\|^2 p(x) dV(x) \\ \text{subject to} \quad & \int_{\mathcal{M}} \lambda(x) p(x) dV(x) = 1, \end{aligned} \quad (38)$$

we define the optimal dual variables μ_ν^*, λ_ν^* of \tilde{D}_ν^* be such that,

$$\tilde{P}_\nu^* = \min_{\phi} \tilde{L}_\nu(\phi, \mu_\nu^*, \lambda_\nu^*) \quad (39)$$

Coming back to the parametric dual problem (6), we know that,

$$D^* \geq \min_{\theta} L(\theta, \mu, \lambda), \text{ for all } \mu, \lambda \quad (40)$$

We thus utilize the optimal dual variables of the functional problem with constraints $\epsilon - M\nu$, i.e. μ_ν^*, λ_ν^* , as follows,

$$D^* \geq \min_{\theta, \rho} L(\theta, \mu_\nu^*, \lambda_\nu^*) \geq \min_{\phi} L(\phi, \mu_\nu^*, \lambda_\nu^*) \geq \min_{\phi} \tilde{L}_\nu(\phi, \mu_\nu^*, \lambda_\nu^*) - \mu_\nu^* M\nu = \tilde{P}_\nu^* - \mu_\nu^* M\nu. \quad (41)$$

Given that $\mathcal{H} \subseteq \bar{\mathcal{H}}$, and Problem 36 is strongly dual. To complete the proof we need to show that,

$$\tilde{P}_\nu^* \geq P^* - 2\nu P_{\phi_\nu^*}^* - \nu^2. \quad (42)$$

Denoting ϕ_ν^* the solution to problem 36, by Assumption 2, we know there is a parameterization $\tilde{\theta}_\nu^*$ such that simultaneously,

$$\sup_{z \in \mathcal{M}} |\phi_\nu^*(z) - f_{\tilde{\theta}_\nu^*}(z)| \leq \nu \quad (43)$$

$$\sup_{z \in \mathcal{M}} \|\nabla \phi_\nu^*(z) - \nabla f_{\tilde{\theta}_\nu^*}(z)\| \leq \nu \quad (44)$$

therefore,

$$\left| \mathbb{E}[\ell(\phi_\nu^*(x), y)] - \mathbb{E}[\ell(f_{\tilde{\theta}_\nu^*}(x), y)] \right| \leq \mathbb{E} \left[|\ell(\phi_\nu^*(x), y) - \ell(f_{\tilde{\theta}_\nu^*}(x), y)| \right] \quad (45)$$

$$\leq M \mathbb{E} \left[|\phi_\nu^*(x) - f_{\tilde{\theta}_\nu^*}(x)| \right] \leq M\nu \quad (46)$$

Since ϕ_ν^* is feasible for Problem 36, it implies that $f_{\tilde{\theta}_\nu^*}$ is feasible for Problem 3. We can now define the $\rho_{\tilde{\theta}_\nu^*}^*$ as the minimum ρ obtained with $f_{\tilde{\theta}_\nu^*}$ as follows,

$$\begin{aligned} \rho_{\tilde{\theta}_\nu^*}^* = & \min_{\rho \geq 0} \rho, \\ \text{subject to} \quad & \mathbb{E}_{p(x,y)}[\ell(f_{\tilde{\theta}_\nu^*}(x), y)] \leq \epsilon, \\ & \|\nabla_{\mathcal{M}} f_{\tilde{\theta}_\nu^*}(z)\|^2 \leq \rho, \quad p(z)\text{-a.e.}, \quad z \in \mathcal{M}. \end{aligned} \quad (47)$$

Returning to (41), and by optimality, $P^* \leq P_{\hat{\theta}_\nu^*}^*$,

$$D^* \geq P_{\phi_\nu^*}^* - \mu_\nu^* M \nu \quad (48)$$

$$\geq P_{\phi_\nu^*}^* + P^* - P_{\hat{\theta}_\nu^*}^* - \mu_\nu^* M \nu \quad (49)$$

Now we need to bound the difference $P_{\phi_\nu^*}^* - P_{\hat{\theta}_\nu^*}^*$. We know that there exist $z_1, z_2 \in \mathcal{M}$ such that $\|\nabla_{\mathcal{M}} \phi_\nu^*(z_1)\|^2 = P_{\phi_\nu^*}^*$ and $\|\nabla_{\mathcal{M}} f_{\hat{\theta}_\nu^*}(z_2)\|^2 = P_{\hat{\theta}_\nu^*}^*$, by optimality,

$$P_{\phi_\nu^*}^* - P_{\hat{\theta}_\nu^*}^* = \|\nabla_{\mathcal{M}} \phi_\nu^*(z_1)\|^2 - \|\nabla_{\mathcal{M}} f_{\hat{\theta}_\nu^*}(z_2)\|^2 \quad (50)$$

$$\geq \|\nabla_{\mathcal{M}} \phi_\nu^*(z_2)\|^2 - \|\nabla_{\mathcal{M}} f_{\hat{\theta}_\nu^*}(z_2)\|^2 \quad (51)$$

$$\geq \|\nabla_{\mathcal{M}} \phi_\nu^*(z_2)\|^2 - (\|\nabla_{\mathcal{M}} f_{\hat{\theta}_\nu^*}(z_2) - \nabla_{\mathcal{M}} \phi_\nu^*(z_2)\| + \|\nabla_{\mathcal{M}} \phi_\nu^*(z_2)\|) \quad (52)$$

$$(\|\nabla_{\mathcal{M}} f_{\hat{\theta}_\nu^*}(z_2) - \nabla_{\mathcal{M}} \phi_\nu^*(z_2)\| + \|\nabla_{\mathcal{M}} \phi_\nu^*(z_2)\|) \quad (53)$$

$$\geq -2\|\nabla_{\mathcal{M}} f_{\hat{\theta}_\nu^*}(z_2) - \nabla_{\mathcal{M}} \phi_\nu^*(z_2)\| \|\nabla_{\mathcal{M}} \phi_\nu^*(z_2)\| - \|\nabla_{\mathcal{M}} \phi_\nu^*(z_2)\|^2 \quad (54)$$

$$\geq -2\nu P_{\phi_\nu^*}^* - \nu^2 \quad (55)$$

where B is a bound on the norm of $\|\nabla_{\mathcal{M}} \phi(z)\|$. Putting (49) and (55) together, we attain the desired result.

Proof of Proposition 1:

This proof follows the lines of (Chamon et al., 2021, Proposition III.4). We begin by considering μ^*, λ^* , and $\hat{\mu}^*, \hat{\lambda}^*$, solutions of 6, and 7, and we define the set of optimal dual minimizers as,

$$\Theta(\mu^*, \lambda^*) = \operatorname{argmin}_{\theta \in \Theta} L(\theta, \mu^*, \lambda^*) \quad (56)$$

$$\hat{\Theta}(\hat{\mu}^*, \hat{\lambda}^*) = \operatorname{argmin}_{\theta \in \Theta} \hat{L}(\theta, \hat{\mu}^*, \hat{\lambda}^*) \quad (57)$$

where L , and \hat{L} are as defined in 6, and 7. We can proceed to bound the difference between the values of the dual problems as follows,

$$D^* - \hat{D}^* = \min_{\theta \in \Theta} L(\theta, \mu^*, \lambda^*) - \min_{\theta \in \Theta} \hat{L}(\theta, \hat{\mu}^*, \hat{\lambda}^*) \quad (58)$$

$$\leq \min_{\theta \in \Theta} L(\theta, \mu^*, \lambda^*) - \min_{\theta \in \Theta} \hat{L}(\theta, \mu^*, \bar{\lambda}^*) \quad (\text{by optimality}) \quad (59)$$

$$\leq L(\theta^\dagger, \mu^*, \lambda^*) - \hat{L}(\hat{\theta}^\dagger, \mu^*, \bar{\lambda}^*) \quad (60)$$

where $\bar{\lambda}^*(z) = \lambda^*(z) / \sum_{n=1}^N \lambda^*(z_n)$, and we define $\hat{\theta}^\dagger \in \hat{\Theta}(\mu^*, \bar{\lambda}^*)$. By utilizing the definition of the Lagrangian, we can obtain,

$$D^* - \hat{D}^* \leq |\mu^*| \left| \mathbb{E}[\ell(f_{\hat{\theta}^\dagger}(x), y)] - \frac{1}{N} \ell(f_{\hat{\theta}^\dagger}(x_n), y_n) \right| \quad (61)$$

$$+ \left| \mathbb{E}_{\lambda^*} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^\dagger}(x_n)\|^2 - \frac{1}{N} \sum_{n=1}^N \bar{\lambda}^*(x_n) \|\nabla_{\mathcal{M}} f_{\hat{\theta}^\dagger}\|^2 \right| \quad (62)$$

Note that the sampled dual variables, converge to the continuous values as follows,

$$\lim_{N \rightarrow \infty} \frac{\lambda^*(z)}{\frac{1}{N} \sum_{n=1}^N \lambda^*(z_n)} = \lambda^*(z) \quad (63)$$

Given that $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \lambda^*(z_n) = 1$ Utilizing the same argument on the other direction of the inequality it yields,

$$D^* - \hat{D}^* \geq L(\theta^\dagger, \hat{\mu}^*, \hat{\lambda}^*) - \hat{L}(\hat{\theta}^\dagger, \hat{\mu}^*, \hat{\lambda}^*) \quad (64)$$

$$\geq |\hat{\mu}^*| \left| \mathbb{E}[\ell(f_{\hat{\theta}^\dagger}(x), y)] - \frac{1}{N} \ell(f_{\hat{\theta}^\dagger}(x_n), y_n) \right| \quad (65)$$

$$+ \left| \mathbb{E}_{\bar{\lambda}^*} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^\dagger}(x_n)\|^2 - \frac{1}{N} \sum_{n=1}^N \lambda^*(x_n) \|\nabla_{\mathcal{M}} f_{\hat{\theta}^\dagger}\|^2 \right| \quad (66)$$

Where $\bar{\lambda}^* = \frac{1}{N} \sum_{n=1}^N \hat{\lambda}^*(x_n) B_\alpha(x_n)$, where $B_\alpha(x_n)$ is a ball of center x_n , and radius α . By tending $\alpha \rightarrow 0$, and $N \rightarrow \infty$, this integral converges, given the compactness of the manifold. Therefore, with probability $1 - 2\delta$, it holds,

$$\left| D^* - \hat{D}^* \right| \leq \max \left\{ |L(\theta^\dagger, \mu^*, \lambda^*) - \hat{L}(\hat{\theta}^\dagger, \mu^*, \bar{\lambda}^*)|, |L(\theta^\dagger, \hat{\mu}^*, \hat{\lambda}^*) - \hat{L}(\theta^\dagger, \hat{\mu}^*, \hat{\lambda}^*)| \right\} \quad (67)$$

Now we utilize the Lipschitzness of the gradient of $\nabla_{\mathcal{M}} f_\theta$ by assumption 2, and given that the set Θ is compact, and the manifold \mathcal{M} is compact,

$$\left| \|\nabla_{\mathcal{M}} f_{\theta_1}(x)\|^2 - \|\nabla_{\mathcal{M}} f_{\theta_2}(x)\|^2 \right| \quad (68)$$

$$\leq 2 \max_{\theta \in \Theta, z \in \mathcal{M}} \|\nabla_{\mathcal{M}} f_\theta(z)\| \|\nabla_{\mathcal{M}} f_{\theta_1}(x) - \nabla_{\mathcal{M}} f_{\theta_2}(x)\| \text{ by triangle inequality,} \quad (69)$$

$$\leq 2RG|f_{\theta_1}(x) - f_{\theta_2}(x)| \text{ by compactness and lipschitz.} \quad (70)$$

Where $\max_{\theta \in \Theta, z \in \mathcal{M}} \|\nabla_{\mathcal{M}} f_\theta(z)\| \leq R < \infty$ by compactness of Θ , and \mathcal{M} . To complete the proof, we leverage Talagrand's lemma (Mohri et al., 2018, Lemma 5.7), finally obtaining,

$$\left| D^* - \hat{D}^* \right| \leq \max \left\{ |\mu^*|, |\hat{\mu}^*| \right\} \zeta(N, \delta) + 2RG\hat{\zeta}(N, \delta) \quad (71)$$

By leveraging Lemma 3 we attain the desired result.

D PROOF OF PROPOSITION 2

Proof of Proposition 2:

To prove that $\hat{\theta}^*$ is approximately feasible, we use the same argument that in Lemma 3. By contradiction, if there exists no $\hat{\theta}^\dagger \in \hat{\Theta}(\hat{\mu}^*, \hat{\lambda}^*)$ that is feasible, the supergradient,

$$\frac{1}{N} \sum_{n=1}^N \ell(f_{\hat{\theta}^\dagger}(x_n), y_n) > \epsilon \quad (72)$$

and therefore $\mathbf{0} \notin \partial \hat{\Theta}(\hat{\theta}^*, \hat{\lambda}^*)$ which contradicts the optimality of $\hat{\theta}^*$, and $\hat{\lambda}^*$. Therefore, there exists $\hat{\theta}^* \in \hat{\Theta}(\hat{\mu}^*, \hat{\lambda}^*)$ such that,

$$\mathbb{E}[\ell(f_{\hat{\theta}^\dagger}(x), y)] \leq \epsilon + \zeta(N, \delta) \quad (73)$$

with probability $1 - \delta$. We can analyze the term given by the summation of the dual variables λ , noting that if there is no solution $\hat{\lambda}^*$, such that

$$\frac{1}{N} \sum_{n=1}^N \lambda(x_n) \|\nabla_{\mathcal{M}} f_{\theta}(x_n)\|^2 = \max_{n \in [N]} \|\nabla f_{\theta}(x_n)\|^2 \quad (74)$$

then, utilizing the same argument, $\mathbf{0} \notin \partial \hat{\Theta}(\hat{\theta}^*, \hat{\lambda}^*)$, contradicting the optimality of $\hat{\theta}^*$.

Now, we know that the norm of the gradient $\|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(x_n)\|$, is smaller than a value $\sqrt{\rho_{\hat{\theta}^*}}$ for all x_n . To bound the maximum gradient over all $z \in \mathcal{M}$, we can leverage the fact that the manifold is compact, and that the gradients are Lipschitz as follows,

$$\|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(x_n)\| \leq \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| + \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(x_n) - \nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| \text{ for all } z \in \mathcal{N}(x_n) \quad (75)$$

$$\leq \sup_{z \in \mathcal{N}(x_n)} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| + \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(x_n) - \nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| \quad (76)$$

$$\leq \sup_{z \in \mathcal{M}} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| + \sup_{z \in \mathcal{N}(x_n)} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(x_n) - \nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| \quad (77)$$

$$\leq \sup_{z \in \mathcal{M}} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| + \sup_{z \in \mathcal{N}(x_n)} d(x_n, z) \quad (78)$$

$$\leq \sup_{z \in \mathcal{M}} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\| + G\tilde{\zeta}(N) \quad (79)$$

where $\{\mathcal{N}(x_n) : x_n = \operatorname{argmin}_{n \in [N]} d_{\mathcal{M}}(x_n, z)\}$, which corresponds to the set of points in \mathcal{M} that are closer to x_n . Function $\tilde{\zeta}(N)$ is a decreasing function that measures the maximum distance between a point sampled by $z \in \mathcal{M}$, and a point in $\mathcal{N}(x_n)$, this number decreases with N given that \mathcal{M} is compact. Now, we can evaluate \hat{D}^* at $\hat{\theta}^*$ as follows,

$$\hat{D}^* = \mu^* \left(\sum_{n=1}^N \ell(f_{\hat{\theta}^*}(x_n), y_n) - \epsilon \right) + \frac{1}{N} \sum_{n=1}^N \hat{\lambda}^*(x_n) \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(x_n)\|^2 \quad (80)$$

$$= \mu^* \left(\sum_{n=1}^N \ell(f_{\hat{\theta}^*}(x_n), y_n) - \epsilon \right) + \max_{n \in [N]} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(x_n)\|^2 \text{ (by optimality of } \hat{\lambda}^*) \quad (81)$$

To conclude the proof, we leverage Proposition 1, in conjunction with (79) and we get that,

$$|P^* - \sup_{z \in \mathcal{M}} \|\nabla_{\mathcal{M}} f_{\hat{\theta}^*}(z)\|^2| \leq |\hat{\mu}^*| \left(\sum_{n=1}^N \ell(f_{\hat{\theta}^*}(x_n), y_n) - \epsilon \right) \quad (82)$$

$$+ G\tilde{\zeta}(N) + \mathcal{O}(\nu) + \max \left\{ |\mu^*|, |\hat{\mu}^*| \right\} \zeta(N, \delta) + 2GR\hat{\zeta}(N, \delta) \quad (83)$$

E PROOF OF PROPOSITION 3

The proof of Proposition 3 utilizes the following Definitions.

Definition 4. The total volume of the smooth Manifold \mathcal{M} is \mathbf{V} , i.e.,

$$\int_{\mathcal{M}} \text{vol}(x) = \mathbf{V}. \quad (84)$$

Definition 5. The function f_θ , parameterized by $\theta \in \Theta$, is universally upper bounded by \mathbf{F} , i.e.

$$\max_{x \in \mathcal{M}, \theta \in \Theta} \|f_\theta(x)\| \leq \mathbf{F} \quad (85)$$

Definition 6. Given a probability distribution $\lambda \in \Lambda$ defined over the manifold \mathcal{M} , we define the degree $d_t(x)$ as,

$$d_t(p) = \int_{\mathcal{M}} G_t(p, z) \lambda(z) \text{vol}(z) \quad (86)$$

Definition 7. We define the continuous heat kernel laplacian $\tilde{\mathbf{L}}_\lambda^t f_\theta(p)$ associated with dual function $\lambda \in \Lambda$, function $f_\theta, \theta \in \Theta$ and point $p \in \mathcal{M}$ as,

$$\tilde{\mathbf{L}}_\lambda^t f_\theta(p) = \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(p, z)}{\sqrt{d_t(p)} \sqrt{d_t(z)}} (f_\theta(p) - f_\theta(z)) \lambda(z) \text{vol}(z) \quad (87)$$

where d_t is the degree function defined in 6.

Lemma 4. For a given $p \in \mathcal{M}$, and any open set $\mathcal{B} \subset \mathcal{M}, p \in \mathcal{B}$, for any function $f_\theta, \theta \in \Theta$, such that $\sup_{\theta \in \Theta, x \in \mathcal{M}} |f(x)| \leq \mathbf{F}$, and any $\lambda \in \Lambda$ with Λ as in Proposition 3, we have that

$$\left| \int_{\mathcal{B}} e^{-\frac{\|p-z\|^2}{4t}} \lambda(z) f_\theta(z) \text{vol}(z) - \int_{\mathcal{M}} e^{-\frac{\|p-z\|^2}{4t}} \lambda(z) f_\theta(z) \text{vol}(z) \right| \leq b \mathbf{V} \mathbf{F} e^{-\frac{r^2}{4t}}, \quad (88)$$

where $r = \inf_{x \notin \mathcal{B}} \|x - p\|$, and b in Λ .

Proof of Lemma 4: The proof is similar to Lemma 4.1 in Belkin and Niyogi (2005). The proof follows from bounding λ by b from assumption, $\|f_\theta\|$ by \mathbf{F} from definition 5, the volume of $\mathcal{M} - \mathcal{B}$ by the volume of \mathcal{M} given by \mathbf{V} , and bounding $e^{-\frac{\|p-z\|^2}{4t}}$ by $e^{-\frac{d}{4t}}$. ■

Lemma 5. The limit of the degree function $d_t(x)$ (defined in Definition 6) when $t \rightarrow 0$ is lower bounded for probabilities $\lambda \in \Lambda$ as in Proposition 3, i.e.,

$$\min_{p \in \mathcal{M}, \lambda \in \Lambda} \lim_{t \rightarrow 0} d_t(x) = a, \quad (89)$$

with $a = \min_{z \in \mathcal{M}} \lambda(z)$ as given by $\lambda \in \Lambda$.

Proof of Lemma 5: To begin with, we can write down the definition of the degree function $d_t(x)$,

$$\min_{p \in \mathcal{M}, \lambda \in \Lambda} \lim_{t \rightarrow 0} d_t(x) = \min_{p \in \mathcal{M}, \lambda \in \Lambda} \lim_{t \rightarrow 0} \int_{\mathcal{M}} G_t(x, z) \lambda(z) \text{vol}(z) \quad (90)$$

$$= \min_{p \in \mathcal{M}, \lambda \in \Lambda} \lim_{t \rightarrow 0} \int_{\mathcal{M}} \frac{1}{(4\pi t)^{d/2}} e^{-\frac{\|x-z\|^2}{4t}} \lambda(z) \text{vol}(z) \quad (91)$$

$$= \min_{p \in \mathcal{M}} \lim_{t \rightarrow 0} a \int_{\mathcal{M}} \frac{1}{(4\pi t)^{d/2}} e^{-\frac{\|x-z\|^2}{4t}} \text{vol}(z) \quad (92)$$

$$= \min_{p \in \mathcal{M}} \lim_{t \rightarrow 0} a \int_{\mathcal{M}} \frac{1}{(4\pi t)^{d/2}} e^{-\frac{(x-z)^\top (\frac{1}{2t} \mathbf{I})(x-z)}{2}} \text{vol}(z) \quad (93)$$

$$= \min_{p \in \mathcal{M}} \lim_{t \rightarrow 0} a \int_{\mathcal{B}} \frac{1}{(4\pi t)^{d/2}} e^{-\frac{(x-z)^\top (\frac{1}{2t} \mathbf{I})(x-z)}{2}} \text{vol}(z) \quad (94)$$

$$= \min_{p \in \mathcal{M}} \lim_{t \rightarrow 0} a \int_{\mathcal{B}} \frac{1}{(4\pi t)^{d/2}} e^{-\frac{(v)^\top (\frac{1}{2t} \mathbf{I})(v)}{2}} \text{vol}(z) \quad (95)$$

$$= \frac{a}{(4\pi t)^{d/2}} ((2\pi)^d (2t)^d)^{\frac{1}{2}} = a \quad (96)$$

Lemma 6. For any two sufficiently close points $p, q \in \mathcal{M}$, such that $q = \exp_p(v) = \mathbb{R}^k$, the relationship between the Euclidean distance and geodesic distance is given by,

$$d_{\mathcal{M}}^2(p, q) = \|v\|_{\mathbb{R}^k}^2 \quad (97)$$

Proof of Lemma 6: Consider a geodesic curve γ , that goes from $\gamma(0) = p$ to $\gamma(1) = q$, this geodesic can be obtained by $\gamma(t) = \exp_p(tv)$. The distance between p, q can be expressed as,

$$d_{\mathcal{M}}(p, q) = \int_0^1 \|\gamma'(t)\| dt = \|v\|. \quad (98)$$

given that the geodesic has constant derivative, and it is given by v by the definition of exponential map (see (Do Carmo and Flaherty Francis, 1992, Proposition 3.6)). ■

Lemma 7. For any point $p \in \mathcal{M}$, the probability of the difference between the point cloud Laplacian operator $\mathbf{L}_{\lambda, N}^t f_{\theta}(p)$ defined in equation 2 and the continuous heat kernel laplacian is given by

$$P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \mathbf{L}_{\lambda, N} f_{\theta}(p)| \geq \epsilon) \leq 2e^{-\frac{\epsilon^2 n}{8\mathbf{B}_t^2}} + 2e^{-\frac{\epsilon^2(n-1)}{4\mathbf{B}_t a^{3/2}}}, \quad (99)$$

where B_t is an upper bound on the norm of the random variable being integrated given by $\frac{1}{t} \frac{1}{(4\pi t)^{k/2}} \frac{2\mathbf{F}}{a}$.

Proof of Lemma 7: The proof is similar to section 5.1 of Theorem 5.2 from Belkin and Niyogi (2005). To begin with, we define an intermediate operator $\hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)$ as,

$$\hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) = \frac{1}{nt} \sum_{i=1}^N \frac{G_t(p, x_i)}{\sqrt{d_t(p)} \sqrt{d_t(x_i)}} (f_{\theta}(p) - f_{\theta}(x_i)) \quad (100)$$

Now we consider equation (99), adding and subtracting the intermediate operator $\hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)$ as follows,

$$P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \mathbf{L}_{\lambda, N} f_{\theta}(p)| \geq \epsilon) \quad (101)$$

$$= P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) + \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) - \mathbf{L}_{\lambda, N} f_{\theta}(p)| \geq \epsilon) \quad (102)$$

$$= 1 - P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) + \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) - \mathbf{L}_{\lambda, N} f_{\theta}(p)| \leq \epsilon) \quad (103)$$

$$= 1 - P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)| \leq \epsilon/2 \cap |\hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) - \mathbf{L}_{\lambda, N} f_{\theta}(p)| \leq \epsilon/2) \quad (104)$$

$$= 1 - (1 - P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)| > \epsilon/2 \cup P(|\hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) - \mathbf{L}_{\lambda, N} f_{\theta}(p)| > \epsilon/2)) \quad (105)$$

$$\leq P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)| > \epsilon/2) + P(|\hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p) - \mathbf{L}_{\lambda, N} f_{\theta}(p)| > \epsilon/2) \quad (106)$$

where equation (103) holds by taking the complement, equation (104) holds given that it is a subset of the total probability, equation (105) by taking complement again, and equation (106) by the union bound.

We will now focus on the first term of (106), i.e. $P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)| > \epsilon/2)$, and by considering the event as follows,

$$|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)| \quad (107)$$

$$= \left| \int_{\mathcal{M}} \frac{1}{t} \frac{G_t(p, z)}{\sqrt{d_t(p)} \sqrt{d_t(z)}} (f_{\theta}(p) - f_{\theta}(z)) \lambda(z) \text{vol}(z) - \frac{1}{tn} \sum_{i=1}^N \frac{G_t(p, x_i)}{\sqrt{d_t(p)} \sqrt{d_t(x_i)}} (f_{\theta}(p) - f_{\theta}(x_i)) \right|$$

Now we can note that for every $t > 0$, we can bound,

$$\left| \frac{1}{t} \frac{G_t(p, z)}{\sqrt{d_t(p)} \sqrt{d_t(z)}} (f_{\theta}(p) - f_{\theta}(z)) \right| \leq \frac{1}{t} \frac{1}{(4\pi t)^{k/2}} \frac{2\mathbf{F}}{a} = \mathbf{B}_t \quad (108)$$

where a as in Λ , and the degree is lower bounded by Lemma 5. By the application of Hoeffding's inequality, we obtain,

$$P(|\tilde{\mathbf{L}}_{\lambda}^t f_{\theta}(p) - \hat{\mathbf{L}}_{\lambda, n} f_{\theta}(p)| > \epsilon/2) \leq 2e^{-\frac{\epsilon^2 n}{8\mathbf{B}_t^2}} \quad (109)$$

We will now focus on the second term of (106), i.e. $P(|\hat{\mathbf{L}}_{\lambda,N}^t f_\theta(p) - \mathbf{L}_{\lambda,N} f_\theta(p)| > \epsilon/2)$, and by considering the event as follows,

$$\begin{aligned}
& |\hat{\mathbf{L}}_{\lambda,N}^t f_\theta(p) - \mathbf{L}_{\lambda,N} f_\theta(p)| \tag{110} \\
&= \left| \frac{1}{tn} \sum_{i=1}^N \frac{G_t(p, x_i)}{\sqrt{d_t(p)}\sqrt{d_t(x_i)}} (f_\theta(p) - f_\theta(x_i)) - \frac{1}{tn} \sum_{i=1}^N \frac{G_t(p, x_i)}{\sqrt{\hat{d}_t(p)}\sqrt{\hat{d}_t(x_i)}} (f_\theta(p) - f_\theta(x_i)) \right| \\
&\leq \frac{1}{tn} \frac{2\mathbf{F}}{(4\pi t)^{k/2}} \sum_{i=1}^N \left| \frac{1}{\sqrt{d_t(p)}\sqrt{d_t(x_i)}} - \frac{1}{\sqrt{\hat{d}_t(p)}\sqrt{\hat{d}_t(x_i)}} \right| \\
&\leq \frac{1}{tn} \frac{2\mathbf{F}}{(4\pi t)^{k/2}} \sum_{i=1}^N \left| \frac{\sqrt{d_t(p)}\sqrt{d_t(x_i)} - \sqrt{\hat{d}_t(p)}\sqrt{\hat{d}_t(x_i)}}{\sqrt{d_t(p)}\sqrt{d_t(x_i)}\sqrt{\hat{d}_t(p)}\sqrt{\hat{d}_t(x_i)}} \right| \\
&\leq \frac{1}{tn} \frac{2\mathbf{F}}{(4\pi t)^{k/2}} \sum_{i=1}^N \left| \frac{\sqrt{d_t(p)}\sqrt{d_t(x_i)} - \sqrt{\hat{d}_t(p)}\sqrt{\hat{d}_t(x_i)} + \sqrt{d_t(p)}\sqrt{\hat{d}_t(x_i)} - \sqrt{d_t(p)}\sqrt{\hat{d}_t(x_i)}}{\sqrt{d_t(p)}\sqrt{d_t(x_i)}\sqrt{\hat{d}_t(p)}\sqrt{\hat{d}_t(x_i)}} \right| \\
&\leq \frac{1}{tn} \frac{2\mathbf{F}}{(4\pi t)^{k/2}} \sum_{i=1}^N \left| \frac{\sqrt{d_t(p)}(\sqrt{\hat{d}_t(x_i)} - \sqrt{d_t(x_i)}) + \sqrt{\hat{d}_t(x_i)}(\sqrt{d_t(p)} - \sqrt{\hat{d}_t(p)})}{\sqrt{d_t(p)}\sqrt{d_t(x_i)}\sqrt{\hat{d}_t(p)}\sqrt{\hat{d}_t(x_i)}} \right| \tag{111}
\end{aligned}$$

Where we can use the Hoeffding's inequality given that the random variables are bounded, as follows,

$$P[|\hat{d}_t(x_i) - d_t(x_i)| > \epsilon] \leq 2e^{-\frac{\epsilon^2(n-1)}{8\mathbf{B}_t}} \tag{112}$$

Finally, we obtain,

$$P[|\hat{\mathbf{L}}_{\lambda,N}^t f_\theta(p) - \mathbf{L}_{\lambda,N} f_\theta(p)| > \epsilon/2] \leq 2e^{-\frac{\epsilon^2(n-1)}{4\mathbf{B}_t a^{3/2}}} \tag{113}$$

All together, we obtain,

$$P(|\tilde{\mathbf{L}}_\lambda^t f_\theta(p) - \mathbf{L}_{\lambda,N} f_\theta(p)| \geq \epsilon) \tag{114}$$

$$\leq P(|\tilde{\mathbf{L}}_\lambda^t f_\theta(p) - \hat{\mathbf{L}}_{\lambda,n} f_\theta(p)| > \epsilon/2) + P(|\tilde{\mathbf{L}}_\lambda^t f_\theta(p) - \hat{\mathbf{L}}_{\lambda,n} f_\theta(p)| > \epsilon/2) \tag{115}$$

$$\leq 2e^{-\frac{\epsilon^2 n}{8\mathbf{B}_t^2}} + 2e^{-\frac{\epsilon^2(n-1)}{4\mathbf{B}_t a^{3/2}}} \tag{116}$$

Completing the proof.

Lemma 8. Consider f_θ with $\theta \in \Theta$, and $\lambda \in \Lambda$, for any point $p \in \mathcal{M}$, there is a uniform bound between the continuous heat kernel laplacian (cf. definition 7) and the laplace-beltrami operator $\Delta_\lambda f_\theta = \Delta f_\theta + \langle \nabla_{\mathcal{M}} f_\theta, \nabla_{\mathcal{M}} \lambda \rangle$, where Δ is the Laplace-Beltrami operator, i.e.

$$|\tilde{\mathbf{L}}_\lambda^t f_\theta(p) - \Delta_\lambda f_\theta(p)| \leq \mathcal{O}(t^{1/2}) \tag{117}$$

with $\tilde{\mathbf{L}}_\lambda^t f_\theta(p) = \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(p, z)}{\sqrt{d_t(p)}\sqrt{d_t(z)}} (f_\theta(p) - f_\theta(z)) \lambda(z) \text{vol}(z)$.

Proof of Lemma 8: The proof is as follows, first we will compare the heat kernel laplacian $\tilde{\mathbf{L}}_\lambda^t f_\theta(p)$ with the heat kernel laplacian with λ instead of d_t . Next, we will consider the integral over a ball \mathcal{B} , as opposed to the integral over the whole manifold \mathcal{M} . Then, we will exploit the euclidean properties of the manifold at point p , which will allow us to convert an integral over a manifold into an integral over the low dimensional structure of the manifold. In this low dimensional manifold, we can compute the integrals, and obtain the desired result.

To begin with, we will introduce the continuous heat kernel laplacian, but using the probability distribution λ as follows,

$$\tilde{\mathbf{L}}_\lambda^t f_\theta(p) = \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(p, z)}{\sqrt{d_t(p)}\sqrt{d_t(z)}} (f_\theta(p) - f_\theta(z)) \lambda(z) \text{vol}(z) \quad (118)$$

$$\leq \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(p, z)}{\sqrt{\lambda(p)}\sqrt{\lambda(z)}} (f_\theta(p) - f_\theta(z)) \lambda(z) \text{vol}(z) \quad (119)$$

$$+ \left| \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(p, z)}{\sqrt{d_t(p)}\sqrt{d_t(z)}} (f_\theta(p) - f_\theta(z)) \lambda(z) \text{vol}(z) - \right. \quad (120)$$

$$\left. \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(p, z)}{\sqrt{\lambda(p)}\sqrt{\lambda(z)}} (f_\theta(p) - f_\theta(z)) \lambda(z) \text{vol}(z) \right| \quad (121)$$

$$\leq \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(p, z)}{\sqrt{\lambda(p)}\sqrt{\lambda(z)}} (f_\theta(p) - f_\theta(z)) \lambda(z) \text{vol}(z) \quad (122)$$

$$+ \frac{1}{t} \int_{\mathcal{M}} G_t(p, z) \left| \frac{1}{\sqrt{d_t(p)}\sqrt{d_t(z)}} - \frac{1}{\sqrt{\lambda(p)}\sqrt{\lambda(z)}} \right| |f_\theta(p) - f_\theta(z)| \lambda(z) \text{vol}(z) \quad (123)$$

By compactness of manifold \mathcal{M} , and the fact that $\lim_{t \rightarrow 0} d_t(p) = \lambda(p)$ the following holds for any point $p \in \mathcal{M}$,

$$d_t(p) = \lambda(p) + \mathcal{O}(tg(p)) \quad (124)$$

where g is a smooth function depending upon higher derivatives of λ , which is bounded by c . Therefore, we can write

$$d_t^{-1/2}(z) = (\lambda(z) + \mathcal{O}(tg(z)))^{-1/2} = \lambda(z)^{-1/2} + \mathcal{O}(t), \quad (125)$$

where $\mathcal{O}(t) \leq c|t|$. Therefore, we can bound the right hand side of 123 by,

$$\frac{1}{t} \int_{\mathcal{M}} G_t(p, z) \left| \frac{1}{\sqrt{d_t(p)}\sqrt{d_t(z)}} - \frac{1}{\sqrt{\lambda(p)}\sqrt{\lambda(z)}} \right| |f_\theta(p) - f_\theta(z)| \lambda(z) \text{vol}(z) \quad (126)$$

$$\leq \frac{1}{t} \int_{\mathcal{M}} G_t(p, z) ct |f_\theta(p) - f_\theta(z)| \lambda(z) \text{vol}(z) \quad (127)$$

$$\leq \int_{\mathcal{M}} G_t(p, z) c |f_\theta(p) - f_\theta(z)| \lambda(z) \text{vol}(z) \leq \mathcal{O}(t) \quad (128)$$

where the bound holds uniformly given the compactness of the manifold \mathcal{M} , and the upper bound on the derivative of f_θ . We now return to equation (123), by virtue of Lemma 4, we can convert the integral over the manifold \mathcal{M} , to an integral over the ball \mathcal{B} as follows,

$$\frac{1}{t} \int_{\mathcal{M}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \quad (129)$$

$$= \frac{1}{t} \int_{\mathcal{B}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \quad (130)$$

$$+ \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \quad (131)$$

$$- \frac{1}{t} \int_{\mathcal{B}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \quad (132)$$

$$\leq \frac{1}{t} \int_{\mathcal{B}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \quad (133)$$

$$+ \left| \frac{1}{t} \int_{\mathcal{M}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \right. \quad (134)$$

$$\left. - \frac{1}{t} \int_{\mathcal{B}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \right| \quad (135)$$

$$\leq \frac{1}{t} \int_{\mathcal{B}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) + b \mathbf{V} \mathbf{F} e^{-\frac{r^2}{4t}} \quad (136)$$

We now consider the exponential coordinates \exp around p , after introducing a ball \mathcal{B} of radius r . Letting the change of variables be $\exp_p : T_p \mathcal{M} \rightarrow \mathcal{M}$. Letting $\tilde{\mathcal{B}}$ be a ball in $T_p \mathcal{M}$, then \mathcal{B} is the image of the ball under the exponential map. We will use the exponential coordinates as $v = \exp_p(z)$, and $\tilde{f}_\theta(v) = f_\theta(\exp_p(v))$, and $\tilde{f}_\theta(0) = f_\theta(\exp_p(0))$.

$$\frac{1}{t} \int_{\mathcal{B}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) = \quad (137)$$

$$\frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \int_{\tilde{\mathcal{B}}} \frac{e^{\frac{\|\exp_p(v) - \exp_p(0)\|^2}{4t}}}{(4\pi t)^{k/2}} \sqrt{\tilde{\lambda}(v)} (\tilde{f}_\theta(0) - \tilde{f}_\theta(v)) \det(g(v)) dv \quad (138)$$

where $\det(g(v))$ is the determinant of the metric tensor in exponential coordinates. We can now expand $\sqrt{\tilde{\lambda}(v)}$, and $\tilde{f}_\theta(v)$ by their Taylor expansions as follows,

$$\tilde{\lambda}^{1/2}(v) \leq \tilde{\lambda}^{1/2}(0) + \frac{1}{2} \tilde{\lambda}^{-1/2}(0) v^\top \nabla \tilde{\lambda}(0) + \mathcal{O}(\|v\|^2) \quad (139)$$

$$\tilde{f}_\theta(v) = \tilde{f}_\theta(0) + v^\top \nabla \tilde{f}_\theta(0) + \frac{1}{2} v^\top \mathbf{H} v + \mathcal{O}(\|v\|^3) \quad (140)$$

$$\tilde{f}_\theta(0) - \tilde{f}_\theta(v) \leq -v^\top \nabla \tilde{f}_\theta(0) - \frac{1}{2} v^\top \mathbf{H} v + \mathcal{O}(\|v\|^3) \quad (141)$$

Where $\nabla \tilde{f} = (\frac{\partial \tilde{f}}{\partial x_1}, \dots, \frac{\partial \tilde{f}}{\partial x_k})^\top$, and \mathbf{H} its corresponding hessian matrix. Now note that given that the norm of the second derivative λ and third derivative of f_θ are uniformly bounded for all $p \in \mathcal{M}$, $\theta \in \Theta$ and $\lambda \in \Lambda$. With these two approximations in hand, we also approximate the heat kernel $e^{\frac{\|\exp_p(v) - \exp_p(0)\|^2}{4t}}$ as follows,

$$e^{-\frac{d_{\mathcal{M}}(\exp_p(v) - p)^2}{4t}} = e^{-\frac{\|v\|^2}{4t}}, \quad (142)$$

given that y is close enough to p . We also approximate the determinant of the metric tensor in exponential coordinates $\det(g(v))$ using the fact that

$$\det(g(v)) = 1 - \frac{1}{6} z^\top \mathbf{R} v + \mathcal{O}(\|v\|^3) \quad (143)$$

$$\leq 1 + k_0 \|v\|^2. \quad (144)$$

where \mathbf{R} is the Ricci curvature tensor which is uniformly bounded given the compactness of \mathcal{M} . The last inequality holds uniformly for all $p \in \mathcal{M}$ given the compactness of \mathcal{M} . Returning to 138, putting all the pieces together we obtain,

$$\frac{1}{t} \int_{\mathcal{B}} \frac{G_t(x, y)}{\sqrt{\lambda(x)}\sqrt{\lambda(y)}} (f_\theta(x) - f_\theta(y)) \lambda(y) \text{vol}(y) \quad (145)$$

$$\leq \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} (\tilde{\lambda}^{1/2}(0) + \frac{1}{2} \tilde{\lambda}^{-1/2}(0) v^\top \nabla \tilde{\lambda}(v) + k_1 \|v\|^2) \quad (146)$$

$$(v^\top \nabla \tilde{f}_\theta(0) + \frac{1}{2} v^t \mathbf{H} v + k_2 \|v\|^3) dv \quad (147)$$

$$+ \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} (\tilde{f}(0) - \tilde{f}(v)) k_0 \|v\|^2 dv \quad (148)$$

Which can be split into 5 integrals as follows:

$$A_t = \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} \tilde{\lambda}^{1/2}(0) v^\top \nabla \tilde{f}_\theta(0) dv \quad (149)$$

$$B_t = \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} \tilde{\lambda}^{1/2}(0) \frac{1}{2} v^t \mathbf{H} v dv \quad (150)$$

$$C_t = \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} \frac{1}{2} \tilde{\lambda}^{-1/2}(0) v^\top \nabla \tilde{\lambda}(0) \nabla \tilde{f}_\theta(0)^\top v dv \quad (151)$$

$$D_t = \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} \frac{1}{2} \tilde{\lambda}^{-1/2}(0) v^\top \nabla \tilde{\lambda}(0) \frac{1}{2} v^\top \mathbf{H} v dv \quad (152)$$

$$E_t = \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} (\tilde{f}(0) - \tilde{f}(v)) k_0 \|v\|^2 d\mathbb{R}^k \quad (153)$$

$$\leq \frac{1}{t} \frac{1}{\sqrt{\tilde{\lambda}(0)}} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} e^{-\frac{\|v\|^2}{4t}} k_0 \|v\|^3 d\mathbb{R}^k \quad (154)$$

Where for E_t we have used the fact that the first derivative has a uniform bound. Note that given the Gaussian kernel integration, the following conditions are true,

$$\int_{\tilde{\mathcal{B}}} v_i e^{-\frac{\|v\|^2}{4t}} dv = 0 \quad (155)$$

$$\int_{\tilde{\mathcal{B}}} v_i v_j e^{-\frac{\|v\|^2}{4t}} dv = 0 \quad (156)$$

$$\int_{\tilde{\mathcal{B}}} v_i v_j^2 e^{-\frac{\|v\|^2}{4t}} dv = 0 \quad (157)$$

$$\frac{1}{t} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} \|v\|^3 e^{-\frac{\|v\|^2}{4t}} dv = \mathcal{O}(t^{1/2}) \quad (158)$$

$$\frac{1}{t} \frac{1}{(4\pi t)^{k/2}} \int_{\tilde{\mathcal{B}}} v_i^2 e^{-\frac{\|v\|^2}{4t}} dv = 2 \quad (159)$$

for $i \neq j$ given the zero mean gaussian distribution with zero mean and diagonal covariance matrix. Using these conditions, we can conclude that $A_t = D_t = 0$. From B_t and C_t , the only non-zero elements are given by the diagonal elements, therefore, we obtain,

$$A_t + B_t + C_t + D_t + E_t = - \sum_{i=1}^k \frac{\partial^2 \tilde{f}(0)}{\partial x_i^2} - \frac{1}{\tilde{\lambda}(0)} \sum_{i=1}^k [\nabla \tilde{\lambda}(0)]_i [\nabla \tilde{f}_\theta(0)]_i + \mathcal{O}(t^{1/2}) \quad (160)$$

$$= -\Delta f - \frac{1}{\lambda} \langle \nabla \lambda, \nabla f_\theta \rangle_{T_p} + \mathcal{O}(t^{1/2}) \quad (161)$$

Where we have used that $\Delta_{\mathcal{M}} f(p) = \Delta_{\mathbb{R}^k} \tilde{f}(0) = -\sum_{i=1}^k \frac{\partial^2 \tilde{f}}{\partial x_i^2}(0)$ (see Chapter 3 of Rosenberg (1997)).

Proof of Proposition 3: To begin with, utilizing Green's identity, for the laplacian $\Delta_\lambda = \Delta f_\theta + \frac{1}{\lambda} \langle \nabla \lambda, \nabla f_\theta \rangle$, the following holds

$$\int_{\mathcal{M}} f_\theta(z) \Delta_\lambda f_\theta(z) \lambda(z) dV(z) = \int_{\mathcal{M}} \|\nabla_{\mathcal{M}} f_\theta(z)\|^2 \lambda(z) dV(z) \quad (162)$$

An equivalent statement of the proof is that for any $\epsilon > 0$, and $\delta > 0$, there exist N_0 such that, for $N \geq N_0$

$$P\left(\sup_{\lambda \in \Lambda, \theta \in \Theta} \left| \frac{1}{N} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) \mathbf{L}_N^{t_N} f_\theta(x_i) - \int_{\mathcal{M}} f_\theta(z) (-\Delta_1 f_\theta(z)) \lambda(x) p(z) dV(z) \right| > \epsilon \right) \leq \delta \quad (163)$$

By considering the complement of the event in (163), we obtain,

$$P\left(\sup_{\lambda \in \Lambda, \theta \in \Theta} \left| \frac{1}{N} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) \mathbf{L}_N^{t_N} f_\theta(x_i) - \int_{\mathcal{M}} f_\theta(z) (-\Delta_1 f_\theta(z)) \lambda(x) p(z) dV(z) \right| \leq \epsilon \right) \geq 1 - \delta \quad (164)$$

We can now add and subtract $\int_{\mathcal{M}} f_\theta(z) \mathbf{L}_N^{t_N}(z) \lambda(z) p(z) dV(z)$,

$$\begin{aligned} P\left(\sup_{\lambda \in \Lambda, \theta \in \Theta} \left| \frac{1}{n} \sum_{i=1}^n f_\theta(x_i) \lambda(x_i) \mathbf{L}_N^{t_N} f_\theta(x_i) - \int_{\mathcal{M}} f_\theta(z) \mathbf{L}_N^{t_N}(z) \lambda(z) p(z) dV(z) \right. \right. \\ \left. \left. + \int_{\mathcal{M}} f_\theta(z) \mathbf{L}_N^{t_N}(z) \lambda(z) p(z) dV(z) - \int_{\mathcal{M}} f_\theta(z) (-\Delta_1 f_\theta(z)) \lambda(x) p(z) dV(z) \right| \leq \epsilon \right) \\ \geq 1 - \delta \end{aligned} \quad (165)$$

By rearranging we obtain,

$$\begin{aligned} P\left(\sup_{\lambda \in \Lambda, \theta \in \Theta} \left| \frac{1}{N} \sum_{i=1}^n \mathbf{L}_N^{t_N} f_\theta(x_i) f_\theta(x_i) \lambda(x_i) - \int_{\mathcal{M}} f_\theta(z) \mathbf{L}_N^{t_N}(z) \lambda(z) p(z) dV(z) \right. \right. \\ \left. \left. + \int_{\mathcal{M}} f_\theta(z) \left(\mathbf{L}_N^{t_N}(z) - (-\Delta_1 f_\theta(z)) \right) \lambda(x) p(z) dV(z) \right| \leq \epsilon \right) \\ \geq 1 - \delta \end{aligned} \quad (166)$$

We denote C_N the event that,

$$\begin{aligned} C_N = \left\{ |\mathbf{L}_N^{t_N} f_\theta(x_i)| |\mathbf{L}_N^{t_N}(z) - (-\Delta_1 f_\theta(z)) f_\theta(z)| \leq \frac{\epsilon}{2} \right. \\ \left. \cap \left| \mathbf{L}_N^{t_N}(z) \right| \left| \frac{1}{N} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) - \int_{\mathcal{M}} f_\theta(z) \lambda(z) p(z) dV(z) \right| \leq \frac{\epsilon}{2} \right\} \end{aligned} \quad (167)$$

Given that we are imposing fixed conditions, C_N is included in the event described in (164),

$$C_N \subseteq \left\{ \left| \frac{1}{N} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) \mathbf{L}_N^{t_N} f_\theta(x_i) - \int_{\mathcal{M}} f_\theta(z) (-\Delta_1 f_\theta(z)) \lambda(x) p(z) dV(z) \right| \leq \epsilon \right\}. \quad (168)$$

Therefore, the probability satisfies,

$$P\left(\left| \frac{1}{N} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) \mathbf{L}_N^{t_N} f_\theta(x_i) - \int_{\mathcal{M}} f_\theta(z) \Delta_1 f_\theta(z) \lambda(x) p(z) dV(z) \right| \leq \epsilon \right) \geq P(C_N) \quad (169)$$

Now we will use the fact that the manifold is compact, and that f_θ is continuous, and denote $\max_{z \in \mathcal{M}} |f_\theta(z)| = F$, we can obtain the following bounds:

$$|\mathbf{L}_N^{t_N} f_\theta(x_i)| \leq 2F \quad (170)$$

$$\left| \int_{\mathcal{M}} f_\theta(z) \lambda(x) p(z) dV(z) \right| \leq F. \quad (171)$$

Using the same logic, and considering the event B_N as,

$$B_N = \left\{ |\mathbf{L}_N^{t_N}(z) - (-\Delta_1 f_\theta(z)) f_\theta(z)| \leq \frac{\epsilon}{2F} \right. \\ \left. \cap |\mathbf{L}_N^{t_N}(z)| \left| \frac{1}{n} \sum_{i=1}^n f_\theta(x_i) \lambda(x_i) - \int_{\mathcal{M}} f_\theta(z) \lambda(z) p(z) dV(z) \right| \leq \frac{\epsilon}{4F} \right\} \quad (172)$$

As $B_N \subset C_N$, it follows that $P(C_N) \geq P(B_N)$. Now, we can use Hoeffding's inequality and consider N to be the number of samples such that,

$$P\left(\left| \frac{1}{N} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) - \int_{\mathcal{M}} f_\theta(z) \lambda(z) p(z) dV(z) \right| \geq \frac{\epsilon}{4F}\right) \leq \frac{\delta}{2}. \quad (173)$$

where N should be such that $\frac{\delta}{2} \geq 2e^{-\frac{\epsilon^2 N}{(8bF^2)^2}}$, that is

$$N \geq \frac{64b^2 F^4}{\epsilon^2} \log(\delta/4). \quad (174)$$

Now, we should evaluate the Laplacian term. To do so, we add and subtract the continuous $\tilde{\mathbf{L}}_\lambda^t f_\theta(z)$ (cf. definition 7) as follows,

$$P\left(|\mathbf{L}_N^{t_N}(z) + \tilde{\mathbf{L}}_\lambda^t f_\theta(z) - \tilde{\mathbf{L}}_\lambda^t f_\theta(z) - (-\Delta_1 f_\theta(z))| \leq \frac{\epsilon}{2F}\right) \geq 1 - \frac{\delta}{2} \quad (175)$$

Utilizing the same argument as before, with $\epsilon/4$, and $\delta/4$ we get,

$$P\left(|\mathbf{L}_N^{t_N}(z) - \tilde{\mathbf{L}}_\lambda^t f_\theta(z)| \leq \frac{\epsilon}{4F}\right) \leq \frac{\delta}{4} \quad (176)$$

and,

$$P\left(|\tilde{\mathbf{L}}_\lambda^t f_\theta(z) - (-\Delta_1 f_\theta(z))| \geq \frac{\epsilon}{4F}\right) \leq \frac{\delta}{4} \quad (177)$$

Now, (176) can be bounded utilizing Lemma 7. That is, N needs to verify

$$\delta/4 \geq 2e^{-\frac{\epsilon^2 N}{8\mathbf{B}_t^2}} + 2e^{-\frac{\epsilon^2(N-1)}{4\mathbf{B}_t a^{3/2}}} \quad (178)$$

Note that $\frac{N}{\mathbf{B}_t^2}$ needs to tend to ∞ , as $N \rightarrow \infty$, and $t \rightarrow 0$. This is analogous to $Nt^{2+d} \rightarrow \infty$, therefore, $t = N^{-\frac{1}{2+d+\alpha}}$, for any $\alpha > 0$ satisfies the condition. With this choice of t , all the integral approximations hold by (Belkin and Niyogi, 2005, Lemma 7.1).

For term (177) can be bounded utilizing Lemma 8, and therefore N needs to verify,

$$t^{1/2} \leq \frac{4}{\delta}, \quad (179)$$

$$N^{\frac{-1}{2(2+d+\alpha)}} \leq \frac{4}{\delta} \quad (180)$$

We have therefore obtained the number of samples N required to a lower bound each probability by $1 - \delta/2$. Therefore, the probability of B_N verifies that,

$$\begin{aligned} & P\left(\sup_{\lambda \in \Lambda, \theta \in \Theta} \left| \frac{1}{n} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) \mathbf{L}_N^{t_N} f_\theta(x_i) - \int_{\mathcal{M}} f_\theta(z) \Delta_1 f_\theta(z) \lambda(x) p(z) dV(z) \right| \leq \epsilon \right) \\ & \geq P(C_N) \geq P(B_N) = 1 - P(B_N^c) \end{aligned} \quad (181)$$

$$\begin{aligned} & \geq 1 - P\left(|\mathbf{L}_N^{t_N}(z) - (-\Delta_1 f_\theta(z))| \geq \frac{\epsilon}{2\Phi}\right) \\ & - P\left(\left| \frac{1}{N} \sum_{i=1}^N f_\theta(x_i) \lambda(x_i) - \int_{\mathcal{M}} f_\theta(z) \lambda(z) p(z) dV(z) \right| \geq \frac{\epsilon}{4\Phi}\right) \end{aligned} \quad (182)$$

$$\geq 1 - \delta/2 - \delta/2 = 1 - \delta \quad (183)$$

Where we first used the fact that the complement of an intersection is the union of the complements. Then, we used the take the maximum over N for the events, so as to verify (173) and (175) respectively. Therefore, by taking the complement, we obtain that for any $\epsilon > 0$, and any $\delta > 0$, there exists a number of samples N (173) and (175), such,

$$P\left(\sup_{\lambda \in \Lambda, \theta \in \Theta} \left| \frac{1}{n} \sum_{i=1}^n f_\theta(x_i) \lambda(x_i) \mathbf{L}_n^{t'_n} f_\theta(x_i) - \int_{\mathcal{M}} f_\theta(z) (-\Delta_1 f_\theta(z)) \lambda(x) p(z) dV(z) \right| \geq \epsilon \right) \leq \delta \quad (184)$$

it suffices to pick the total number of samples N_0 , such that it verifies all three conditions (174), (178), and (180). Moreover, t needs to go to zero as $N \rightarrow \infty$, that is why $t = N^{-\frac{1}{d+2+\alpha}}$ for any $\alpha > 0$.

F DUAL ASCENT ALGORITHM

In this section, we explore the two algorithms that we present to solve problem 6; the gradient based primal dual Algorithm 1 in subsection F.1, and the pointcloud laplacian based primal dual 2 in subsection F.2. The laplacian variant, presents a computational advantage, but requires conditions on the dual variables λ that might be difficult to secure in practice. The gradient based method on the other hand, presents a less restrictive algorithm at the expense of a higher computational cost. It is important to note that the sole difference between the two relies on the gradient of the lagrangian with respect to θ , that is to say, the update on the dual variables λ and μ remains the same. In this appendix section, we provide an verbose explanation of the two procedures. Moreover, we elaborate on the estimator of the norm of the gradient F.3.

F.1 GRADIENT BASED PRIMAL DUAL ASCENT

On this subsection we elaborate on the gradient based method to update the primal variable θ . Upon showing that under certain conditions problem (3) and problem (7) are close (see Proposition 1, and 2), we will introduce a dual ascent algorithm to solve the later. The problem that we seek to solve is given by,

$$\max_{\hat{\lambda}, \hat{\mu}} \hat{d}_G(\hat{\lambda}, \hat{\mu}) := \min_{\theta} \frac{1}{N} \sum_{n=1}^N \left(\ell(f_{\theta}(x_n), y_n) - \mu\epsilon \right) + \hat{\mu}^{-1} \frac{1}{N} \sum_{n=1}^N \hat{\lambda}(x_n) \|\nabla_{\mathcal{M}} f_{\theta}(x_n)\| \quad (185)$$

$$\text{subject to } \frac{1}{N} \sum_{n=1}^N \lambda(x_n) = 1 \quad (186)$$

We will procede with an iterative process that minimizes over θ , and maximizes over λ and μ . In short, for each value of μ , λ , we seek to minimize the Lagrangian $\hat{L}(\theta, \hat{\mu}, \hat{\lambda})$ by taking gradient steps as follows,

$$\theta_{k+1} = \theta_k + \eta_{\theta} \nabla_{\theta} \hat{L}(\theta, \hat{\mu}, \lambda) \quad (187)$$

$$= \theta_k + \eta_{\theta} \nabla_{\theta} \left(\hat{\mu} \frac{1}{N} \sum_{n=1}^N \ell(f_{\theta}(x_n), y_n) + \frac{1}{N} \sum_{n=1}^N \hat{\lambda}(x_n) \|\nabla_{\mathcal{M}} f_{\theta}(x_n)\|^2 \right) \quad (188)$$

The gradient base approach computes gradients with respect to the loss function ℓ , and with respect to the norm of the gradient. After the dual function $\hat{d}(\hat{\lambda}, \hat{\mu})$ is minimized, in order to solve 6, we require to maximize the dual function $\hat{d}(\hat{\lambda}, \hat{\mu})$ over both $\hat{\lambda}$, and $\hat{\mu}$, which can be done by evaluating the constraint violation as follows,

$$\lambda_n \leftarrow \lambda_n + \eta_{\lambda} \partial_{\hat{\lambda}_n} \hat{\mathcal{L}}(\theta, \hat{\mu}, \hat{\lambda}), \text{ with } \partial_{\hat{\lambda}_n} \hat{\mathcal{L}}(\theta, \hat{\mu}, \hat{\lambda}) = \|\nabla_{\mathcal{M}} \phi(x_n)\|^2 \quad (189)$$

$$\hat{\mu} \leftarrow \hat{\mu} + \eta_{\hat{\mu}} \hat{\mathcal{L}}(\theta, \hat{\mu}, \hat{\lambda}), \text{ with } \partial_{\hat{\mu}} \hat{\mathcal{L}}(\theta, \hat{\mu}, \hat{\lambda}) = \frac{1}{N} \sum_{n=1}^N \ell(f_{\theta}(x_n), y_n) - \epsilon, \quad (190)$$

where stepsizes $\eta_{\lambda}, \eta_{\mu}$ are positive numbers. Note that to evaluate the norm of the gradient on a particular point $\|\nabla_{\mathcal{M}} \phi(x_n)\|$ we can look at its neighboring points $\mathcal{N}(x_i)$, and estimate its norm. After updating the dual variables λ using gradient ascent, we require $|\lambda|_1 = 1$, which can be done by either normalizing, or projecting to the simplex Wang and Carreira-Perpinán (2013). How to compute the norm of the gradient will be explained in the sequel in subsection F.3. The overall procedure is explained in Algorithm 1.

F.2 POINT CLOUD LAPLACIAN DUAL ASCENT

Our algorithm will take advantage of the Laplacian formulation given in Proposition 3. That is, we will estimate the gradient of the lagrangian with respect to θ utilizing the point cloud laplacian formulation. Formally, for a vector $\hat{\lambda} \in \mathbb{R}_+^N$, such that $\frac{1}{N} \sum_{n=1}^N \hat{\lambda}_n = 1$, and a

Algorithm 1 Gradient Based Smooth Learning on Data Manifold

1: Initialize parametric function θ , define neighborhoods $\mathcal{N}(p)$ for every $p \in \mathcal{D}$
2: **repeat** for epochs $0, 1, \dots$
3: **for** primal steps k **do**
4: Estimate maximum norm of gradient for all x_i : $\|\nabla_{\mathcal{M}} f_{\theta}(x_i)\| = \max_{z \in \mathcal{N}(x_i)} \frac{\|f_{\theta}(x_i) - f_{\theta}(z)\|}{d(x_i, z)}$
5: Update θ : $\theta \leftarrow \theta - \eta_{\theta} \nabla_{\theta} \left(\mu \frac{1}{N} \sum_{i=1}^N \ell(f_{\theta}(x_i), y_i) + \frac{1}{N} \sum_{i=1}^N \lambda(x_i) \|\nabla_{\mathcal{M}} f_{\theta}(x_i)\|^2 \right)$
6: **end for**
7: Update dual variable μ : $\mu \leftarrow [\mu + \eta_{\mu} \frac{1}{N} \sum_{i=1}^N \ell(f_{\theta}(x_i), y_i) - \epsilon]_+$
8: Update dual variable $\lambda(x_i)$: $\lambda(x_i) \leftarrow [\lambda(x_i) + \eta_{\lambda} (\max_{z \in \mathcal{N}(x_i)} \frac{\|f_{\theta}(x_i) - f_{\theta}(z)\|^2}{d(x_i, z)^2})]_+$
9: Project λ : $\lambda = \arg \min_{0 \leq \tilde{\lambda}} \|\tilde{\lambda} - \lambda\|$ s.t. $|\tilde{\lambda}|_1 = N$
10: **until** convergence

Algorithm 2 Point-Cloud Laplacian Smooth Learning on Data Manifold

1: Initialize parametric function θ , fix temperature t
2: **repeat** for epochs $0, 1, \dots$
3: **for** primal steps k **do**
4: Update θ : $\theta \leftarrow \theta - \eta_{\theta} \nabla_{\theta} \left(\mu \frac{1}{N} \sum_{i=1}^N \ell(f_{\theta}(x_i), y_i) + \frac{1}{N} \sum_{i=1}^N f_{\theta}(x_i) \lambda(x_i) \mathbf{L}_{\lambda p, N}^t f_{\theta}(x_i) \right)$
5: **end for**
6: Update dual variable μ : $\mu \leftarrow [\mu + \eta_{\mu} \frac{1}{N} \sum_{i=1}^N \ell(f_{\theta}(x_i), y_i) - \epsilon]_+$
7: Update dual variable $\lambda(x_i)$: $\lambda(x_i) \leftarrow [\lambda(x_i) + \eta_{\lambda} (\max_{z \in \mathcal{N}(x_i)} \frac{\|f_{\theta}(x_i) - f_{\theta}(z)\|^2}{d(x_i, z)^2})]_+$
8: Project λ : $\lambda = \arg \min_{0 \leq \tilde{\lambda}} \|\tilde{\lambda} - \lambda\|$ s.t. $|\tilde{\lambda}|_1 = N$
9: **until** convergence

constant $\hat{\mu} \in \mathbb{R}^+$, the empirical dual function associated with the Lagrangian $\hat{L}(\theta, \hat{\mu}, \hat{\lambda}_n)$ of the empirical dual function associated with (6) is defined as,

$$\hat{d}(\hat{\lambda}, \hat{\mu}) = \min_{\theta} \frac{1}{N} \sum_{n=1}^N \ell(f_{\theta}(x_n), y_n) + \hat{\mu}^{-1} \frac{1}{N} \sum_{n=1}^N f_{\theta}(x_n) \hat{\lambda}_n \mathbf{L}_{\lambda p, N}^{t_N} f_{\theta}(x_n), \quad (191)$$

Note that in (191), we have omitted the term $-\hat{\mu}\epsilon$ as it is a constant term for a given $\hat{\mu}$, and we have divided over $\hat{\mu}$ which renders an equivalent problem as long as $\hat{\mu} > 0$. By taking the maximum of the dual function over $\hat{\mu}$, and $\hat{\lambda}$, we recover the dual problem 7. For a given choice of dual variables, $\hat{\lambda}, \hat{\mu}$, the dual function (191) is an unconstrained problem that only depends on the parameters θ . Now, the link with Manifold Regularization (5) is seen; as considering $\tilde{\lambda}(x_n) = 1/N$, and $\tilde{\mu} = \gamma$, the problems become equivalent. In order to minimize (191), we can update the parameters θ following the gradient,

$$\theta \leftarrow \theta + \eta_{\theta} \nabla_{\theta} \left(\frac{1}{N} \sum_{n=1}^N \ell(f_{\theta}(x_n), y_n) + \hat{\mu}^{-1} \frac{1}{N} \sum_{n=1}^N f_{\theta}(x_n) \hat{\lambda}_n \mathbf{L}_{\lambda p, N}^{t_N} f_{\theta}(x_n) \right), \quad (192)$$

where $\eta_{\theta} > 0$ is a step-size. Note that the updates in 192 have two parts, one that relies on the loss ℓ which utilizes labeled data, and another term given by $\hat{\lambda}_n \mathbf{L}_{\lambda p, N}^{t_N}$, that penalizes the Lipschitz constant of the function, and can be computed utilizing unlabeled data. A more succinct explanation of the algorithm is described in Algorithm 2.

F.3 GRADIENT NORM ESTIMATE

For both Algorithms 1, and 2, we require to compute the norm of the gradient of f_{θ} at sample point x_n with respect to the manifold \mathcal{M} . Leveraging the Lipschitz constant definition 1, in order to estimate the maximum norm of the gradient, we require to evaluate the Lipschitz over the sample points in our dataset. To do so, we require to set a metric distance over the

manifold $d_{\mathcal{M}}$, with which we will define the neighborhood of sample x_n as the samples that are sufficiently close, i.e.

$$\mathcal{N}(x_n) = \{z \in \mathcal{M} : d_{\mathcal{M}}(z, x_n) \leq \delta\}. \quad (193)$$

Note that our definition of neighborhood 193, requires a maximum distance δ . In practice, we can either choose δ , or set the degree of the neighborhood, i.e. choose the k nearest neighbors. Upon the selection of the neighborhood, the norm of the gradient is computed as follows.

$$\|\nabla_{\mathcal{M}} f_{\theta}(x_i)\| = \max_{z \in \mathcal{N}(x_i)} \frac{\|f_{\theta}(x_i) - f_{\theta}(z)\|}{d(x_i, z)} \quad (194)$$

The selection of distance over the manifold is application dependent. In controls problems that involve physical systems, the metric might involve the work required from two states to reach each other. On computer vision applications, common choices of metrics can be perceptual losses Zhang et al. (2018), or distances over embedding Khrulkov et al. (2020).

G DETAILS ON NUMERICAL EXAMPLES

G.1 NAVIGATION CONTROLS PROBLEM

In this section, we consider the give futher details on the problem of continuous navigation of an agent. The agent’s objective is to reach a goal while avoiding obstacles. The state space of the agent is $\mathcal{S} = [0, 20] \times [0, 10]$, which represents the x and y axis respectively. The agent navigates by taking actions on the velocity $v \in \mathbb{R}^2$, and the state evolves according to the dynamics $s_{t+1} = s_t + v_t T_s$ where $T_s = 0.1s$. We construct a square grid of points in the environment that are on the free space i.e. outside of the obstacles, and utilize Dijkstra’s algorithm to find the shortest path for two starting $[1, 9]^T$, $[14, 1]^T$ positions, and goal $[19, 1]^T$ along the grid. For those two grid trajectories, we compute the optimal actions to be taken at each point in order to follow the trajectory. The learner is equipped with both the labeled trajectories, as well as the unlabeled point grid. To leverage the manifold structure of the data, we consider the grid of points, and we construct the point cloud Laplacian considering adjacent points in the grid. We train a two layer neural network using the mean square error loss over the optimal set of points and actions for ERM, ERM with ambient Lipschitz regularizer, Manifold Regularization, and our method Manifold Lipschitz method.

To evaluate the performance, we randomly chose 100 starting points and compute the trajectories generated by each learned function. A trajectory is successful if it reaches the goal without colliding with the obstacles or the walls. The results are summarized in table 1, and showcase the benefit of implementing manifold lipschitz. Our method outperforms the 3 other methods due to the fact that it minimizes the gradient of the function over the domain of the data. As opposed to ERM, our method generates a smooth function outside of the labeled trajectory. Ambient Lipschitz regularization fails due to the fact that the euclidean distance ignores the real distance between samples across wall, forcing similar outputs for points that should take different actions. Manifold regularization is able to capture the similarity between points, but it fails to properly capture the sharp turns near the edges off the obstacles. The success of Manifold Lipschitz, can be explained by its dual variables λ shown in figure 2c. In this figure, the radius of each ball represents the value of the dual variable, which is larger close to the corners of the obstacles due to the fact that the problem requires larger gradients to make sharp turns over it. Besides, the fact that we can disentangle the loss on the labeled data, from the Lipschitz constant, allows us to overfit the data as much as we require.

G.2 TWO-MOONS DATASET

In this subsection we provide the details of the experiment with the Two-moons data set utilized in 1.

To generate the data we utilized `sklearn` library, and we utilize 1 labeled, and 200 unlabeled samples per class (i.e. moon), and we added noise $\sigma = \{0.05, 0.1\}$. For the neural network, we utilized a two layer fully connected neural network with 64 hidden neurons with bias term, and hyperbolic tangent as the non-linearity. For the optimizer, we utilized a learning rate of 0.9, and no momentum. For the ambient regularization, we added a weight decay of 0.1. For the construction of the Laplacian of Figure 1 we utilized a heat kernel temperature of $t = 0.005$, and we normalize it. For Laplacian regularization, we set $\gamma = 0.5$. For Manifold Lipschitz (our method), a μ dual step of 0.5, and a λ dual step of 0.1. An ablation study over different values of temperature coefficient t can be found in sections I.

G.3 GROUND ROBOTIC VEHICLE RESIDUAL LEARNING

In this section we explain the residual learning experiment that involves a ground robot vehicle. The data acquisition of this experiment involves an *iRobot* Packbot equipped with high resolution camera. The setting of the data acquisition, is a robot making turns on both *pavement*, and *grass*. The dynamics of the system, are govern by the discrete-time nonlinear state-space system of equations

$$x_{k+1} = f(x_k, u_k) + g(u_k), \quad (195)$$

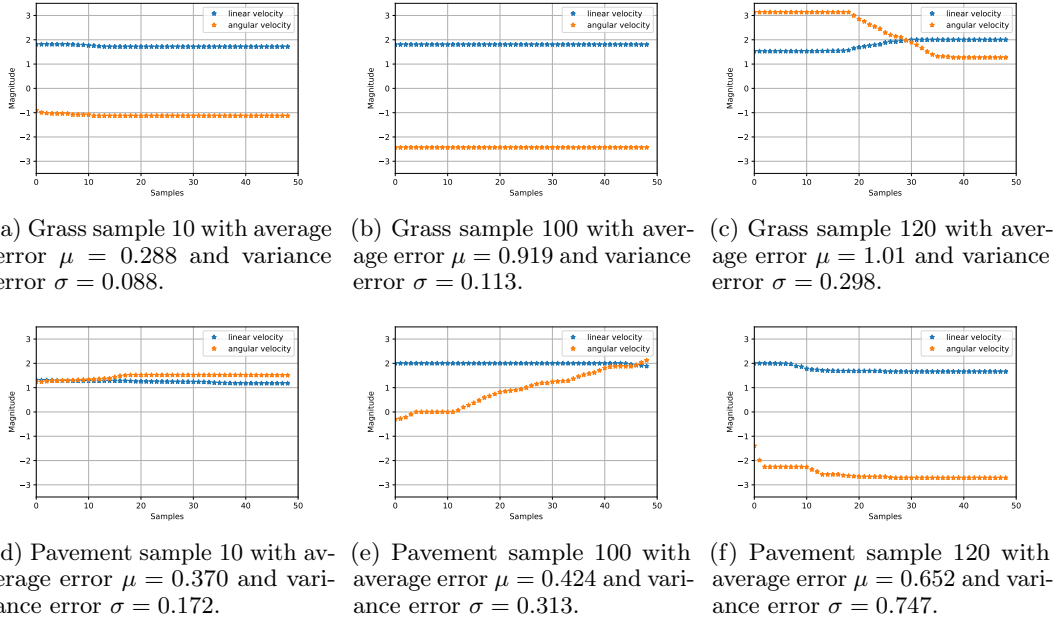


Figure 3: Sample trajectories for grass, and pavement, with their corresponding average error, and variance error.

where x_k is the state of the system, and u_k is the action taken, $f(x_k, u_k)$ is the model prediction, and $g(u_k)$ is a non-modellable error of the prediction. In this particular setting, the robots state involves the position, and the action taken is given by its linear, and angular velocities. The dynamics of the system are modeled by $f(x_k, u_k)$, and they involve the mass of the robot, the radius of the wheel, among other known parameters of the robot. In practice, the model $f(x_k, u_k)$ is not perfect. There exists a model mismatch given the difference in friction with the ground, delays in communications with the sensors, and discrepancies in the robot specifications. To make matters worse, some of the discrepancies, are difficult to model, or intractable to compute in practice. Therefore, we seek to learn the model mismatch $g(u_k)$.

Figure 3 shows examples of trajectories given by time series x_k , each of which is associated with an average error model mismatch, and its corresponding variance,

$$\mu_i = \frac{1}{K} \sum_{k=1}^K \|x_{k+1} - f(x_k, u_k)\|, \quad (196)$$

$$\mu_i = \sqrt{\frac{1}{K} \sum_{k=1}^K \|\mu_i - (x_{k+1} - f(x_k, u_k))\|^2}. \quad (197)$$

In practice, samples from the grass dataset tend to have larger disturbance mean μ , given that pavement presents a more uniform setting. For example, if we compare sample 3b to 3f, we can see that for similar trajectories, the error is larger in the case of grass. Moreover, trajectories that involve larger magnitudes in velocities, also present larger disturbance means, and variances (cf. 3a vs 3c). This is due to the fact that at high speeds, the non-modeled effects such as drift, and friction are more significant.

For more details on the data acquisition, please refer to the original paper Koppel et al. (2016).

G.4 NUMERICAL IMPLEMENTATION

We are equipped with two datasets, one for grass with 195 samples, and one for pavement with 224 samples. We partition the datasets into 170 samples for training and 25 samples

for testing, and 200 samples for training and 24 samples for testing, for grass and pavement respectively. Each sample has associated a disturbance mean, and disturbance variance vector $[\mu, \sigma_i] \in \mathbb{R}^2$.

Regarding the optimization, we utilize the mean square error as a loss, and we train a 2 layer fully connected neural network with 256 hidden dimensions and hyperbolic tangent as the non-linearity. We train for 10000 epochs, with a learning rate of 0.0015 in the case of pavement, and 0.00015 in the case of grass. For batch size, we utilize the whole training set. For the case of ambient Lipschitz, we utilize weight decay $w_D = 0.3$, and $w_D = 0.1$ for grass and pavement respectively.

For the point cloud laplacian, we utilize heat kernel with temperature $t = 15$, and $t = 35$ for grass and pavement respectively. This gives us a point cloud laplacian with one connected component. For the Laplacian regularization we utilize $1e^{-3}$, and $1e^{-4}$ for grass and pavement respectively.

As for the Manifold Lipschitz, we utilize $\eta_\mu = 0.01$, and $\eta_\lambda = 0.1$, and $\epsilon = 0.005$, and 0.01 for grass and pavement respectively. We initialize $\mu = 5$, and ρ uniform.

The final results are summarized in Table 2.

H FURTHER REFERENCES

Since we introduce the Lipschitz constant as a constraint to the learning problem our reformulation and solution methodologies are framed within the constrained learning paradigm (Chamon and Ribeiro, 2020; Chamon et al., 2021; Yang, 2019). Central to the solution of constrained learning problems is the use of dual formulations and dual ascent learning algorithms. These are finding increasing applicability as evidenced by their use in, e.g., adversarial robustness (Robey et al., 2021), graph neural networks (Cervino et al., 2022; Arghal et al., 2021), federated learning (Shen et al., 2021), active learning (Elenter et al., 2022), reinforcement learning (Paternain et al., 2019; 2022; Castellano et al., 2021; Bai et al., 2021; Hasanbeig et al., 2018), and wireless communications (Eisen et al., 2019).

In the context of adversarial attacks to neural networks, manifold based regularization techniques have shown a vast amount of empirical and theoretical evidence of its utility, improving its adversarial robustness (Zhang et al., 2021; Khoury and Hadfield-Menell, 2018; Ma et al., 2018; Moosavi-Dezfooli et al., 2019; Jin and Rinard, 2020; Lassance et al., 2021). Some works seek to obtain manifold attacks, which are more realistic attacks than utilizing the norm- ∞ ball, given the high dimensionality of the input and the low dimensional structure of the data (Stutz et al., 2019). Smooth function have also been studied in the context of robustness (Rosca et al., 2020; Bubeck and Sellke, 2021).

Our work, is based on previous results that show convergence of graph laplacians to Laplace-Beltrami operators. There exists a vast amount of work on that validates the convergence results for point cloud operators over Manifolds (Hein et al., 2005; 2007; Dunson et al., 2021; Wu and Wu, 2018).

Regarding Lipschitz constant estimation for neural networks, Fazlyab et al. (2019) has formulated the problem as a convex optimization problem.

I ABLATION STUDY ON LAPLACIAN CONSTRUCTION

In this section we study the impact of the temperature coefficient t in the construction of the Laplacian. To do so we repeat the setting of Figure 1, and we vary the value of the temperature t . We consider the two moons dataset problem with 1 labeled, and 200 unlabeled samples per class. We vary the value of the temperature coefficient, which varies the number of cross-manifold edges, and therefore makes the problem more challenging.

Value of Heat Kernel	Connected Components	Number of Cross-Manifold Edges	Manifold Regularization	Manifold Gradient (Ours)
0.0040	2	0	100%	100%
0.0050	1	2	N/A	100%
0.0060	1	3	N/A	100%
0.0070	1	10	N/A	100%
0.0080	1	20	N/A	100%
0.0090	1	27	N/A	100%
0.0100	1	42	N/A	100%
0.0150	1	124	N/A	100%
0.0175	1	176	N/A	100%
0.0180	1	192	N/A	100%
0.0190	1	217	N/A	100%
0.0200	1	251	N/A	N/A

Table 3: Ablation study on the temperature of the heat kernel t . We plot the accuracy when it achieves 100%, and N/A otherwise, given that an accuracy of less than 100% is not representative as the method fails to capture the manifold structure of the problem.

As seen in table 3, Laplacian regularization fails to achieve a perfect accuracy when the number of connected components is less than 2. That is to say, manifold regularization achieves a perfect accuracy when each class has a connected component. However, once the components become connected, Laplacian regularization smoothness the integral of the gradient, and therefore does not properly identify the transition between components.

As can be seen in table 3, our method is more robust to non-exact manifolds. Which means that if the manifold is not calculated perfectly, and as a result we obtain 1 connected component as opposed to 2 separate moons, our method still works.

It is important to remark that our method still works when the connected components have cross-manifold edges in different places of the manifold. As an example, take the Laplacian with heat kernel $t = 0.007$, there are edges on both side of the manifold. Moreover, our method is able to distinguish between the two classes even with edges in the middle of the two manifold as can be seen with $t = 0.0150$, and $t = 0.0175$ (cf. 4h).

Our method brakes once the manifold structure vanishes and most of the points are inter-connected, as can be seen with $t = 0.02$. In this case there are 251 cross manifold edges, and the low-dimensional structure of the problem disappears.

In all, our proposed solution is more robust to imperfect estimation of the manifold. Even when the number of cross-edges is large, our method is able to create a partition between classes given that it finds the points with maximal separation, and allows the function to change values between them.

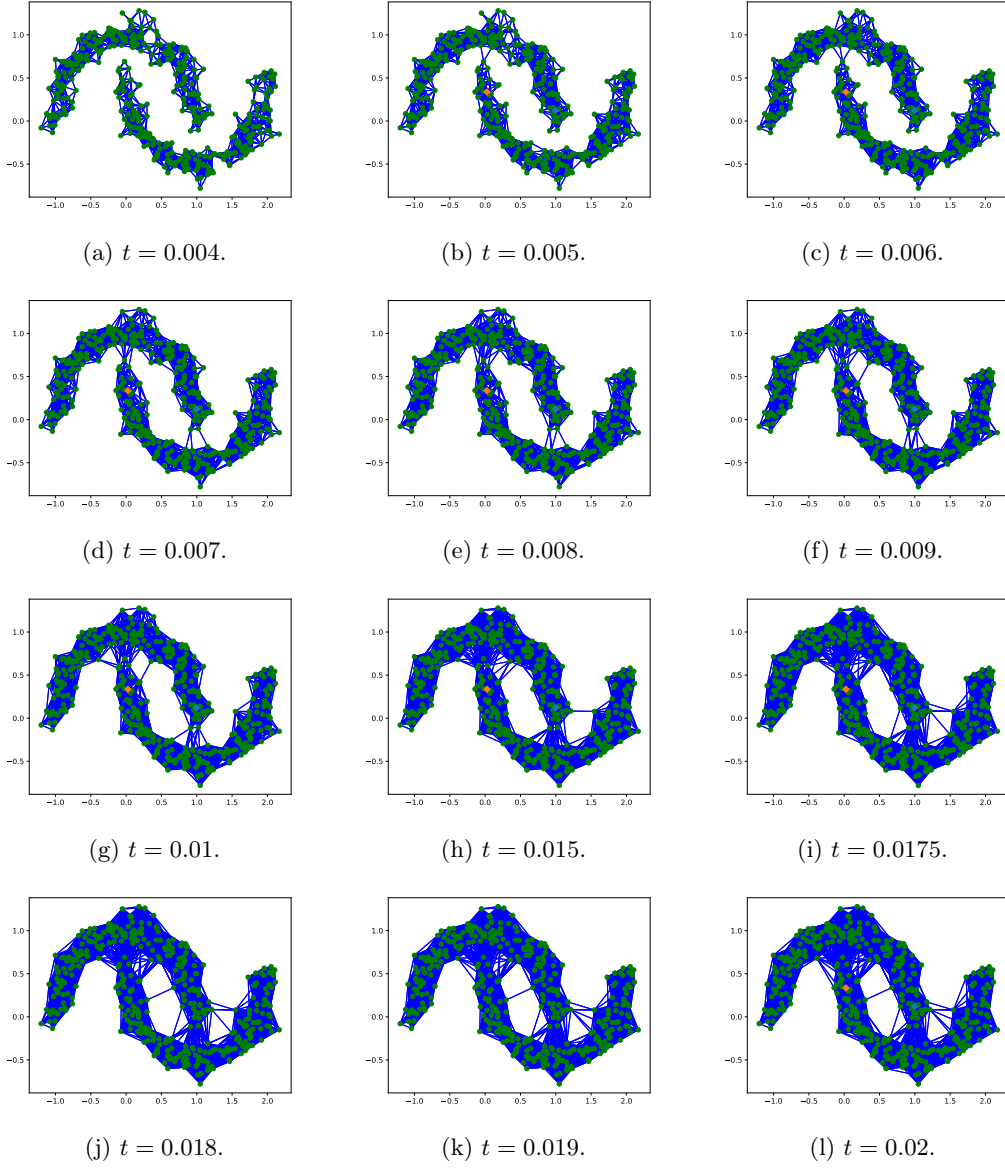


Figure 4: Laplacian for different values of temperature coefficient t .

J FURTHER NUMERICAL EXPERIMENTS - QUADROTOR STATE PREDICTION

In this section we present a state prediction problem based on a real world collected from a quadrotor taking off and flying in a circle for a 15 seconds. The aerial robot is the open-source Crazyflie 2.1 quadrotor (<https://www.bitcraze.io/products/crazyflie-2-1/>), which has a mass of 32 g, and a mass of size 9 cm². The quadrotor communicates with a computer running on Intel i7 CPU, and the communication is established with the Crazyradio PA and at a nominal rate of 500 Hz ($T_s = 1/500$). To measure the position of the quadrotor a VICON is utilized. From the position is obtained from the VICON, whereas the accelerations are obtained from the on-board accelerometers and gyroscope sensors. For further information on this setup, please refer to Jiahao et al. (2022); Chee et al. (2022); Jiahao et al. (2021).

The experimental setup is to target a speed of 0.4m/s and to track a circular trajectory of radius 0.5m. We consider 2 trajectories of 12 seconds each (each trajectory has 6000 time stamps), and the starting position of the quadrotor is the same for both trajectories. For each time stamp t , we have measurements of position, velocity and acceleration in \mathbb{R}^3 , i.e. $[x_t, y_t, z_t, \dot{x}_t, \dot{y}_t, \dot{z}_t, \ddot{x}_t, \ddot{y}_t, \ddot{z}_t]$.




Figure 5: Quadrotor Sample Trajectories

The learning problem consists of learning the dynamical system, i.e. we consider the dynamics given by the equation,

$$[x_{t+1}, y_{t+1}, z_{t+1}, \dot{x}_{t+1}, \dot{y}_{t+1}, \dot{z}_{t+1}] = f([x_t, y_t, z_t, \dot{x}_t, \dot{y}_t, \dot{z}_t, \ddot{x}_t, \ddot{y}_t, \ddot{z}_t]) \quad (198)$$

For the learning procedure we utilize the 6000 samples, and we seek to minimize the mean square error loss between the next state and the prediction given the current state (cf. equation 198). We train a two layer neural network with different methods as seen in Table 4. To test the neural network, we compute the difference between the predicted state, and the next state on the test trajectory. Note that even though the training, and testing trajectories are not the same, there is a resemblance between the two of them.

To begin with, we can conclude that adding regularization not always helps, as ambient regularization does not improve upon the ERM prediction method. A salient conclusion of the results shown in Table 4, is that adding regularization on the manifold space always improves upon ERM. In particular, our method is almost 3 times better than standard ERM, and more than 2 times better than standard Laplacian Regularization.

Method	Error on Test Trajectory
ERM	0.00666
Ambient Reg.	0.00734
Manifold Reg.	0.00625
Manifold Lipschitz	0.00237

Table 4: State prediction error on test trajectory.

To conclude with, we show that our method obtains an improvement over all the techniques considered. This allows us to conclude that in predicting the next state of a quadrotor from the current state under noisy measurements utilizing smooth functions improves generalization.

J.1 DETAILS ON NUMERICAL IMPLEMENTATION

For the state prediction problem of a quadrotor we utilized a two layer fully connected neural network with 8192 hidden units, hyperbolic tangent as the non-linearity, and bias term. For the optimizer, we utilized a learning rate of 10^{-5} , and the full dataset per batch. We trained until convergence in all cases with number of epochs $e = 1000$. For the ambient regularization we used weight decay 0.1. For the construction of the Laplacian, we utilized temperature coefficient $t = 0.01$. For Laplacian regularization we utilized $\gamma = 10^{-6}$. For our method, we utilized μ dual step 0.5, $\epsilon = 0.003$, and λ dual step 0.1. In all cases we trained until convergence with $e = 10000$ epochs.