



Contents lists available at ScienceDirect

Journal of Computational Mathematics and Data Science

journal homepage: www.elsevier.com/locate/jcmds

Exploring singularities in data with the graph Laplacian: An explicit approach



Martin Andersson, Benny Avelin*

Department of Mathematics, Uppsala University, S-751 06 Uppsala, Sweden

ARTICLE INFO

MSC:

primary 58K99
secondary 68R99
60B99

Keywords:

Graph Laplacian
Geometry
Singularities

ABSTRACT

We develop theory and methods that use the graph Laplacian to analyze the geometry of the underlying manifolds of datasets. Our theory provides theoretical guarantees and explicit bounds on the functional forms of the graph Laplacian when it acts on functions defined close to singularities of the underlying manifold. We use these explicit bounds to develop tests for singularities and propose methods that can be used to estimate geometric properties of singularities in the datasets.

1. Introduction

High dimensional data is common in many research problems across academic fields, and such data can often be represented as points collected or sampled from \mathbb{R}^N . A common assumption is that the dataset $X = \{X_i\}_i^n \subset \mathbb{R}^N$ lies on a lower-dimensional set Ω and is in fact a sample from some probability distribution over Ω . A further assumption, that makes our model of data more tractable, is that Ω can be represented as the union of several well-behaved manifolds, i.e. $\Omega = \cup_i \Omega_i$. Here, each Ω_i could represent a different class in a classification problem: For instance, if a dataset contains two classes, i and j , class i might be contained in Ω_i and class j in Ω_j , with the two classes potentially being disjoint. However, classification is not always so clear-cut: For instance, in the MNIST dataset, handwritten digits of $\varepsilon 1 \varepsilon \in \Omega_1$ and $\varepsilon 7 \varepsilon \in \Omega_7$ can appear very similar, suggesting that $\Omega_1 \cap \Omega_7 \neq \emptyset$. Therefore, understanding geometric situations such as intersections is of interest in classification problems.

In our manifold model of data, an intersection between two different manifolds Ω_i, Ω_j is either represented just as such, or it can be viewed as a singularity if we consider $\Omega = \Omega_i \cup \Omega_j$ as a single manifold. Other regions in Ω that can be viewed as singular, such as boundaries and edges, may also be of interest as they can signify important features in the data. Regions that are part of just one manifold Ω_i , and in its interior, we consider as non-singular.

To study such singularities, we use the graph Laplacian $L_{n,t}$. This operator, which depends on the number of data points n and a parameter t , can act on functions defined on the dataset X . In non-singular regions, as n tends to infinity and t tends to 0, $L_{n,t}$ converges to the Laplace–Beltrami operator [1]. In this work, we primarily study the behavior of $x \rightarrow L_{n,t}f(x)$ for functions f , when x is close to singular points.

1.1. Motivation

Consider data points situated in the ambient space, where the underlying manifold structure is unknown. We examine the graph Laplacian operator applied to the linear function $f(x) = v \cdot x$, where v is a unit vector. On our discrete dataset, this function reduces

* Corresponding author.

E-mail addresses: martin.andersson@math.uu.se (M. Andersson), benny.avelin@math.uu.se (B. Avelin).

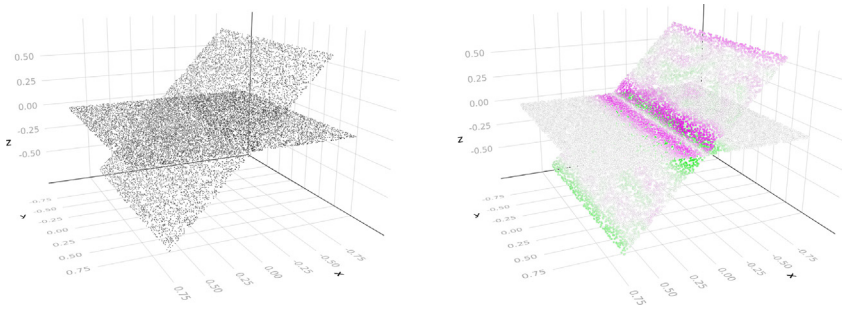


Fig. 1. Graph Laplacian $L_{n,t}$ acting on a linear function f . Purple color showing positive, and green color negative values of $L_{n,t}f$, where lack of color indicates values near 0. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

to the vector $\mathbf{f} = (v \cdot X_1, \dots, v \cdot X_n) \in \mathbb{R}^n$. The contribution of this paper is that we calculate explicitly how the graph Laplacian acts on such linear functions f , which proves to be valuable for estimation purposes.

The choice of f is motivated by the convergence properties of $L_{n,t}$ to the Laplace–Beltrami operator. In the interior of Ω , we expect that $L_{n,t}f(x) \approx 0$. However, for singular points like intersections, the limit operator is of first order [2], and $L_{n,t}f(x) \neq 0$, which can be seen in Fig. 1. Our results show how $x \rightarrow L_t f(x)$ and, through a finite-sample bound, how $x \rightarrow L_{n,t}f(x)$ behaves. More specifically, given $x_0 \in \Omega_i$ near some singularity, and x in the ball $B_R(x_0)$, including the case when $x \notin \Omega_i$, we show how the function $x \rightarrow L_{n,t}f(x)$ deviates from being constantly 0 and has specific functional forms. These forms depend on the type of singularity.

1.2. Overview of results

In Section 4.1, we consider the scenario where $\Omega = \cup_{i=1}^m \Omega_i$ is a union of compact d -dimensional flat (linear) submanifolds of \mathbb{R}^N . This geometric configuration, illustrated in Fig. 3, is particularly relevant to the neural network architecture discussed in Section 2.1.

To set up the results, we start with an $x_0 \in \Omega$ which is assumed to be a flat manifold, and let $x \in B_R(x_0)$, where $R = \sqrt{t}r_0 > 0$, and use \hat{x} to denote the projection of x to the tangent space of Ω . We also define $v_{n,\Omega}$ as the projection of v onto $x - \hat{x}$, and $v_{n,\partial\Omega}$ is the projection of v onto the outwards normal of $\partial\Omega$, which is assumed to be flat (locally linear) close to x . Then we show the following:

- In Theorem 2, we let $\|x - x_0\| = r\sqrt{t}$ and θ is the angle between vectors $x - x_0$ and $\hat{x} - x_0$. If x is not close to $\partial\Omega$, then

$$L_t f(x) = A(x)t^{\frac{d+1}{2}} v_{n,\Omega} \sin(\theta) r e^{-\sin^2(\theta)r^2} + t^{\frac{d+1}{2}} B(x) e^{-r^2}.$$

The function A is close to being constantly equal to $\pi^{d/2}$, and B is uniformly bounded. Note that for small t we have r_0 can be taken large and thus the last term is very small. Both functions have explicit bounds.

- Theorem 3 shows what happens when x is close to $\partial\Omega$:

$$L_t f(x) = \hat{A}_1(x)t^{\frac{d+1}{2}} v_{n,\Omega} \sin(\theta) r e^{-\sin^2(\theta)r^2} + \hat{A}_2(x)t^{\frac{d}{2}} v_{n,\partial\Omega} e^{-\sin^2(\theta)r^2} + t^{\frac{d+1}{2}} B(x) e^{-r^2},$$

where functions \hat{A}_1, \hat{A}_2 and B have explicitly computable bounds.

In Sections 4.2 and 4.3 we prove more general results:

- In Theorem 4 we relax the conditions on Ω , considering non-flat manifolds, and prove a weaker version of Theorem 2.
- In Theorem 5 we relax the conditions further, and allow for noise when sampling from Ω .

In Section 5 we show how the results can be used to construct hypothesis tests for intersections in data. We test this method in Section 6.1. We also show how such hypothesis tests can be used to detect singularities in zero sets of neural networks in Section 6.1.1.

Finally, in Section 6.2 we propose methods to find intersections in data and estimate the angle of such intersections, which are motivated by the aforementioned theorems and Corollary 4.4. We also provide numerical experiments, in Section 6, to test these methods.

2. Earlier work

The framework of assuming an underlying low-dimensional manifold of data, in conjunction with graph-related tools and, in particular, the graph Laplacian, has been used extensively. Some examples include work in clustering [3–7], dimensionality reduction [8,9], and semi-supervised learning [10].

Several of the approaches to study datasets that use the graph Laplacian leverage that if the manifold is smooth enough and well-behaved, then the graph Laplacian approximates some well-understood operator (for instance, the Laplace–Beltrami operator [11]), which has useful mathematical properties.

Therefore, convergence properties of the graph Laplacian becomes important, and it has been studied in [1,9,12,13]. In particular, and highly influential of this paper, is what the asymptotic convergence looks like near singularities of the manifold, which was shown in [2].

The special case that Ω is composed of linear manifolds has received earlier attention. For instance, trying to recover facts about the underlying geometry has been explored in [14,15]. However, these methods have not utilized the graph Laplacian, but other means.

2.1. Main motivating application

There is a lot of interest in understanding the geometry of the loss landscape of neural networks. In particular, zero sets of the loss function.

As a motivating example for the methodology we develop in this paper, consider the following problem. Consider the following class of neural networks:

$$f_W(x) = \sum_{i=1}^k a_i(w_i \cdot x)_+, \quad x \in \mathbb{S}^1, \quad a_i = \pm 1, \quad w_i \in \mathbb{R}^2 \quad (2.1)$$

where \mathbb{S}^1 is the unit circle. For general spheres this is what is called a spherical neural network and has in essence all the same properties as a standard single hidden layer network. This class of networks has been studied in for instance [16–18]. Let a target function $g(x) = f_{W^*}(x)$ be given, but with unknown parameter W^* . We are interested in the set of parameters Ω_δ such that

$$\Omega_\delta := \{W \in \mathbb{R}^{2k} : |f_W(x) - g(x)| < \delta, \quad x \in \mathbb{S}^1\}.$$

We know by the definition of such networks that Ω_δ is the sublevelset of a piecewise linear function. It is thus clear that the sublevelset is itself a union of linear manifolds.

3. Basic mathematical objects and theory

In this section, we provide more precise definitions and introduce the basic mathematical theory we will be using to present and prove our results.

3.1. Conditions on manifolds

We will consider sets of the form $\Omega = \cup_i^m \Omega_i$, where each Ω_i is a smooth and compact d -dimensional Riemannian submanifold of \mathbb{R}^N . We will assume that if any pair Ω_i, Ω_j , and $i \neq j$, have a non-empty intersection, then this intersection will have dimension lower than d .

Our analysis will assume that we only have access to samples from Ω , and so associated to Ω there is a probability measure with density $p : \Omega \rightarrow \mathbb{R}$. For clarity of presentation of our results, we assume that the restriction of p to Ω_i is a uniform distribution. Relaxing this would not significantly alter the main arguments, but the analysis, however, would become more cumbersome.

For any point $x \in \Omega_i$ in the interior, we can consider the tangent space $T_{\Omega_i, x} \simeq \mathbb{R}^d$, which we will identify as a subspace of the ambient space \mathbb{R}^N . More precisely, given open subsets $U \subset \mathbb{R}^d$ and $W \subset \Omega_i$ (W is open in the subspace topology of Ω_i), and a coordinate chart $\alpha : U \rightarrow W$ such that $\alpha(0) = x$, we define $T_{\Omega_i, x}$ as the image of \mathbb{R}^d under the action of the Jacobian. We denote the Jacobian $D\alpha : U \rightarrow \mathbb{R}^{N \times d}$, evaluated at 0, by $D\alpha(0)$. The best linear approximation to $u \mapsto \alpha(u)$ is, of course, given by $u \mapsto x + D\alpha(0)u$, and $x + T_{\Omega_i, x}$ is a first order approximation to Ω_i at x .

The definition of Ω implies that an interior point $x \in \Omega$ can have more than one associated tangent space. For example, if $x \in \Omega_i \cap \Omega_j$ and $i \neq j$, then both $T_{\Omega_i, x}$ and $T_{\Omega_j, x}$ exist, and can be different.

A note on notation is that we denote the interior of a manifold Ω_i by $\text{Int } \Omega_i$, and the boundary by $\partial\Omega_i$.

3.2. Types of singularities

The following are what we will refer to as singular points, which will be of four different kinds. Given $x \in \Omega = \cup_i \Omega_i$, we have the following types:

- (Type 1) There is a submanifold Ω_i such that $x \in \partial\Omega_i$.
- (Type 2) There are submanifolds $\Omega_i \neq \Omega_j$ such that $x \in \text{Int } \Omega_i \cap \text{Int } \Omega_j$.
- (Type 3) There are submanifolds $\Omega_i \neq \Omega_j$ such that $x \in \partial\Omega_i \cap \text{Int } \Omega_j$.
- (Type 4) There are submanifolds $\Omega_i \neq \Omega_j$ such that $x \in \partial\Omega_i \cap \partial\Omega_j$.

The different types above can of course have non-empty intersection with each other, and a *non-singular* point is simply a point $x \in \text{Int } \Omega_i$ such that if $j \neq i$, then $x \notin \Omega_j$. See Fig. 2 for two examples of singularities.




Fig. 2. There is a singularity in the intersection of the lines above. The left figure shows a point of Type 4, and the right figure shows a point of Type 2.

3.2.1. (L, r) -Regular manifolds

To formulate our results for general non-flat manifolds, we will need some measure of how regular, with regard to curvature, our set Ω is. The following definition captures the necessary information.

Definition 3.1. Let $\Omega = \cup \Omega_i$ be a union of compact submanifolds in \mathbb{R}^N . We also let $r > 0$ be the largest radius such that any point $x \in \text{Int } \Omega$, such that $B_r(x) \cap \partial\Omega = \emptyset$, allows coordinate charts $\alpha : U \rightarrow B_r(x) \cap \Omega_i$, where $U \subset \mathbb{R}^d$ and $B_r(x) \subset \mathbb{R}^N$ is an open ball of radius r around x . Further, assume also that conditions (A.3) and (A.4) hold over all tuples $(U, W, x, y, \pi, \Omega_i)$ for some $L > 0$. Then we say that Ω is (L, r) -regular.

Example 3.2. Any smooth and compact submanifold is (L, r) -regular for some $r > 0$. For instance the graph of the function $x \rightarrow x^2$ over the compact interval $[-1, 1]$ is $(1, 1)$ -regular.

3.3. Graph Laplacian

In this section we introduce the graph Laplacian and how it acts on real-valued functions defined on \mathbb{R}^N .

Given n i.i.d. random samples $X = \{X_1, \dots, X_n\}$ from the distribution with density p on Ω , we build a weighted fully connected graph $G = (V, E)$ as follows: We let each sample X_i represent a vertex i , and for vertices $i, j \in V$ the weight on $(i, j) \in E$ is given by

$$W_{n,t}(i, j) := W_{n,t}(X_i, X_j) = \frac{1}{n} K_t(X_i, X_j) = \frac{1}{n} e^{-\frac{\|X_i - X_j\|^2}{t}}.$$

The function $W_{n,t}$ is naturally viewed as an $n \times n$ matrix, and the variable t is in the literature often referred to as the *bandwidth* of the kernel K_t .

Remark 3.3. In the limit analysis as $n \rightarrow \infty$, it is useful also to normalize by $\frac{1}{\mu^{d/2+1/2}}$. But, since a priori we do not know the dimension d , we will work without this normalization.

We define the diagonal weighted degree matrix as

$$D_{n,t}(i, i) = \sum_j W_{n,t}(i, j),$$

and the *graph Laplacian* $L_{n,t}$ as

$$L_{n,t} = D_{n,t} - W_{n,t}.$$

Remark 3.4. This is often referred to as the *unnormalized graph Laplacian*. There are other normalizations of this matrix which are used, for example, in [6,7,13]. One difference between these normalizations are their limit properties.

Given the fully connected graph $G = (V, E)$, the graph Laplacian above can be seen as an operator acting on arbitrary functions $f : V \rightarrow \mathbb{R}$ in the following way:

$$L_{n,t}f(X_i) = \frac{1}{n} \sum_j K_t(X_i, X_j)(f(X_i) - f(X_j)), \quad (X_i, X_j) \in E.$$

We extend this operator to acting on functions that live in the ambient space $f : \mathbb{R}^N \rightarrow \mathbb{R}$ by the canonical choice

$$L_{n,t}f(x) = \frac{1}{n} \sum_j K_t(x, X_j)(f(x) - f(X_j)), \quad x \in \mathbb{R}^N. \tag{3.1}$$

Several of our results will be stated in terms of the expected operator:

$$L_t f(x) = \mathbb{E}_p[L_{n,t} f(x)] = \int_{\Omega} K_t(x, y)(f(x) - f(y))p(y) dy. \tag{3.2}$$

Further, an immediate consequence of the linearity of the integral is that

$$L_t f(x) = \int_{\Omega} K_t(x, y)(f(x) - f(y))p(y) dy = \sum_i \int_{\Omega_i} K_t(x, y)(f(x) - f(y))p(y) dy. \tag{3.3}$$

This is useful to us, and in our approach we work with the *restricted Laplacian* L_t^i , which we define as

$$L_t^i f(x) = \int_{\Omega_i} K_t(x, y)(f(x) - f(y))p(y) dy. \tag{3.4}$$

3.4. The class of functions

As mentioned in the introduction, we apply the graph Laplacian to linear functions of the form $f(x) = v \cdot x$, where v is a unit vector, and all our results are stated with this assumption.

3.5. A concentration inequality

As stated earlier, many of our results pertain to the expected operator L_t . To connect these results to the operator $L_{n,t}$, we prove the following concentration inequality:

Theorem 1. *Let X_1, \dots, X_n be i.i.d. samples from a density p on the union of manifolds $\Omega = \cup \Omega_i$. Let $f(x) = v \cdot x$ for $x \in \Omega$ and v is a unit vector. Then*

$$\mathbb{P} \left(\max_i \left| L_{n,t} f(X_i) - \frac{n-1}{n} L_t f(X_i) \right| > \varepsilon \right) \leq 2n \exp \left(-\frac{4\varepsilon(n-1)\varepsilon^2}{t} \right).$$

This concentration can be improved if we have some further information about the manifold. For instance, if we know that the manifold is flat, and we know the dimension. This allows us to get tight control over the variance of the graph Laplacian, which is what improves the concentration, but we do not explore this further.

4. Main results

Now that we have the necessary definitions and mathematical background, we are ready to present our main results. Before stating the theorems, we will provide a brief section that explains the geometry of some terms that will be used.

General structure of results

By Eq. (3.3) it is enough to understand the restricted Laplacian, L_t^i defined in Eq. (3.4). Because of this, our results are formulated to show the behavior of L_t^i . Depending on what type of singularity being examined, it is straight forward to extend the results to the full Laplacian. In [Corollary 4.4](#) we give one example of how to extend the results to the sum $\sum_{i=1}^2 L_t^i$ when one is close to an intersection of two manifolds.

Geometry and notation for Section 4.1

We will in several theorems also formulate the function $x \rightarrow L_t^i f(x)$ partly in terms of new coordinates (r, θ) . Here r is defined by the relation $\|x - x_0\| = \sqrt{t}r$, and given the projection \hat{x} of x to the tangent plane of the flat manifold Ω_i , we define $\theta \in [0, \pi/2]$ to be the angle between vectors $x - x_0$ and $x_0 - \hat{x}$, as the schematic in [Fig. 3](#). By simple geometry, it follows that $\|\hat{x} - x\| = r\sqrt{t} \sin \theta$.

Given a vector $v \in \mathbb{R}^N$, we will have reason to write the expression $v \cdot (\hat{x} - x)$ as

$$v \cdot (\hat{x} - x) = r\sqrt{t} \sin(\theta)v \cdot \frac{\hat{x} - x}{\|\hat{x} - x\|} = r\sqrt{t} \sin(\theta)v_{n,\Omega_i}(x),$$

where we have defined

$$v_{n,\Omega_i}(x) := \begin{cases} v \cdot \frac{\hat{x} - x}{\|\hat{x} - x\|} & \text{if } x \neq \hat{x} \\ 0 & \text{if } x = \hat{x}. \end{cases}$$

Here v_{n,Ω_i} is the projection of v onto a unit normal vector of Ω_i , and which can vary with x . We will later use a natural coordinate system parallel to Ω_i , see [Appendix A.1](#), such that when $x \neq \hat{x}$, this function is constant up to its sign. This implies that evaluating $r\sqrt{t} \sin(\theta)v_{n,\Omega_i}(x)$ is the same as letting v_{n,Ω_i} be fixed, but allowing θ to change sign depending on which side of Ω_i x is, i.e. as if we have fixed the coordinate system in which we measure the angle θ . In the following, to increase readability, we suppress the x -dependency of v_{n,Ω_i} .

Additionally, in [Theorem 3](#) we will have a term $v_{n,\partial\Omega_i}$ that is specific to that theorem. This will be defined in the case where there is a boundary close to x . In [Fig. 3](#), this would imply there is a boundary of Ω_1 nearby. To give the definition of this term first assume that Ω_i is a flat manifold and assume that close to the point x , $\partial\Omega_i$ is flat. Then, we first let $\hat{x}_{\partial\Omega_i}$ be the projection of \hat{x} to




Fig. 3. Schematic picture of the geometry of Theorem 2, where Ω_1 is the object of interest and $x \in \Omega_2$ for visualization purposes.

$\partial\Omega_i$. We can now define a unit normal at $\hat{x}_{\partial\Omega_i}$, denoted by $n_{\partial\Omega_i}$. Two choices are natural, a normal pointing either towards, or away from Ω_i . We define $n_{\partial\Omega_i}$ as the latter. Given a vector $v \in \mathbb{R}^N$, we define

$$v_{n,\partial\Omega} := v \cdot n_{\partial\Omega}.$$

Since we assume that close to x , $\partial\Omega$ is flat, $n_{\partial\Omega}$ can be considered constant.

Geometry and notation for Section 4.2

To help with the geometric picture for general manifolds, the situation is as explained in Section 4 and Fig. 3: the terms x, x_0, \hat{x}, θ and v_{n,Ω_i} are in the same relation to each other as in Section 4. The main difference is that the tangent space now changes with x as the manifold is no longer flat. In that sense the geometry for more general manifolds is not more difficult, but handling error terms is more involved.

4.1. Flat manifolds

In this section, in addition to the assumptions presented in Section 3.1, we assume that each Ω_i is a flat manifold, i.e. a d dimensional compact linear submanifold of \mathbb{R}^N . This implies that each coordinate chart around a point $x \in \text{Int } \Omega_i$ is an isometry with an open neighborhood U of x , where U is a ball in \mathbb{R}^d .

In Theorem 2 we give a result concerning the behavior of $x \rightarrow L_i^j f(x)$ when we are not close to the boundary $\partial\Omega$. This case is easier to prove, and we give explicit bounds of all terms involved, and express them with elementary functions.

In Theorem 3 we show what happens when we are close to $\partial\Omega$, and here our results contain more complex expressions for some terms.

In the following theorems, we think of x_0 as the potentially singular point, see Fig. 3. The appearance of terms different from zero either comes from the fact that we are evaluating the restricted operators outside of their corresponding manifolds, or because that we are close to the boundary. For example, if $\Omega = \Omega_1 \cup \Omega_2$ and $x_0 \in \Omega_1 \cap \Omega_2$, the extra contribution thus comes from each L_i^j , $i \in \{1, 2\}$, when evaluated on the other manifolds, i.e. $L_i^j f(x)$ when $x \in \Omega_2$ and vice versa. By combining Theorems 2 and 3, it is possible to consider several types of singularities defined in Section 3.2.

Theorem 2. Let $f(x) = v \cdot x$ for some unit vector $v \in \mathbb{R}^N$ and assume that p is the uniform density over Ω_i . Let $x_0 \in \Omega_i$ and assume that $\partial\Omega_i \cap B_{2R}(x_0) = \emptyset$ for $R = r_0\sqrt{t}$, where $r_0 > 2$. Further, $x \in B_R(x_0)$, and v_{n,Ω_i} , r and θ are as described in Section 4. If $t \leq \frac{R^2}{d/2+1}$, $d \geq 1$ and $r < r_0/2$, then we have that

$$L_i^j f(x) = t^{d/2+1/2} \left(A(\theta, r_0, d) v_{n,\Omega_i} \sin \theta r e^{-\sin^2 \theta r^2} + B(x) e^{-r_0^2} \right),$$

where A, B are real-valued functions. The function B depends on x , but it is uniformly bounded by $|B(x)| \leq \frac{d+1}{4} r_0^d \frac{|\mathbb{S}^{d-1}|}{|\Omega_i|}$; and A depends on x only through θ , and is bounded by

$$\frac{1}{2} \max(\pi^{d/2}, 2\pi^{d/2} - |\mathbb{S}^{d-1}| 2^{d/2} r_0^{d-1} e^{-r_0^2+1}) \leq A(\theta, r_0, d) \leq \pi^{d/2}.$$

The following theorem is an extension of Theorem 2 to the case when the ball $B_R(x_0) \cap \partial\Omega_i \neq \emptyset$, which gives rise to an additional term in the expression of $L_i^j f(x)$. We again refer to the schematic picture of Fig. 3 and comments in Section 4 for explanation of the coordinates (r, θ) , function v_{n,Ω_i} and constant $v_{n,\partial\Omega_i}$.

Theorem 3. Let $f(x) = v \cdot x$ for some unit vector $v \in \mathbb{R}^N$, and assume that p is the uniform density over Ω_i . Let $x_0 \in \Omega_i$ and assume that $\partial\Omega_i \cap B_{2R}(x_0)$ is part of a $d - 1$ dimensional plane for $R = r_0\sqrt{t}$, where $r_0 > 2$. Further, $x \in B_R(x_0)$, and v_{n,Ω_i} , $v_{n,\partial\Omega_i}$, r and θ are as

described in Section 4. If $t \leq \frac{R^2}{d/2+1}$, $d \geq 1$ and $r < r_0/2$, then we have that

$$L_t^i f(x) = \widehat{A}_1(x)t^{\frac{d+1}{2}}v_{n,\Omega_i} \sin(\theta)r e^{-\sin^2(\theta)r^2} + \widehat{A}_2(x)t^{\frac{d}{2}}v_{n,\partial\Omega_i} e^{-\sin^2(\theta)r^2} + B(x)t^{\frac{d+1}{2}}e^{-r^2},$$

for explicitly computable function \widehat{A}_2 , and with explicitly computable bounds of function \widehat{A}_1 . The function B has the same bounds as in Theorem 2.

Remark 4.1. The function \widehat{A}_1 is bounded by

$$\frac{1}{2\delta_0} \left(e^{-k_0^2} \gamma \left(\frac{d-1}{2}, \delta_0^2 - k_0^2 \right) - \frac{2(\delta_0^2 - k_0^2)^{\frac{d-1}{2}}}{d-1} \right) \leq \widehat{A}_1 \leq \Gamma \left(\frac{d-1}{2} \right) \sqrt{\pi}$$

and \widehat{A}_2 is given by

$$\widehat{A}_2 = \frac{|\mathbb{S}^{d-2}|}{2} \left(e^{-\delta_0^2} \frac{(\delta_0^2 - k_0^2)^{(d-1)/2}}{d-1} + \frac{1}{2} e^{-k_0^2} \gamma \left(\frac{d-1}{2}, \delta_0^2 - k_0^2 \right) \right),$$

where $\Gamma(\cdot)$ is the Gamma function, and $\gamma(\cdot, \cdot)$ is the incomplete lower Gamma function, see Appendix A.3 for details. To define k_0 and δ_0 , we recall the geometric picture of Section 4. Then K is the projection of $(\hat{x} - \hat{x}_{\partial\Omega_i})$ to $n_{\partial\Omega_i}$, $k_0 = K/\sqrt{t}$, and $\delta_0 = \sqrt{r_0^2 - r^2 \sin^2 \theta}$.

4.2. General manifolds

In this section we will not assume the manifolds are flat, instead we assume that Ω is $(L, 2R)$ -regular, see Definition 3.1. The type of singularity we deal with for a more general manifold will be a Type 2, and we will assume we are not too close to any boundary. Although we achieve a similar result as in Theorem 2, the constants depend on the curvature of the manifold which in many practical applications would be unknown. This comes from the non-local nature of the t -regularized Laplacian L_t . It is however useful if we are in a situation where some a-priori information about the curvature is known, for instance if we are interested in exploring level sets of certain functions.

Theorem 4 (General Manifold). Let $f(x) = v \cdot x$ for some unit vector $v \in \mathbb{R}^N$ and assume that p is the uniform density over Ω_i , a $(L, 2R)$ -regular manifold. Let $x_0 \in \Omega_i$ and assume that $\partial\Omega_i \cap B_{2R}(x_0) = \emptyset$ for $R = r_0\sqrt{t}$, where $r_0 > 2$. Further, $x \in B_R(x_0)$, and v_{n,Ω_i} , r and θ are as described in Section 4. If $LR^2/\sqrt{t} \leq \frac{1}{2}$, $t \leq \frac{R^2}{d/2+1}$, $d \geq 1$ and $r < r_0/2$, then we have that

$$L_t^i f(x) = t^{d/2+1/2} \widehat{A}(x)v_{n,\Omega_i} r \sin \theta e^{-r^2 \sin^2 \theta} + t^{d/2} C_{L,R}(x) 4p\pi^{d/2} + D(x)e^{-r^2}.$$

In the above, $\widehat{A} > 0$ is a function such that

$$|A(d, r, \theta) - \widehat{A}(x)| \leq (1 + 3C_{L,R})A(d, r, \theta)$$

where $A(d, r, \theta)$ as in Theorem 2; $C_{L,R}$ is a function such that

$$|C_{L,R}(x)| \leq 4(L4R^2)^2 R/t + LR^2(1 + 4LR^2)$$

and $|D(x)| \leq \text{diam}(\Omega)$.

The next lemma gives useful bounds on $L_t^i f(x)$ when x is non-singular.

Lemma 4.2. Given the conditions of Theorem 4 and the additional assumption that $x \in \Omega_i$, we have that

$$L_t^i f(x) = t^{d/2} \widehat{A}(x)LR^2 + t^{d/2} C_{L,R}(x) 4p\pi^{d/2} + D(x)e^{-r^2},$$

Remark 4.3. The result in Lemma 4.2 can be used together with both Theorem 2 and Theorem 4 to analyze the behavior of the mapping $x \rightarrow L_t f(x)$ around intersections.

The following corollary is a consequence of Theorem 4 and Lemma 4.2. The geometry is as in , projecting x specifically to the tangent plane T_{Ω_i, x_0} .

Corollary 4.4. Let $f(x) = v \cdot x$ for some vector $x \in \mathbb{R}^N$ and assume that p is the uniform density over a $(L, 2R)$ -regular manifold $\Omega = \cup_{i=1}^2 \Omega_i$. Let $x_0 \in \Omega_1 \cap \Omega_2$ and assume that $\partial\Omega_i \cap B_{2R}(x_0) = \emptyset$ for $i \in \{1, 2\}$ and $R = r_0\sqrt{t}$, where $r_0 > 2$. Assuming R and t satisfies the conditions of Theorem 4, then for $x \in B_R(x_0) \cap \Omega_2$ such that $\|x - x_0\| = r\sqrt{t}$ for $r < r_0/2$, we have that

$$L_t f(x) = t^{d/2+1/2} \widehat{A}(x)v_{n,\Omega_1} r \sin \theta_1 e^{-r^2 \sin^2 \theta_1} + t^{d/2} \widehat{A}(x)LR^2 + t^{d/2} C_{L,R}(x) 8p\pi^{d/2} + 2e^{-r^2} D(x)$$

In the above, θ and v_{n,Ω_1} are as in Section 4, with $\Omega_i = \Omega_1$. Functions \widehat{A} , $C_{L,R}$ and D are as in Theorem 4.

4.3. Manifolds with noise

In earlier sections, we have assumed that the samples used to evaluate $L_{n,t}f(x)$ are taken directly from Ω . However, in many real-world applications we can only expect that we access the points X_1, X_2, \dots, X_n approximately, for instance due to measurements uncertainties. In this section we propose a scenario where our results can have application also in this more general case. For instance, we could have the following situation:

- (1) We first obtain a set observations x_1, x_2, \dots, x_n of X_1, X_2, \dots, X_n from Ω .
- (2) We cannot observe these points directly, but instead we are able to take multiple, noisy, measurements from the random variables $x_1 + \varepsilon_1, x_2 + \varepsilon_2, \dots, x_n + \varepsilon_n$, where each $\varepsilon_i \sim \mathcal{N}(0, \sigma^2 I)$

To model this is, we replace the operator

$$L_{n,t}f(x) = \frac{1}{n} \sum_{j=1}^n K_t(x, X_j)(f(x) - f(X_j)),$$

by

$$L_{n,t,\varepsilon}f(x) = \frac{1}{n} \sum_{j=1}^n K_t(x, X_j + \varepsilon_j)(f(x) - f(X_j + \varepsilon_j)).$$

The following theorem shows that with our results, and with regards to the second source randomness, the uncertainty in measurements, we can understand the expectation of $L_{n,t,\varepsilon}$.

Theorem 5 (Stochastic Version). *Let $L_{n,t,\varepsilon}$ be as above, and the operator $\mathbb{E}_\varepsilon[\cdot] = \mathbb{E}[\cdot \mid X_1, \dots, X_n]$ be expectation with regard to the random variables $(\varepsilon_1, \dots, \varepsilon_n)$. Then*

$$\mathbb{E}_\varepsilon L_{n,t,\varepsilon}f(x) = \frac{2t^{N/2+1}}{(2\sigma^2 + t)^{N/2+1}} \frac{1}{n} \sum_{j=1}^n K_{2\sigma^2+t}(x, X_j)(f(x) - f(X_j)).$$

Thus, for $t' = t + \sigma^2$ and up to normalization, $\mathbb{E}_\varepsilon L_{n,t,\varepsilon}$ is the operator as $L_{n,t'}$.

5. Hypothesis test for singularities in flat manifolds

In this section we will develop a hypothesis test for singularities in flat manifolds. We begin by assuming that $\Omega = \cup_i \Omega_i$ is a union of flat manifolds, and that we have a set of samples $\{X_1, \dots, X_n\}$ from the uniform density p on Ω . In this section we wish to consider the following hypothesis for a fixed point $x_0 \in \Omega$:

- H_0 : $B_{2R}(x_0) \cap \Omega = B_{2R}(x_0) \cap \Omega_i$ for some i and $B_{2R}(x_0) \cap \partial\Omega_i = \emptyset$.
- H_1 : Hypothesis H_0 is false.

Theorem 6. *Let $\Omega = \cup \Omega_i \subset \mathbb{R}^N$ be a union of flat manifolds and let $f(x) = v \cdot x$ for some unit vector $v \in \mathbb{R}^N$. Assume that p is the uniform density over Ω . Let X_1, \dots, X_n i.i.d. samples from p , define the statistic*

$$T = \max_{m: X_m \in B_1(x_0)} |L_{n,t}f(X_m)|,$$

and for a fixed test-level $\alpha \in (0, 1)$ the rejection region $\Theta := \{T > \delta\}$, where

$$\delta = \sqrt{\frac{t}{e(n-1)} \log\left(\frac{2n}{\alpha}\right)}.$$

If

$$t \leq \min \left\{ 1, \frac{2}{\log\left(\frac{e(n-1)}{2 \log\left(\frac{2n}{\alpha}\right)}\right)} \right\},$$

and we reject H_0 in the rejection region Θ , then this is a test of level α .

Remark 5.1. Under the null hypothesis that Ω_i is (L, R) -regular, we can use the same test statistic T , but we change our rejection region by redefining δ using [Theorem 4](#). This δ will necessarily be larger and thus if we wish to detect singularities in more general manifolds, we need more samples to achieve the same power.

As with all hypothesis tests we can only safely reject the null hypothesis, but why we reject it is not clear. Looking at [Theorem 3](#) it could either be that we have a singularity of the types [Type 1–Type 4](#) if we assume that all submanifolds are flat. However, the

reason for rejection could also be that the manifold is not flat. What we will outline next is a way to estimate the power of the test in [Theorem 6](#) in the case when the alternative hypothesis is that there is an intersection between two flat manifolds. Similar results can be obtained for a singularity of any of the types [Type 1–Type 4](#).

Theorem 7. *Let $\Omega = \Omega_1 \cup \Omega_2$ and otherwise as in [Theorem 6](#). Let H'_1 be an alternative hypothesis that there is an intersection at x_0 between Ω_1 and Ω_2 where both are flat manifolds of intrinsic dimension d , where x_0 is a singular point of [Type 2](#) and $\partial\Omega_i \cap B_{2R}(x_0) = \emptyset$, $i = 1, 2$. Let T be the test statistic in [Theorem 6](#) and δ be as in [Theorem 6](#). Then the probability of rejecting H_0 under H'_1 (power) is bounded by*

$$\mathbb{P}(T > \delta \mid H'_1) \geq 1 - \alpha - \mathbb{P}(X_m \in \mathbf{B}_{R/3}^c(x_0))^n,$$

when

$$t(n, \alpha) = \min \left\{ 1, \frac{2}{\log \left(\frac{e(n-1)}{\log \left(\frac{2n}{\alpha} \right)} \right)} \right\},$$

and n is so large that

$$t \left((d+1) \log(1/t) + \log \left(\frac{2}{\pi^d |v_{n,\Omega_1}|^2} \right) \right) \leq 2(1 - \sin^2 \theta_1), t \leq \frac{1}{d/2 + \log \left(\frac{16^2 e^2}{\pi^d |v_{n,\Omega_1}|^2 \sin^2 \theta_1} \right)}. \tag{5.1}$$

Remark 5.2. Asymptotically as $|v_{n,\Omega_1}|$ and $\sin \theta$ tends to zero, the number of samples needed to satisfy [Eq. \(5.1\)](#) is

$$\frac{n}{\log n} = O \left(\frac{1}{|v_{n,\Omega_1}|^4 \sin^4 \theta_1} \right).$$

As such we see that we drastically lose power if the angle θ_1 is small, or if the normal vector v_{n,Ω_1} is small. This is intuitive, as the smaller $|v_{n,\Omega_1}|$ is, the less we see the intersection in the chosen direction v and if it is zero, the Laplacian applied to f will not see the intersection at all. On the other hand, if θ_1 is small, the intersection becomes less distinct. When $\theta = 0$, the two manifolds are tangent, and from the perspective of the Laplacian (for small t), they are effectively the same manifold, as such there is no singularity. In a practical application we can search for a vector v such that T is maximized (on a subsample). Thus, the only real constraint is the actual angle of intersection between the manifolds.

6. Numerical experiments

6.1. Hypothesis test for singularities

In this section we will demonstrate the hypothesis test in [Theorem 6](#) for a flat manifold. We will assume that we have a set of samples (X_1, X_2, \dots, X_n) distributed according to the uniform density on Ω , and that we have a fixed point $x_0 \in \Omega$. We will assume that the null hypothesis H_0 is true, and that x_0 is not close to any boundary. Specifically we will assume that $N = 3$ and that Ω is piece of a hyperplane in \mathbb{R}^3 . We will then evaluate $L_{n,t} f(x)$ on (X_1, X_2, \dots, X_n) and calculate the test statistic T as in [Theorem 6](#), with t as in [Theorem 7](#). We will then calculate the rejection region Θ and see if T is in Θ . You will find the results in [Fig. 4](#).

To test the power of [Theorem 7](#) we construct an intersection between two hyperplanes in \mathbb{R}^3 at different angles. The result is in [Fig. 4](#).

6.1.1. Neural networks

As mentioned in the introduction (see [Section 2.1](#)) we are interested in the zero set of a simple, non-trivial, neural network. Specifically, let us consider the zero set Ω_δ of a neural network as in [Section 2.1](#), where $k = 3$ and the target function f_{W^*} is chosen such that all three weights w_i are the same, $a_1 = -a_2 = -a_3$ and $\delta = 10^{-16}$. The choice of $k > 2$ ensures that the zero set Ω_δ is large, as due to cancellations in the sum in [Eq. \(2.1\)](#), there should be many weights that produce the same function f_{W^*} . Note that the dimension of the weights W is $2k = 6$, in this example.

As described in [Appendix F](#), the dataset we use is an approximation of

$$\hat{\Omega}_\delta := \{W \in \mathbf{B} : |f_W(x_i) - f_{W^*}(x_i)| < \delta, \quad x_i \in \mathcal{D}\},$$

where \mathbf{B} is the box $[-10, 10]^6$ and \mathcal{D} is a dataset of 100 points on the unit circle. The method used to generate the dataset, as described in [Appendix F](#), is called constraint propagation. It can be computationally intensive, but leads to a dataset of 12,753,597 points X with the following key properties:

Samples	Rejections under H_0	Rejections under $H'_1 (\theta = \pi/4)$	$H'_1 (\theta = \pi/2)$
$2 \cdot 10^4$	0	0	0.04
$3 \cdot 10^4$	0	0	0.73
$4 \cdot 10^4$	0	0	1
$5.5 \cdot 10^4$	0	0.2	1
$6 \cdot 10^4$	0	0.28	1
$6.5 \cdot 10^4$	0	0.57	1
$7 \cdot 10^4$	0	0.81	1

Fig. 4. Results of hypothesis test for singularities in flat manifolds, each experiment run 100 times.




Fig. 5. Two visualizations of neural network analysis: (left) PCA projection of a neural network’s zero set from \mathbb{R}^6 to \mathbb{R}^3 ; (right) Graph Laplacian analysis showing $L_t f(x)$ on the y -axis and $f(x)$ on the x -axis, where $f(x) = v \cdot x$. The dashed line in the right plot represents the rejection region θ specified in Theorem 6.

- (1) Each point in X is within a small distance ($0.005\sqrt{6}$) of a point in $\hat{\Omega}_\delta$.
- (2) If we put a box around each point in X with side-length 0.01, then the union of these boxes will cover $\hat{\Omega}_\delta$.

A larger number of nodes (k) or a smaller side-length of the boxes leads to an increase in the computational cost. Our particular choice of box-size (0.01) and $k = 3$ were made to balance computational tractability, accuracy and the possibility of a complex enough zero set.

The zero set itself $\hat{\Omega}_\delta$, by its construction and the piecewise linearity of the activation function, constitutes a union of flat (linear) manifolds, as highlighted in Section 4.1.

In Fig. 5 we have for illustration purposes projected a subsample of our X of $\hat{\Omega}_\delta$ to \mathbb{R}^3 , using PCA, and this plot suggests that these manifolds are also intersecting. We wish to apply the hypothesis test in Theorem 6 to verify this.

We can see the result of applying the graph Laplacian evaluated on a subsample of X to a linear function $f(x) = v \cdot x$ in Fig. 5 (v was chosen to maximize the absolute value of the operator $L_t f$). The dashed line is the rejection threshold from Theorem 6, with parameter t chosen according to Theorem 7. The direction v was chosen on an independently sampled subset of X and the test performed on another independent subset of 30000 samples. The conclusion is that we can reject the null hypothesis that X is a union of flat manifolds that do not intersect at scale $R = 1$ where $x_0 = 0$.

As the above hypothesis test rejects the null hypothesis, we can look closer at the behavior of $L_t f$ in the vicinity of x_0 . Note that, in the case where close to x_0 , the manifold $\Omega = \cup_{i=1}^m \Omega_i$ is an intersection between flat manifolds but without their boundaries intersecting $B_2(x_0)$, see Theorem 2. Then each expected operator L_t^i will have the form

$$L_t^i f(x) = t^{d/2+1/2} \left(A(d, r_0, \theta_i) v_{n,\Omega_i} \sin \theta_i r e^{-\sin^2 \theta_i r^2} + B(x) e^{-r^2} \right).$$

If the vector v found above, makes one of $|v_{n,\Omega_i}|$ large (recall the definition from Section 4) and the others small, we would expect to see the first term in the above expression present in the full operator $L_t f$. Indeed, since $v \cdot x \approx \sin \theta_i r$ for x close to x_0 the first term of this operator exhibits behavior similar to the function ye^{-y^2} with respect to $y = \sin \theta_i r$ for small enough t . By our choice of v , it is clear in the right panel of Fig. 5 that such a term is present in the operator $L_t f$ if we compare with the ideal signal in Fig. 6. This is a strong indication that the zero set of the neural network close to x_0 contains an intersection of flat manifolds.




Fig. 6. Graph of $y = xe^{-x^2}$.

6.2. Estimating singularities

In the following experiments we demonstrate how to estimate the location of an intersection and the intersecting angle θ of a union of two manifolds, given that we know that $\Omega = \Omega_1 \cup \Omega_2$, and $\Omega_1 \cap \Omega_2$ is non-empty and we have a [Type 2](#) singularity. Further, we assume that we have both a set of samples $X \subset \Omega$ distributed according to the associated density on Ω , and an additional set of points Y from a curve $\Gamma \subset \Omega_i$, for some $i \in \{1, 2\}$. The curve Γ intersects $\Omega_1 \cap \Omega_2$, and we assume that no other singularity is very close, which is a situation like in [Figs. 7](#) and [10](#).

This particular setup could, for instance, be understood in the context of urban traffic flow analysis, where Ω_1 and Ω_2 represent two intersecting roads in a city, or some kind of interchange. In this scenario,

- The set X represents GPS data points from different vehicles traveling on the roads or interchange.
- The set $Y \subset \Gamma$ represents samples from the path of a specific vehicle traveling through the intersection or interchange.
- The intersection $\Omega_1 \cap \Omega_2$ is the point of intersection between the roads, and the properties of $\Omega_1 \cap \Omega_2$ could be of interest for understanding traffic flow.

Now let $Y = \{y_1, \dots, y_m\} \subset \Gamma$ be as above. When we evaluate $x \mapsto L_{n,t}f(x)$ on Y , we get the values $P = \{p_1, \dots, p_m : p_i = L_{n,t}f(y_i), p_i \in \Gamma\}$. From our previous results, we know that these values, with enough samples, are close to

$$L_t f(y_i) = t^{d/2+1/2} \left(A(d, r_0, \theta) v_{n,\Omega_i} \sin \theta_i r e^{-\sin^2 \theta r_i^2} \right) + \text{Error}(y_i, r_0, L). \tag{6.1}$$

The function depends on $r = \|y_i - x_0\|$, and the angle $\theta = \angle(y_i - x_0, \hat{y}_i - x_0)$, where \hat{y}_i is the projection of y_i onto either, in the flat case, a manifold Ω_i or, in the curved case, its tangent plane, as defined in [Theorem 4](#), for some point $x_0 \in \Omega_i$. The error term depends on x , r_0 and L . As before, the constant L bounds the curvature of Ω . The error term can be bounded using the results in [Sections 4.1](#) and [4.2](#). By choosing a sufficiently small t , we can make r_0 arbitrarily large, causing the function $A(d, r_0, \theta)$ to approach $2\pi^{d/2}$. For our purposes in this section, we assume that $A(d, r_0, \theta) \approx 2\pi^{d/2}$ and that $\text{Error}(x, r_0, L) \approx 0$ are good enough approximations. With these assumptions,

$$L_t f(y_i) = t^{d/2+1/2} 2\pi^{d/2} v_{n,\Omega_i} \sin \theta_i r e^{-\sin^2 \theta r_i^2}. \tag{6.2}$$

Remark 6.1. [Theorem 4](#) uses $\hat{A}(x)$ instead of $A(d, r_0, \theta)$, but $\hat{A}(x)$ can be made arbitrarily close to $A(d, r_0, \theta)$ by choosing t or L small enough. How small we can make L , however, depends on how well-behaved the manifold is.

We will use the functional form of equation [Eq. \(6.2\)](#) and our two sets of samples to get estimators to the location and the angle of intersection of our manifold. We do not give explicit probabilistic bounds here, but assume $L_{n,t}f$ is sufficiently close to L_t .

Remark 6.2. Let us note that a similar concentration bound as in [Theorem 1](#) should be possible for this type of scenario, which would allow us to get explicit bounds for these experiments. That and not using the approximations above, would allow us to get explicit probabilistic intervals for all estimators.

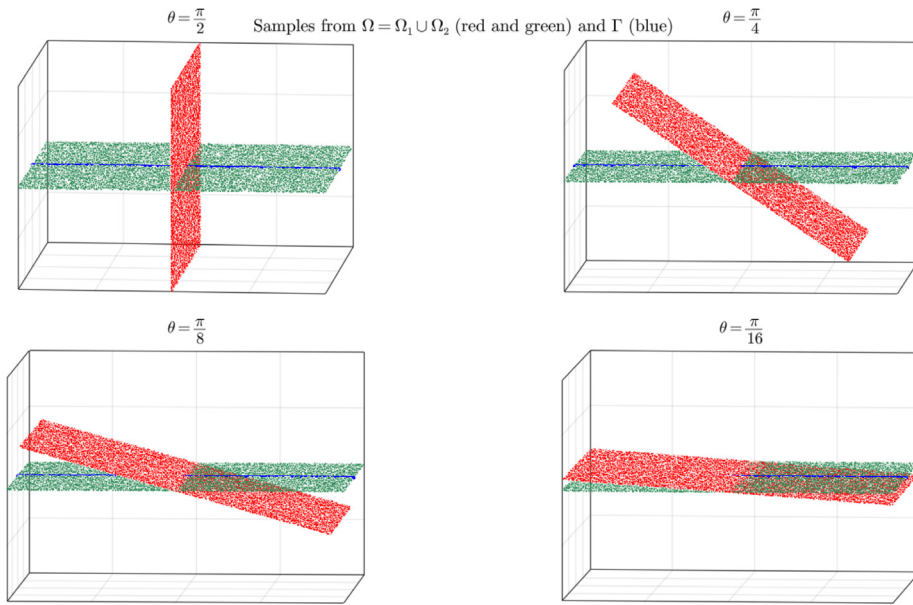


Fig. 7. Samples of $\Omega = \Omega_1 \cup \Omega_2$ and Γ (blue), where Ω_1 (green) and Ω_2 (red) have no curvature. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 8. $L_{n,i}f$ evaluated on Γ . Flat manifolds.

6.2.1. Estimators for intersection location and angle

In this section we derive estimators for the location of the intersection and angle of intersection, based on Eq. (6.2).

Let $g(r, \theta) = v_{n,\Omega} r \sin \theta e^{-r^2 \sin^2 \theta}$. We observe that g only depends on $x = r \sin \theta$ (up to the sign of $v_{n,\Omega}$), allowing us to write $g(r, \theta) = g(x)$, with $g(x)$ being a rescaled version of the function $h(x) = xe^{-x^2}$. See Fig. 6 for an illustration of h .




Fig. 9. Estimates of θ and $|s|$ on flat manifolds.

The function h attains its minimum and maximum values at $z_1 = -\frac{1}{\sqrt{2}}$ and $z_2 = \frac{1}{\sqrt{2}}$ respectively. The intersection corresponds to the midpoint of these two points. We can estimate the point s where Γ^i intersects $\Omega_1 \cap \Omega_2$ as:

$$\hat{s} = \frac{\arg \max_{x_i}(P) + \arg \min_{x_i}(P)}{2}.$$

To estimate the intersection angle $\hat{\theta}$ of θ , we first need \hat{r}_{\max} , an estimate of the scaled distance from 0 that maximizes $g(r, \theta)$. This is

$$\hat{r}_{\max} = \frac{\|\hat{s} - \arg \max_{x_i} P\|}{\sqrt{t}}.$$

Since $\max_r g(r, \theta) = \frac{1}{\sqrt{2}}$, we have that $\hat{r}_{\max} \sin \theta \approx \frac{1}{\sqrt{2}}$. Solving for θ , this leads to our estimator:

$$\hat{\theta} = \arcsin\left(\frac{1}{\sqrt{2}\hat{r}_{\max}}\right).$$

6.2.2. Experimental setup

We tested our estimation methods on hypersurfaces in \mathbb{R}^3 intersecting at angles given by $\theta \in \{\pi/2, \pi/4, \pi/8, \pi/16\}$, and with $t = 10^{-3}$. For both flat and curved manifolds, we conducted 100 runs with random chosen functions of $f(x) = v \cdot x$, where v is sampled from the uniform distribution on \mathbb{S}^2 .

For each run:

- We sampled 2×10^4 points, from both Ω_1 and Ω_2 , in a bounded region near the intersection
- We evaluated $L_{n,t}$ on 10^3 uniformly spaced points on Γ to obtain our set P .

Remark 6.3. We chose \mathbb{R}^3 mainly for visualization purposes. The intrinsic dimension of Ω remains low, regardless of the ambient dimension. However, choosing a function $f(x) = v \cdot x$ for L^i to act on becomes more challenging in higher dimensions, as it becomes harder to orient v such that v_{n,Ω_i} is large, especially when v is chosen randomly, as in our experiments.

In the following sections we present our results for each type of manifold and angle of intersection.

6.2.3. Results for flat manifolds

In this section, we apply our estimation methods to the case that Ω_1, Ω_2 are flat. Since the manifolds are flat, the angle $\theta_i = \theta$, where $\theta \in [0, \frac{\pi}{2}]$ is fixed.

In Fig. 7 we see our samples of Ω and Γ , and in Fig. 8 we see an example of the values we get in P . Finally, in Fig. 9 we see how well this approach works in trying to learn both θ and s .

6.2.4. Results for curved manifolds

Here we test these methods in the case that $\Omega = \Omega_1 \cup \Omega_2$ is $(L, 2R)$ -regular, with $L = 0.5$ and R having no upper bound. The setup is the same as what we see in Fig. 10.

If we let θ denote the angle between the tangent spaces of Ω_1 and Ω_2 at x_0 , then since there is curvature, we cannot expect $\theta_i = \theta$ for every $i = 1, \dots, n$, or even between any pair of them. However, we can still estimate the location of the intersection as before, and estimating θ in this way provides some information about the intersection, even if it is not as strong as in the case without curvature. The range of possible values for θ , due to curvature, can be bounded by knowing the bounds of the curvature constant L .

In Fig. 10 we see our samples of Ω and Γ , in Fig. 11 we see an example of the values we get in P . Finally, in Fig. 12 we see how well this approach works in trying to learn both θ and s .

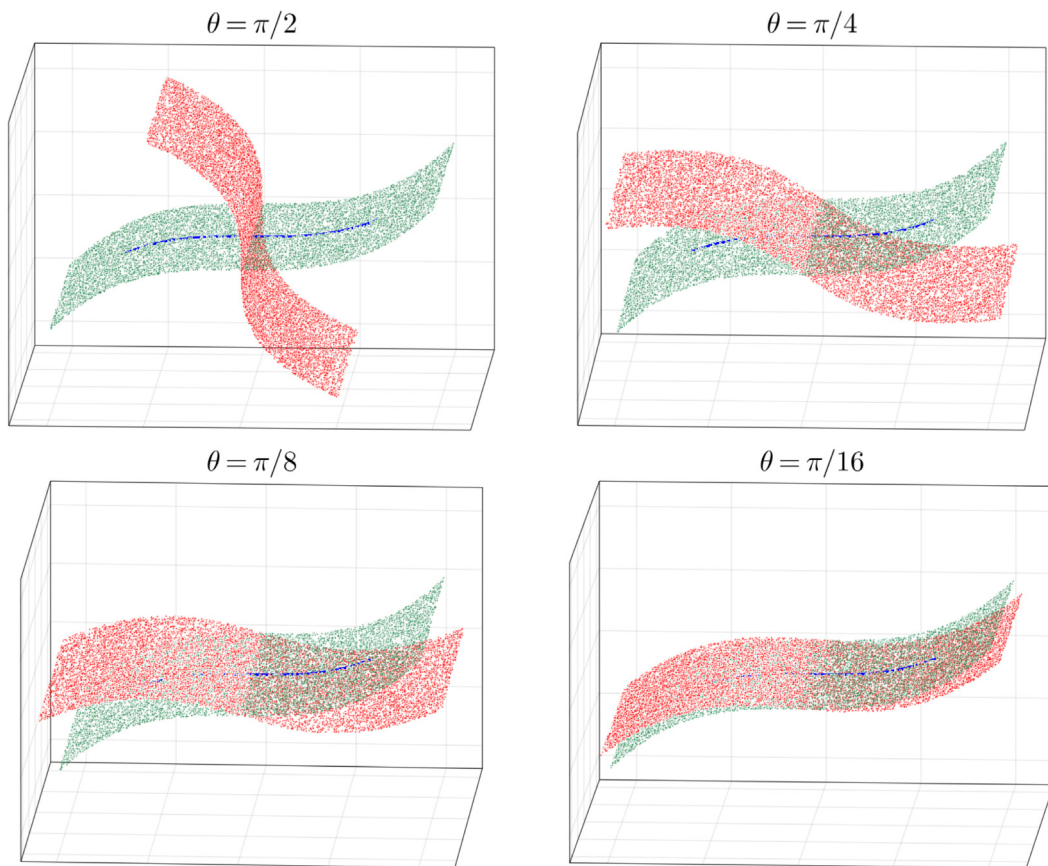


Fig. 10. Samples of $\Omega = \Omega_1 \cup \Omega_2$ and Γ (blue), where Ω_1 (green) and Ω_2 (red) have curvature. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

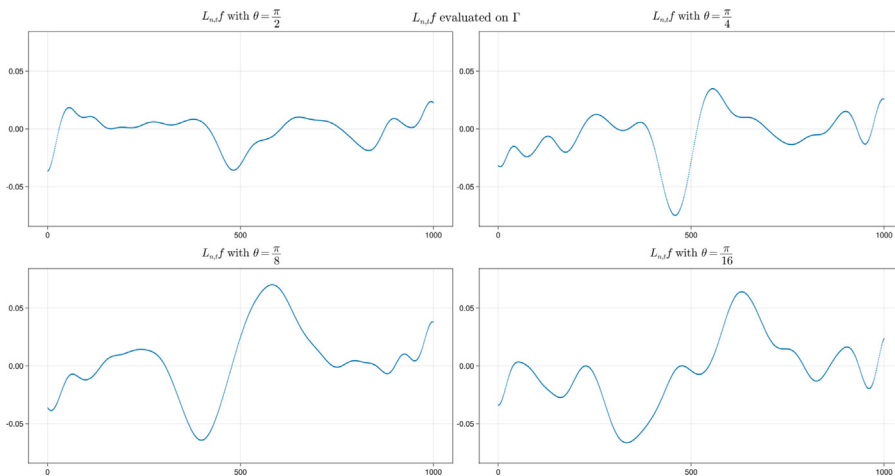


Fig. 11. $L_{n,t}f$ evaluated on Γ . Curved manifolds.

7. Final remarks

In this paper we built upon the work of [2] and developed explicit versions of their asymptotic analysis of $x \rightarrow L_t f(x)$. Our results are the strongest and most useful in the case of flat manifolds, and the motivation to focus on this scenario comes partly from neural networks (see Section 2.1).




Fig. 12. Estimates of θ and $|s|$ on curved manifolds.

We used these explicit bounds to construct hypothesis tests with theoretical guarantees on the level of the test as well as the power under specific conditions. The numerical experiments in Section 6.1 suggest that the power can be close to 1 with many samples as predicted by Theorem 7. The power of these hypothesis tests can clearly be seen in the example of the zero set of a neural network (see Section 6.1.1) where we can safely reject the null hypothesis that the zero set is a single flat manifold.

While the bounds in Theorem 4 are weaker, our numerical experiments suggest that this approach can be useful for gaining geometric information about the union of more general manifolds $\Omega = \cup_i \Omega_i$. In [2], the authors mainly considered sets $\Omega = \cup_{i=1}^n \Omega_i$ with $n \leq 2$. Our approach of splitting L_i into components L_i^l makes it straight forward to directly apply our theorems for $n \geq 2$, allowing us to consider a wider range of singularities. For example, we can extend the framework to examine points that are both of Type 1 and Type 2, or to study intersections of more than two manifolds.

In our numerical experiments in Section 6.2, we assumed that $\Omega = \cup_{i=1}^2 \Omega_i$ and had access to samples of an intersecting curve Γ , which allowed us to estimate geometric properties near intersections. Future work could involve extending our framework to other types of singularities and developing similar tests and estimators. It would also be interesting to explore methods that do not rely on access to such curves.

Similar theorems can be proven for other kernels besides the Gaussian one, as many ideas used in our proofs are not specific to the Gaussian case, but rather rely mainly on symmetries of K_r . Investigating the use of other kernels and comparing their performance in different scenarios is a potential direction for future research.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The first author was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, Sweden. The second author was supported by the Swedish Research Council grant dnr: 2019–04098. The authors would like to thank the anonymous reviewers for their valuable comments and suggestions which helped to improve the quality of the paper.

Appendix A. Preliminaries on manifolds and explicit bounds on Gamma functions

A.1. Integration on Ω

We will integrate scalar-valued functions, $f : \Omega \rightarrow \mathbb{R}$, over Ω . When formulating integration of scalar-valued functions over submanifolds of \mathbb{R}^N , we follow the approach in [19]. Because we need some preliminary results concerning integration on Ω , we make some important definitions explicit.

First, let x_1, \dots, x_k be vectors in \mathbb{R}^N for $k \leq N$. If $I = (i_1, i_2, \dots, i_k)$ is a k -tuple of integers such that $i_1 \leq i_2 \leq \dots \leq i_k$, define $X_I \in \mathbb{R}^{k \times k}$ as the $k \times k$ matrix containing only rows i_1, \dots, i_k of the matrix $X = (x_1, \dots, x_k)$. Now we can define the volume function $V : \mathbb{R}^{N \times k} \rightarrow \mathbb{R}$, by $V(X) = \sqrt{\det^2(X^T X)} = [\sum_I \det^2 X_I]^{1/2}$, where the I 's span over k -tuples as above, see [19, Theorem 21.4].

In general, given a coordinate chart $\alpha : U \rightarrow W$, where $U \subset \mathbb{R}^d$, $W \subset \Omega_i \subset \mathbb{R}^N$ are open subsets, and $D\alpha$ is the Jacobian of α , we can express integration over W as

$$\int_W f \, dV = \int_U f \circ \alpha \, V(D\alpha).$$

In the coming proofs, when integrating around a point $x \in \text{Int } \Omega_i$, we will change coordinates to the standard basis in $T_{\Omega_i, x} = \mathbb{R}^k$. With this we mean that we can find open sets $W \subset \Omega_i$ around x such that the projection map $\pi : W \rightarrow B \subset x + T_{\Omega_i, x}$ is a diffeomorphism, where $x + T_{\Omega_i, x} := \{x + y \mid y \in T_{\Omega_i, x}\}$. To integrate over $T_{\Omega_i, x}$ we use the map π^{-1} precomposed with an inclusion map.

More specifically and without loss of generality, by translation and an orthonormal coordinate change, we can assume that $T_{\Omega_i, x} = \mathbb{R}^d \times \{0\}^{n-d}$. In this coordinate system we can write

$$\alpha : U \xrightarrow{i} x + T_{\Omega_i, x} \xrightarrow{\pi^{-1}} W \subset \Omega_i, \tag{A.1}$$

where i is the natural inclusion map and U an open subset in \mathbb{R}^k .

A.2. Important bounds

The following bounds will be used later in our proofs: First, let $T_{\Omega_i, x}, U, W, \pi$ be as in [Appendix A.1](#). Then for any $y \in W$,

$$\|y - \pi(y)\| \leq O(\|x - \pi(y)\|^2). \tag{A.2}$$

This follows since Ω_i is smooth and the tangent space represents a first order approximation of Ω_i around x .

To formulate the second bound, we need the lemma below.

Lemma A.1. *Let $U, W, x, y, \Omega_i, \pi, i, \alpha$ be as in [Appendix A.1](#). Then the following holds for the volume function V :*

$$V(D\alpha(y)) = 1 + O(\|x - \pi(y)\|^2).$$

Proof. Since $\alpha = \pi^{-1} \circ i$, and the tangent space is a first order approximation of Ω , we can parametrize the W by $\alpha(u) = (u, g(u))$. It is then easy to see that for $i = 1, \dots, d$ we have

$$\partial_i \alpha(y) = (e_i, \partial_i g(u)),$$

where $\partial_i g(0) = 0$ and $\|\partial_i g(u)\| = O(\|u\|)$. Now

$$\det D\alpha_I = \begin{cases} 1 & \text{if } I = (1, 2, \dots, d) \\ O(\|u\|) & \text{otherwise.} \end{cases}$$

If we Taylor expand $x \rightarrow \sqrt{x}$, we get

$$V(D\alpha) = \left(\sum_I (\det D\alpha_I)^2 \right)^{1/2} = 1 + O(\|u\|^2),$$

and by applying the above on $(u, 0) = x - \pi(y)$ we are finished. \square

Further, since we have a finite union $\Omega = \cup_i \Omega_i$, Ω_i is compact, Eq. (A.2), and the previous lemma implies that we can find a uniform bound L such that for all tuples $(U, W, x, y, \pi, \Omega_i)$

$$\|y - \pi(y)\| \leq L\|x - \pi(y)\|^2 \tag{A.3}$$

and

$$|V(D\alpha) - 1| \leq L\|x - \pi(y)\|^2 \tag{A.4}$$

holds.

A.3. Gamma functions

In the proofs of several of our results we will need to handle the *Gamma function* $\Gamma(\cdot)$, and both the *lower and upper incomplete gamma functions*, $\gamma(\cdot, \cdot)$ and $\Gamma(\cdot, \cdot)$ respectively. These are well-known and are defined by the equations

$$\Gamma(a) = \int_0^\infty t^{a-1} e^{-t} dt, \quad \gamma(a, x) = \int_0^x t^{a-1} e^{-t} dt, \quad \Gamma(a, x) = \int_x^\infty t^{a-1} e^{-t} dt.$$

In this paper both a and x are non-negative real numbers.

We will need the following bounds: First, if $a \geq 1$, then $t^{a-1} \geq x^{a-1}$ and

$$\Gamma(a, x) \geq x^{a-1} \int_x^\infty e^{-t} dt = x^{a-1} e^{-x}. \tag{A.5}$$

Secondly, if $e^x > 2^a$, then by [\[20, Theorem 4.4.3\]](#),

$$\Gamma(a, x) \leq ax^{a-1} e^{-x}. \tag{A.6}$$

Finally, we need the lower bound

$$\gamma(a, a) \geq \frac{1}{2} \Gamma(a). \tag{A.7}$$

That this holds can be seen by viewing $\gamma(a, x)$ as an unnormalized version of the cumulative distribution function of the Gamma distribution, for which it is well-known that the median v is less than a , [21].

Appendix B. Proof of Theorems 2 and 3

B.1. Proof of Theorem 2

Since $x \rightarrow L_t^i f(x)$ is translation and rotation invariant, we can without loss of generality assume that Ω_i oriented in \mathbb{R}^N in such a way which makes it a subset of $\mathbb{R}^d \times \{0\}^{N-d}$.

We want to evaluate

$$L_t^i f(x) = \int_{\Omega_i} K_t(x, y)(f(x) - f(y))p \, dy.$$

We begin by splitting the integral above into

$$\begin{aligned} \int_{\Omega_i} K_t(x, y)(f(x) - f(y))p \, dy &= \int_{B_R(x) \cap \Omega_i} K_t(x, y)(f(x) - f(y))p \, dy + \int_{\Omega_i \setminus B_R(x)} K_t(x, y)(f(x) - f(y))p \, dy \\ &= I_1 + I_2. \end{aligned} \tag{B.1}$$

For estimating I_2 , by translation invariance we can without loss of generality assume that $x = 0$. Now we make a change of variables and rescale y , which allows us to say that

$$\begin{aligned} |I_2| &= \left| \int_{\Omega_i \setminus B_R(0)} K_t(0, y)(f(0) - f(y))p \, dy \right| = \left| \int_{\Omega_i \setminus B_{r_0 \sqrt{t}}(0)} e^{-\|y\|^2/t} v \cdot (-y)p \, dy \right| = \left| \int_{\left(\frac{1}{\sqrt{t}}\Omega_i\right) \setminus B_{r_0}(0)} e^{-\|y\|^2} v \cdot (-y\sqrt{t})t^{d/2} p \, dy \right| \\ &\leq t^{d/2+1/2} \int_{\mathbb{R}^d \setminus B_{r_0}} e^{-\|y\|^2} \|y\| p \, dy. \end{aligned}$$

Now, by first changing to spherical coordinates and integrating out the angular parts, we deduce that

$$|I_2| \leq t^{d/2+1/2} \left| \mathbb{S}^{d-1} \right| p \int_{r_0}^{\infty} e^{-s^2} s^d \, ds = \frac{1}{2} p t^{d/2+1/2} \left| \mathbb{S}^{d-1} \right| \Gamma\left(\frac{d+1}{2}, r_0^2\right). \tag{B.2}$$

To finalize the bound of I_2 , we note that it follows from the assumption $t \leq \frac{R^2}{d/2+1}$ that $r_0^2 > \frac{d+1}{2}$, and we can use Eqs. (A.6) and (B.2) to conclude

$$|I_2| \leq B(x) t^{d/2+1/2} e^{-r_0^2}, \tag{B.3}$$

where $B(x)$ is some function such that

$$B(x) \leq \frac{d+1}{4} r_0^d p \left| \mathbb{S}^{d-1} \right|.$$

To bound I_1 , we use the following simple geometric fact:

$$\|x - y\|^2 = \|\hat{x} - y\|^2 + \|\hat{x} - x\|^2 = \|\hat{x} - y\|^2 + \sin^2 \theta r^2 t,$$

which implies that

$$e^{-\|x-y\|^2/t} = e^{-\sin^2 \theta r^2} e^{-\|\hat{x}-y\|^2/t}.$$

From the above we can conclude

$$\begin{aligned} I_1 &= e^{-\sin^2 \theta r^2} \int_{B_R(x) \cap \Omega_i} e^{-\|\hat{x}-y\|^2/t} v \cdot (x - y)p \, dy \\ &= e^{-\sin^2 \theta r^2} \left(\int_{B_R(x) \cap \Omega_i} e^{-\|\hat{x}-y\|^2/t} v \cdot (x - \hat{x})p \, dy + \int_{B_R(x) \cap \Omega_i} e^{-\|\hat{x}-y\|^2/t} v \cdot (\hat{x} - y)p \, dy \right) \\ &= e^{-r^2 \sin^2 \theta} (II + III). \end{aligned} \tag{B.4}$$

It is easier to integrate over a ball centered around \hat{x} , and to this end we define $\delta \geq 0$ by

$$\delta = \sqrt{R^2 - tr^2 \sin^2 \theta}. \tag{B.5}$$

Then since \hat{x} is the orthogonal projection of x , we have that $B_R(x) \cap \Omega_i = B_\delta(\hat{x}) \cap \Omega_i$.

Let us focus on II : We use Eq. (B.5) and change to spherical coordinates, which yields

$$\begin{aligned} II &= v \cdot (\hat{x} - x) t^{d/2} \int_{B_{\delta/\sqrt{t}}(\hat{x}) \cap \Omega_i} e^{-\|\hat{x}-y\|^2} p \, dy. \\ &= v \cdot (\hat{x} - x) t^{d/2} |\mathbb{S}^{d-1}| p \int_0^{\delta/\sqrt{t}} e^{-s^2} s^{d-1} \, ds = v \cdot (\hat{x} - x) t^{d/2} |\mathbb{S}^{d-1}| p \gamma(d/2, \delta^2/t) \\ &= \frac{1}{2} v \cdot \frac{\hat{x} - x}{\|\hat{x} - x\|} t^{d/2+1/2} r \sin \theta |\mathbb{S}^{d-1}| p \gamma(d/2, \delta^2/t). \end{aligned} \tag{B.6}$$

To estimate the RHS of Eq. (B.6) we will bound the γ -function from above and below: Using $r_0^2 \geq \frac{d+2}{2}$, $r < r_0/2$ and the definition of δ , we get

$$\frac{d}{2} \leq r_0^2 - \sin^2 \theta r^2 = \frac{\delta^2}{t}.$$

By Eq. (A.7) we now see that

$$\frac{1}{2} \Gamma(d/2) \leq \gamma(d/2, d/2) \leq \gamma(d/2, \delta^2/t). \tag{B.7}$$

Further, an application of Eq. (A.5) yields

$$\gamma(d/2, \delta^2/t) \leq \gamma(d/2, r_0^2) = \Gamma(d/2) - \Gamma(d/2, r_0^2) \leq \Gamma(d/2) - (r_0^2)^{d/2-1} e^{-r_0^2} = \Gamma(d/2) - r_0^{d-2} e^{-r_0^2}. \tag{B.8}$$

Now Eq. (B.6)–(B.8) together with $|\mathbb{S}^{d-1}| = \frac{2\pi^{d/2}}{\Gamma(d/2)}$ finally gives

$$II = A(d, r_0, \theta) v_{n, \Omega_i} t^{d/2+1/2} r \sin \theta,$$

where

$$\frac{1}{2} \max(p\pi^{d/2}, 2\pi^{d/2}p - p) |\mathbb{S}^{d-1}| r_0^{d-2} e^{-r_0^2} \leq A(d, r_0, \theta) \leq p\pi^{d/2}. \tag{B.9}$$

Finally, $III = 0$. This follows from that $B_R(x) \cap \partial\Omega_i = \emptyset$, the rotational symmetry of K , and the fact that the linear function is odd. Collecting Eq. (B.1), (B.3), (B.4) and (B.6) we get

$$L_t f(x) = t^{d/2+1/2} \left(A(d, r_0, \theta) v_{n, \Omega_i} \sin \theta r e^{-\sin^2 \theta r^2} + B(x) e^{-r_0^2} \right).$$

□

B.2. Proof of Theorem 3

We will follow the proof of Theorem 2 and modify where needed. Let I_2 , II and III be defined as in Eqs. (B.1) and (B.4). Then, since I_2 is bounded like in Eq. (B.3), we only need to find bounds for II and III .

Let δ be defined as in Eq. (B.5) and define $\delta_0 = \delta/\sqrt{t}$. Recall also the fact that $B_R(x) \cap \Omega_i = B_\delta(\hat{x}) \cap \Omega_i$. Now the difference in bounding II and III to the proof of Theorem 2 is that $B_\delta(\hat{x}) \cap \Omega_i$ is nonempty. Since, by assumption, $\partial\Omega_i$ is part of a $d-1$ -dimensional flat space, $B_\delta(\hat{x}) \cap \Omega_i$ is a d -dimensional ball, but missing a spherical cap.

We now use cylindrical coordinates (h, φ, θ) to describe the domain

$B_{\delta/\sqrt{t}}(\hat{x}) \cap \Omega_i$. In these new coordinates we are centered around \hat{x} , and (φ, θ) are coordinates for a $d-1$ -dimensional ball tangential to $\partial\Omega$, while the perpendicular coordinate h is oriented along the outwards normal of $\partial\Omega_i$. Let us denote this unit normal by $n_{\partial\Omega}$, and the projection of \hat{x} to $\partial\Omega$ by $\hat{x}_{\partial\Omega}$. We now set $K = (\hat{x} - \hat{x}_{\partial\Omega}) \cdot n_{\partial\Omega} = \sqrt{t}k_0$, where $-\delta_0 \leq k_0 \leq \delta_0$.

Then, with III defined in Eq. (B.4) we get

$$III = \int_{-\delta}^K \int_0^{\sqrt{\delta^2-h^2}} \int_{\mathbb{S}^{d-2}} K_t(\hat{x}, y) v \cdot (\hat{x} - y) \rho^{d-2} \, d\varphi \, d\theta \, dh.$$

We split v into a normal component $v_n = (v \cdot n_{\partial\Omega})n_{\partial\Omega}$ and a component $v_T = v - v_n$ which is tangential to the boundary $\partial\Omega$. Then, since the function $y \rightarrow v_T \cdot (\hat{x} - y)$ is odd as a function centered around \hat{x} , and the domain of integration is symmetric around \hat{x} , we know that the tangential component of III satisfies

$$III_T := \int_{-\delta}^K \int_0^{\sqrt{\delta^2-h^2}} \int_{\mathbb{S}^{d-2}} K_t(\hat{x}, y) v_T \cdot (\hat{x} - y) \rho^{d-2} \, d\varphi \, d\theta \, dh = 0.$$

By definition of $v_{n, \partial\Omega}$, we have that $v_n \cdot (\hat{x} - y) = v_{n, \partial\Omega}(n_{\partial\Omega} \cdot (\hat{x} - y)) = v_{n, \partial\Omega} h$, which implies that

$$\begin{aligned} III &= v_{n, \partial\Omega} \int_{-\delta}^K \int_0^{\sqrt{\delta^2-h^2}} \int_{\mathbb{S}^{d-2}} K_t(\hat{x}, y) h \rho^{d-2} \, d\varphi \, d\theta \, dh \\ &= v_{n, \partial\Omega} \int_{-\delta}^K \int_0^{\sqrt{\delta^2-h^2}} \int_{\mathbb{S}^{d-2}} e^{-h^2/t - \rho^2/t} h \rho^{d-2} \, d\varphi \, d\theta \, dh \end{aligned}$$

$$= t^{d/2} v_{n,\partial\Omega} \int_{-\delta_0}^{k_0} h e^{-h^2} \int_0^{\sqrt{\delta_0^2 - h^2}} \int_{\mathbb{S}^{d-2}} e^{-\rho^2} \rho^{d-2} d\varphi d\rho dh.$$

Continuing with the two inner integrals,

$$\int_0^{\sqrt{\delta_0^2 - h^2}} \int_{\mathbb{S}^{d-2}} e^{-\rho^2} \rho^{d-2} d\varphi d\rho = \frac{|\mathbb{S}^{d-2}|}{2} \int_0^{\delta_0^2 - h^2} e^{-s} s^{d/2-3/2} ds = \frac{|\mathbb{S}^{d-2}|}{2} \gamma\left(\frac{d-1}{2}, \delta_0^2 - h^2\right).$$

Using this expression in the full integral and applying partial integration in the second equality below yields

$$\begin{aligned} III &= t^{d/2} v_{n,\partial\Omega} \frac{|\mathbb{S}^{d-2}|}{2} \int_{-\delta_0}^{k_0} e^{-h^2} h \gamma\left(\frac{d-1}{2}, \delta_0^2 - h^2\right) dh \\ &= t^{d/2} v_{n,\partial\Omega} \frac{|\mathbb{S}^{d-2}|}{2} \left(\frac{1}{2} \left[-e^{-h^2} \gamma\left(\frac{d-1}{2}, \delta_0^2 - h^2\right) \right]_{-\delta_0}^{k_0} - \frac{1}{2} e^{-\delta_0^2} \int_{-\delta_0}^{k_0} (\delta_0^2 - h^2)^{(d-3)/2} h dh \right) \\ &= t^{d/2} v_{n,\partial\Omega} \frac{|\mathbb{S}^{d-2}|}{2} \left(\frac{1}{2} e^{-k_0^2} \gamma\left(\frac{d-1}{2}, \delta_0^2 - k_0^2\right) + e^{-\delta_0^2} \frac{(\delta_0^2 - k_0^2)^{(d-1)/2}}{d-1} \right). \end{aligned}$$

Thus, we know that

$$III = t^{d/2} v_{n,\partial\Omega} \frac{|\mathbb{S}^{d-2}|}{2} \left(e^{-\delta_0^2} \frac{(\delta_0^2 - k_0^2)^{(d-1)/2}}{d-1} + \frac{1}{2} e^{-k_0^2} \gamma\left(\frac{d-1}{2}, \delta_0^2 - k_0^2\right) \right). \tag{B.10}$$

We now address the integral *II* defined in Eq. (B.4), which means we need to calculate

$$J := \int_{B_R(x) \cap \Omega_i} e^{-\|x-y\|^2/t} p dy.$$

After a change cylindrical coordinates as for *III*, we rewrite this integral as

$$J = \int_{-\delta_0}^{k_0} e^{-h^2} \gamma\left(\frac{d-1}{2}, \delta_0^2 - h^2\right) dh.$$

We can immediately bound *J* from above by

$$\Gamma\left(\frac{d-1}{2}\right) \int_{-\delta_0}^{k_0} e^{-h^2} dh \leq \Gamma\left(\frac{d-1}{2}\right) \int_{-\infty}^{\infty} e^{-h^2} dx = \Gamma\left(\frac{d-1}{2}\right) \sqrt{\pi}. \tag{B.11}$$

Now we bound *J* from below: Since the integrand is positive, we can without loss of generality assume that $k_0 < 0$. Then a change of variables $h = -\sqrt{\delta_0^2 - y}$ yields that

$$J \geq e^{-\delta_0^2} \int_0^{\delta_0^2 - k_0^2} e^y \gamma\left(\frac{d-1}{2}, y\right) \frac{1}{2\sqrt{\delta_0^2 - y}} dy \geq e^{-\delta_0^2} \frac{1}{2\delta_0} \int_0^{\delta_0^2 - k_0^2} e^y \gamma\left(\frac{d-1}{2}, y\right) dy.$$

Using partial integration above we then get

$$J \geq \frac{e^{-\delta_0^2}}{2\delta} \left[e^y \gamma\left(\frac{d-1}{2}, y\right) - \frac{y^{\frac{d-1}{2}}}{\frac{d-1}{2}} \right]_0^{\delta_0^2 - k_0^2} = \frac{e^{-\delta_0^2}}{2\delta_0} \left(e^{\delta_0^2 - k_0^2} \gamma\left(\frac{d-1}{2}, \delta_0^2 - k_0^2\right) - \frac{2(\delta_0^2 - k_0^2)^{\frac{d-1}{2}}}{d-1} \right).$$

Simplifying further gives us

$$J \geq \frac{1}{2\delta_0} \left(e^{-k_0^2} \gamma\left(\frac{d-1}{2}, \delta_0^2 - k_0^2\right) - e^{-\delta_0^2} \frac{2(\delta_0^2 - k_0^2)^{\frac{d-1}{2}}}{d-1} \right). \tag{B.12}$$

Thus, equation Eq. (B.10) and the bounds in Eq. (B.12) and Eq. (B.11) proves the theorem. \square

Appendix C. Proofs of Section 5

C.1. Proof of Theorem 6

Set $R = 2$, let $r_0 = 2/\sqrt{t}$, and under $H_0 \partial\Omega_i \cap B_{4R}(x_0) = \emptyset$, we can apply Theorem 2 to get an expression for $L_t^i f(x)$, i.e. for any $x \in B_1(x_0)$ we have that

$$L_t^i f(x) = t^{\frac{d+1}{2}} B(x) e^{-1/t}.$$

Furthermore, under our assumption of H_0 , we have that

$$|L_t^j f(x)| = \left| \int_{\Omega_j} e^{-\|x-y\|^2/t} v \cdot (x-y) p(y) dy \right| \leq e^{-2/t}$$

whenever $j \neq i$. Now applying [Theorem 1](#) we can see that for

$$\varepsilon = \delta - \max \left\{ \sqrt{t} \frac{d+1}{2} \frac{|\mathbb{S}^{d-1}|}{2^d |\mathbb{S}^d|}, 1 \right\} e^{-2/t},$$

thus choosing $t \leq 1$ such that $\sqrt{t} \frac{d+1}{2} \frac{|\mathbb{S}^{d-1}|}{2^d |\mathbb{S}^d|} e^{-1/t} \leq \delta/2$ we get that

$$\mathbb{P}(\exists m : X_m \in B_R(x_0); |L_{n,t} f(X_m)| > \varepsilon + L_t f(X_m)) \leq 2n \exp \left(-\frac{4e(n-1)\varepsilon^2}{t} \right) \leq \alpha.$$

□

C.2. Proof of [Theorem 7](#)

Under hypothesis H'_1 we have from [Theorem 2](#) that for $r_0 = 1/\sqrt{t}$ that

$$L_t^i f(x) = t^{\frac{d+1}{2}} \left(A(\theta_i(x), r_0, d) v_{n,\Omega_i} r_i(x) \sin \theta_i(x) e^{-r_i(x)^2 \sin^2 \theta_i(x)} + B(x) e^{-r_0^2} \right),$$

whenever $x \in B_1(x_0)$. Also note that $\theta_i(x) = 0$ if $x \in \Omega_i$, see [Fig. 3](#).

The power of rejecting H_0 when H'_1 is true can be calculated by

$$\mathbb{P}(T > \delta).$$

Define the events

$$A := \{ \max_m |L_{n,t} f(X_m) - L_t f(X_m)| < \delta \},$$

$$B := \{ \max_m |L_t f(X_m)| \chi_{B_1(x_0)}(X_m) > 2\delta \}.$$

Then, for any index m satisfying B we have

$$\max_m |L_{n,t} f(X_m)| \geq \max_m (|L_t f(X_m)| - |L_{n,t} f(X_m) - L_t f(X_m)|) \geq \delta,$$

which gives together with the union bound that

$$\mathbb{P}(T > \delta) \geq \mathbb{P}(A \cap B) = 1 - \mathbb{P}(A^c \cup B^c) \geq 1 - \mathbb{P}(A^c) - \mathbb{P}(B^c).$$

We can bound $\mathbb{P}(A^c)$ immediately using [Theorem 1](#):

$$\mathbb{P}(\max_m |L_{n,t} f(X_m) - L_t f(X_m)| > \delta) \leq \alpha.$$

We next bound $\mathbb{P}(B^c)$, using [Theorem 2](#) and the definition of δ :

$$\begin{aligned} \mathbb{P}(B^c) &\leq \mathbb{P}(\max_m |L_t^1 f(X_m)| \chi_{B_1(x_0) \cap \Omega_2}(X_m) < 2\delta + L_t^2 f(X_m) \chi_{B_1(x_0) \cap \Omega_2}(X_m)) \\ &\leq \mathbb{P}(\max_m |L_t^1 f(X_m)| \chi_{B_1(x_0) \cap \Omega_2}(X_m) < 3\delta) \\ &= \mathbb{P}(|L_t^1 f(X_m)| \chi_{B_1(x_0) \cap \Omega_2}(X_m) < 3\delta)^n. \end{aligned}$$

We can now apply [Theorem 2](#) to get that for $x \in B_1(x_0) \cap \Omega_2$

$$|L_t^1 f(x)| \geq t^{\frac{d+1}{2}} 2\pi^{d/2} |v_{n,\Omega_1}| \frac{\|x - x_0\|}{\sqrt{t}} |\sin \theta_1(x)| e^{-\|x - x_0\|^2 \sin^2 \theta_1(x)/t} - \delta.$$

Denote $G(x) = xe^{-x^2}$, then $G^{-1}(y) = \sqrt{-W_0(-2y^2)/2}$ and $G^{-1}(y) = \sqrt{-W_{-1}(-2y^2)/2}$ are the two solutions to $G(x) = y$, and W is the Lambert W function and the subindex denotes the two different branches. Thus, we can estimate

$$\mathbb{P}(|L_t^1 f(X_m)| \chi_{B_1(x_0) \cap \Omega_2}(X_m) < 3\delta) \leq \mathbb{P}(G(\|X_m - x_0\| \sin \theta_1 / \sqrt{t}) \chi_{B_1(x_0) \cap \Omega_2} < \frac{\delta}{t^{\frac{d+1}{2}} C})$$

where $C = \pi^{d/2} |v_{n,\Omega_1}|/2$. Using the inverse of G as above we get the following terms

$$P_1 := \mathbb{P} \left(\left\| X_m - x_0 \right\| < \frac{\sqrt{-tW_0 \left(-2 \frac{\delta^2}{t^{d+1} C^2} \right)}}{\sqrt{2} \sin \theta_1}; X_m \in B_1(x_0) \cap \Omega_2 \right),$$

$$P_2 := \mathbb{P} \left(\left\| X_m - x_0 \right\| > \frac{\sqrt{-tW_{-1} \left(-2 \frac{\delta^2}{t^{d+1} C^2} \right)}}{\sqrt{2} \sin \theta_1}; X_m \in B_1(x_0) \cap \Omega_2 \right),$$

$$P_3 := \mathbb{P}(X_m \in (B_1(x_0) \cap \Omega_2)^c).$$

For P_1 we note that W_0 satisfies

$$-W_0(-\rho) < \epsilon\rho,$$

for $0 \leq \rho \leq 1$, as such

$$P_1 \leq \mathbb{P} \left(\|X_m - x_0\| < \frac{\epsilon\delta}{t^{d/2}C \sin \theta_1}; X_m \in \Omega_2 \right) \leq |\Omega_2| \left(\frac{\epsilon\delta}{t^{d/2}C \sin \theta_1} \right)^d.$$

Thus, if we require that

$$\frac{\epsilon\delta}{t^{d/2}C \sin \theta_1} \leq 1/4,$$

we have that $P_1 < \mathbb{P}(X_m \in B_{1/4}(x_0))$. Secondly, we know that

$$-W_{-1}(-\rho) > \log 1/\rho,$$

for $0 < \rho < 1$, and as such

$$P_2 \leq \mathbb{P} \left(\|X_m - x_0\| > \frac{\sqrt{t \log \left(\frac{t^{d+1}C^2}{2\delta^2} \right)}}{\sqrt{2} \sin \theta_1}; X_m \in B_1(x_0) \cap \Omega_2 \right).$$

Which if

$$t \left((d+1) \log(1/t) + \log \left(\frac{2^3}{C^2} \right) \right) \leq 2(1 - \sin^2 \theta_1),$$

is bounded as

$$P_2 \leq \mathbb{P}(X_m \in (B_1(x_0) \setminus B_{1/2}(x_0)) \cap \Omega_2).$$

In conclusion, we have that for large enough n , $P_1 + P_2 + P_3 < 1$. Assembling the above we have proven our theorem. \square

Appendix D. Proof of Theorems 4 and 5

D.1. Proof of Theorem 4

We begin by splitting up the domain Ω_i :

$$\begin{aligned} L_t^i f(x) &= \int_{\Omega_i} K_t(x, y)(f(x) - f(y))p \, dy = \int_{\Omega_i \cap B_1(x)} K_t(x, y)(f(x) - f(y))p \, dy + \int_{\Omega_i \setminus B_1(x)} K_t(x, y)(f(x) - f(y))p \, dy \\ &= I + II. \end{aligned} \tag{D.1}$$

We first note that

$$II = \int_{\Omega_i \setminus B_R(x)} K_t(x, y)(f(x) - f(y))p \, dy \leq e^{-R^2/t} \text{diam}(\Omega). \tag{D.2}$$

To estimate I we will make a change of variables to the tangent space at x_0 and use arguments similar to those in the proof of Theorem 2. Specifically, let $\pi : \Omega_i \cap B_R(x) \rightarrow T_{\Omega_i, x_0} \cap B_R(x)$ be the projection map, and $\alpha = \pi^{-1} \circ i : \mathbb{R}^d \cap B_R(0) \rightarrow \Omega_i \cap B_R(x)$ a coordinate chart as in (A.1). We will use α to integrate over T_{Ω_i, x_0} .

To simplify notation, we will use \hat{x} and \hat{y} to denote both $\pi(x), \pi(y) \in \mathbb{R}^N$, and sometimes implicitly assume the projection i^{-1} such that $\hat{x}, \hat{y} \in \mathbb{R}^d$. The space in which these points lie should be clear from context.

Before making the coordinate change, we find bounds relating $K(x, y)$ to $K(x, \hat{y})$: We recall that $K_t(x, y) = e^{-\frac{\|x-y\|^2}{t}}$, and by elementary considerations we get

$$e^{-\frac{\|x-\hat{y}\|^2}{t} - 3\frac{\|y-\hat{y}\|^2}{t}} \leq e^{-\frac{\|x-\hat{y}\|^2}{t}} \leq e^{-\frac{\|x-\hat{y}\|^2}{t} + 3\frac{\|y-\hat{y}\|^2}{t}}. \tag{D.3}$$

Since Ω_i is $(L, 2R)$ -regular, we use (A.3) and the fact that $y \in B_{2R}(x)$ to conclude

$$\|y - \hat{y}\| \leq L\|x_0 - \hat{y}\| \leq L4R^2,$$

which together with (D.3) yields

$$e^{-3(L4R^2)^2/t} K_t(x, \hat{y}) \leq K_t(x, y) \leq e^{3(L4R^2)^2/t} K_t(x, \hat{y}).$$

Furthermore, since $(L4R^2)^2/t \leq \frac{1}{4}$ by our assumptions, we have the bounds

$$e^{3(L4R^2)^2/t} \leq 1 + 4(L4R^2)^2/t \quad \text{and} \quad e^{-3(L4R^2)^2/t} \geq 1 - 3(L4R^2)^2/t.$$

Thus,

$$|K_t(x, y) - K_t(x, \hat{y})| \leq 4 \frac{(L4R^2)^2}{t} K_t(x, \hat{y}). \quad (D.4)$$

Replacing $K_t(x, y)$ with $K_t(x, \hat{y})$ in I we get

$$I = \int_{\Omega_t \cap B_R(x)} K_t(x, \hat{y})(f(x) - f(y))p \, dy + E_1, \quad (D.5)$$

and using Eq. (D.4) it holds that

$$|E_1| \leq 4 \frac{(L4R^2)^2}{t} R \left| \int_{\Omega_t \cap B_R(x)} K_t(x, \hat{y})p \, dy \right|. \quad (D.6)$$

We now decompose the integral of the first term in Eq. (D.5) as follows

$$\begin{aligned} \int_{\Omega_t \cap B_R(x)} K_t(x, \hat{y})(f(x) - f(y))p \, dy &= \int_{\Omega_t \cap B_R(x)} K_t(x, \hat{y})(f(x) - f(\hat{y}))p \, dy \\ &+ \int_{\Omega_t \cap B_R(x)} K_t(x, \hat{y})(f(\hat{y}) - f(y))p \, dy = I_1 + I_2 \end{aligned} \quad (D.7)$$

The quantity I_2 will be treated like an error term. Using Eq. (A.3) we see that

$$|I_2| \leq \int_{\Omega_t \cap B_R(x)} K_t(x, \hat{y})L\|\hat{y} - x_0\|^2 p \, dy.$$

Now we make a coordinate change with α and use the bound on the volume form in Eq. (A.4) to get

$$\begin{aligned} &\int_{\Omega_t \cap B_R(x)} K_t(x, \hat{y})L\|\hat{y} - x_0\|^2 p \, dy \\ &\leq LR^2 \int_{T_{\Omega_t, x_0} \cap B_R(x)} K_t(x, \hat{y})(1 + L\|x_0 - \hat{y}\|^2)p \, d\hat{y} \\ &\leq LR^2(1 + L4R^2) \int_{T_{\Omega_t, x_0} \cap B_R(x)} K_t(x, \hat{y})p \, d\hat{y}. \end{aligned}$$

The RHS of the above display can be handled similarly to Eq. (B.6), which means

$$|I_2| \leq LR^2(1 + L4R^2) \left| \mathbb{S}^{d-1} \right| t^{d/2} p \Gamma(d/2) = LR^2(1 + L4R^2) t^{d/2} 2p\pi^{d/2}.$$

We proceed now with I_1 from Eq. (D.7), which we want to estimate as accurately as possible. Using the coordinate change α and Eq. (A.4) we write

$$\begin{aligned} I_1 &= e^{-r^2 \sin^2 \theta} \int_{\Omega_t \cap B_R(x)} K_t(\hat{x}, \hat{y})(f(x) - f(\hat{y}))p \, dy \\ &= e^{-r^2 \sin^2 \theta} C_1(x) \int_{T_{\Omega_t, x_0} \cap B_R(x)} K_t(\hat{x}, \hat{y})(f(x) - f(\hat{y}))p \, d\hat{y}, \end{aligned} \quad (D.8)$$

where $C_1(x) > 0$ is such that $|C_1 - 1| \leq (1 + L4R^2)$.

The integral on the right in Eq. (D.8) is exactly II from Eq. (B.4), which we compute as in Eq. (B.6):

$$\int_{T_{\Omega_t, x_0} \cap B_R(x)} K_t(\hat{x}, \hat{y})(f(x) - f(\hat{y}))p \, d\hat{y} = A(d, r_0, \theta) v_{\Omega_t} t^{d/2+1/2} r \sin \theta, \quad (D.9)$$

where $A(d, r_0, \theta)$ is as in (B.9). Now, from Eqs. (D.7)–(D.9) we have

$$I_1 + I_2 = C_1(x)A(d, r_0, \theta)v_{\Omega_t} t^{d/2+1/2} r \sin \theta e^{-r^2 \sin^2(\theta)} + C_2(x)LR^2(1 + L4R^2)t^{d/2} 2p\pi^{d/2},$$

for a function $|C_2(x)| \leq 1$. This combined with the split in (D.5) and (D.6) gives us for another function $|C_3(x)| \leq 1$

$$\begin{aligned} I &= I_1 + I_2 + E_1 = C_1(x)A(d, r_0, \theta)v_{\Omega_t} t^{d/2+1/2} r \sin \theta e^{-r^2 \sin^2(\theta)} \\ &+ C_3(x)C_{L,R} t^{d/2} 2p\pi^{d/2} \end{aligned}$$

where

$$C_{L,R} := 4 \frac{(L4R)^4}{t} R + LR^2(1 + L4R^2).$$

Finally, the above and (D.2) finishes the proof. \square

D.2. Proof of Lemma 4.2

First applying Theorem 4 $L_t^i f(x)$, and then using the $(L, 2R)$ regularity of Ω_t , we use Eq. (A.3) to get

$$|r \sin \theta| = \frac{\|x - \hat{x}\|}{\sqrt{t}} \leq \frac{LR^2}{\sqrt{t}} < \frac{1}{2^4}.$$

As such we have that $e^{-2^{-8}} |r \sin \theta| \leq |r \sin \theta| e^{-r^2 \sin^2 \theta} \leq |r \sin \theta|$. Now applying [Theorem 4](#) the conclusion follows. \square

D.3. Proof of [Theorem 5](#)

To simplify notation, let $h_j = x - X_j$.

$$\mathbb{E}_\varepsilon L_{n,t} f(x) = \frac{1}{n} \sum_{j=1}^n \mathbb{E}_\varepsilon K_t(x, X_j) (f(x) - f(X_j + \varepsilon_j)) = \frac{1}{n} \sum_{j=1}^n \mathbb{E}_\varepsilon e^{-\|h_j - \varepsilon_j\|^2/t} v \cdot (h_j - \varepsilon_j) \quad (\text{D.10})$$

Let us compute a single term in the sum in [Eq. \(D.10\)](#): Since the expectation is w.r.t ε we can treat $h = h_j$ as fixed, and then algebraic manipulations give us for $z \sim \mathcal{N}(0, \sigma^2 I)$

$$\begin{aligned} & \mathbb{E}_z e^{-\|h+z\|^2/t} v \cdot (h+z) \\ &= (2\pi\sigma^2)^{-N/2} \int_{\mathbb{R}^N} e^{-\|h+z\|^2/t} e^{-\|z\|^2/(2\sigma^2)} v \cdot (h+z) dz \\ &= (2\pi\sigma^2)^{-N/2} \int_{\mathbb{R}^N} e^{-(\|h\|^2/t + 2\langle h, z \rangle/t + \|z\|^2/t + \|z\|^2/(2\sigma^2))} v \cdot (h+z) dz \\ &= (2\pi\sigma^2)^{-N/2} e^{-\frac{\|h\|^2}{2\sigma^2+t}} \int_{\mathbb{R}^N} e^{-\frac{1}{k} \|z+kh\|^2} v \cdot (h+z) dz. \end{aligned}$$

In the second to last step we completed the square and used $k = \frac{2\sigma^2}{2\sigma^2+t}$. This last integral can be viewed as the expectation

$$(\pi kt)^{-N/2} \int_{\mathbb{R}^N} e^{-\frac{1}{k} \|z+kh\|^2} v \cdot (h+z) dz = \mathbb{E}_X [(h+X)] = (1-k)v \cdot h$$

where $X \sim \mathcal{N}(-kh, I \frac{kt}{2})$. Then we can conclude that

$$\mathbb{E}_z e^{-\|h+z\|^2/t} v \cdot (h+z) = v \cdot h \frac{t^{N/2+1}}{(2\sigma^2+t)^{N/2+1}} e^{-\frac{\|h\|^2}{2\sigma^2+t}}.$$

\square

Appendix E. Proof of [Theorem 1](#)

First, using the union bound we get

$$P\left(\max_i \left|L_{n,t} f(X_i) - \frac{n-1}{n} L_t f(X_i)\right| > \varepsilon\right) \leq \sum_{i=1}^n P\left(\left|L_{n,t} f(X_i) - \frac{n-1}{n} L_t f(X_i)\right| > \varepsilon\right). \quad (\text{E.1})$$

Using the definitions of $L_{n,t}$ and L_t , see [Eqs. \(3.1\)](#) and [\(3.2\)](#), and using that the random variables X_1, \dots, X_N are i.i.d., we can replace each X_i by X_1 in each term in the summand of [Eq. \(E.1\)](#). Let Z be an independent copy of X_1 . Then each summand in [Eq. \(E.1\)](#) equals

$$P\left(\left|\frac{1}{n} \sum_{j=1}^n K_t(X_1, X_j) (f(X_1) - f(X_j)) - \frac{n-1}{n} \mathbb{E}_Z [K_t(X_1, Z) (f(X_1) - f(Z))]\right| > \varepsilon\right). \quad (\text{E.2})$$

To simplify notation, we denote

$$W_i(x) = K_t(x, X_i) (f(x) - f(X_i)) \quad \text{and} \quad Y_i(x) = W_i(x) - \mathbb{E}_{X_i} [W_i(x)].$$

We now rewrite [Eq. \(E.2\)](#) as

$$P\left(\left|\frac{1}{n-1} \sum_{i=2}^n Y_i(X_1)\right| > \frac{n}{n-1} \varepsilon\right).$$

Now by the tower property we have that

$$P\left(\left|\frac{1}{n-1} \sum_{i=2}^n Y_i(X_1)\right| > \frac{n}{n-1} \varepsilon\right) = \mathbb{E} \left[P\left(\left|\frac{1}{n-1} \sum_{i=2}^n Y_i(X_1)\right| > \frac{n}{n-1} \varepsilon \mid X_1\right) \right]$$

In order to use Hoeffding's inequality we need to show that $Y_i(x)$ is a bounded random variable for all $x \in \Omega$. First,

$$\begin{aligned} W_i(x) &= K_t(x, X_i) (f(x) - f(X_i)) = e^{-\|x-X_i\|^2/t} v \cdot (x - X_i) \leq e^{-\|x-X_i\|^2/t} \|x - X_i\| \\ &\leq \sup_{\rho} e^{-\rho^2} \rho \sqrt{t} \leq \sqrt{\frac{t}{2e}}. \end{aligned}$$

Now Hoeffding's inequality states that (where $C_n = \frac{n}{n-1}$)

$$\mathbb{P}\left(\left|\frac{1}{n-1} \sum_{i=2}^n Y_i\right| > C_n \varepsilon \mid X_1\right) \leq 2 \exp\left(-\frac{4e(n-1)C_n^2 \varepsilon^2}{t}\right)$$

and the proof is complete after taking expectations. \square

Appendix F. Construction of the neural network dataset

As explained in Section 6.1.1, we consider neural networks of the form given in Section 2.1, where $k = 3$ and $W, W^* \in \mathbb{R}^6$, $a_1 = -a_2 = -a_3$, all three weights w_i^* are the same. We will also set $\mathbf{B} = [-10, 10]^6$.

In order to create our data-set we begin by setting up a regression problem, i.e. we sample $D = \{x_1, \dots, x_{100}\}$ from the uniform density on the unit circle in \mathbb{R}^2 , and our data pairs for the regression is then $(x_i, f_{W^*}(x_i))$. We are interested in the zero set of the neural network with respect to the dataset $(x_i, f_{W^*}(x_i))$, i.e.

$$\hat{\Omega}_\delta := \{W \in \mathbf{B} : |f_W(x_i) - f_{W^*}(x_i)| < \delta, \quad x_i \in D\}.$$

To try to limit computations, we use interval constraint propagation, and more specifically forward–backward propagation in the computational tree of the expressions $|f_W(x_i) - f_{W^*}(x_i)| < \delta$, where $\delta \approx 10^{-16}$ and $x_i \in D$ to produce a paving (a true covering of boxes). For more information about these methods, see [22–24]. The resulting paving is a set of boxes \mathbb{X}^+ , where each box in \mathbb{X}^+ has a maximum width of 0.01. We will again reiterate here that the choice of $k = 3$ and box-size 0.01 is what we found to be the largest k and smallest box size that still allowed us to compute the zero set in a reasonable time, and we ended up with roughly 10^7 boxes in \mathbb{X}^+ .

To construct a point of the zero set we will simply choose the centroid of each box as a representative. By construction, we know that such a centroid is at most $0.005\sqrt{6}$ distance away from a point in $\hat{\Omega}_\delta$. In our experiment, this leads to a dataset of 12,753,597 points X that are within a small distance ($0.005\sqrt{6}$) of $\hat{\Omega}_\delta$ such that if we put a box around each with side-length 0.01, then the union of these boxes will cover $\hat{\Omega}_\delta$. This is in contrast with other search methods, like a grid search (which would require $\sim 10^{19}$ boxes) or a random search (hitting probability is 10^{-13}) that cannot guarantee the coverage of $\hat{\Omega}_\delta$.

References

- [1] Belkin M, Niyogi P. Towards a theoretical foundation for Laplacian-based manifold methods. *J Comput System Sci* 2008;74(8):1289–308.
- [2] Belkin M, Que Q, Wang Y, Zhou X. Toward understanding complex spaces: Graph Laplacians on manifolds with singularities and boundaries. In: *Proceedings of the 25th Annual Conference on Learning Theory*. 23, 2012, p. 36.1–26.
- [3] Belkin M, Niyogi P. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Adv Neural Inf Process Syst* 2001;14.
- [4] Kannan R, Vempala S, Vetta A. On clusterings: Good, bad and spectral. *J ACM* 2004;51(3):497–515.
- [5] Luxburg Uv, Bousquet O, Belkin M. On the convergence of spectral clustering on random samples: the normalized case. In: *International conference on computational learning theory*. Springer; 2004, p. 457–71.
- [6] Shi J, Malik J. Normalized cuts and image segmentation. *IEEE Trans Pattern Anal Mach Intell* 2000;22(8):888–905.
- [7] Ng A, Jordan M, Weiss Y. On spectral clustering: Analysis and an algorithm. *Adv Neural Inf Process Syst* 2001;14.
- [8] Belkin M, Niyogi P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput* 2003;15(6):1373–96.
- [9] Nadler B, Lafon S, Coifman RR, Kevrekidis IG. Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. *Appl Comput Harmon Anal* 2006;21(1):113–27.
- [10] Belkin M, Niyogi P. Semi-supervised learning on Riemannian manifolds. *Mach Learn* 2004;56(1):209–39.
- [11] Belkin M, Niyogi P. Convergence of Laplacian eigenmaps. *Adv Neural Inf Process Syst* 2006;19.
- [12] Bousquet O, Chapelle O, Hein M. Measure based regularization. *Adv Neural Inf Process Syst* 2003;16.
- [13] Lafon SS. *Diffusion maps and geometric harmonics*. Yale University; 2004.
- [14] Vidal R, Ma Y, Sastry S. Generalized principal component analysis (GPCA). *IEEE Trans Pattern Anal Mach Intell* 2005;27(12):1945–59.
- [15] Chen G, Lerman G. Foundations of a multi-way spectral clustering framework for hybrid linear modeling. *Found Comput Math* 2009;9(5):517–58.
- [16] Arora S, Du S, Hu W, Li Z, Wang R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In: *International conference on machine learning*. PMLR; 2019, p. 322–32.
- [17] Du SS, Zhai X, Poczos B, Singh A. Gradient descent provably optimizes over-parameterized neural networks. 2018, arXiv preprint [arXiv:1810.02054](https://arxiv.org/abs/1810.02054).
- [18] Avelin B, Julin V. Approximation of BV functions by neural networks: A regularity theory approach. *Anal Appl* 0; 0(0): 1–51. <http://dx.doi.org/10.1142/S0219530525500046>.
- [19] Munkres JR. *Analysis on manifolds*. CRC Press; 2018.
- [20] Gabcke W. *Neue Herleitung und explizite Restabschätzung der Riemann-Siegel-Formel* (Ph.D. thesis), Georg-August-Universität Göttingen; 1979.
- [21] Chen J, Rubin H. Bounds for the difference between median and mean of gamma and Poisson distributions. *Statist Probab Lett* 1986;4(6):281–3.
- [22] Jaulin L, Desrochers B. Introduction to the algebra of separators with application to path planning. *Eng Appl Artif Intell* 2014;33:141–7.
- [23] Jaulin L, Kieffer M, Braems I, Walter E. Guaranteed non-linear estimation using constraint propagation on sets. *Internat J Control* 2001;74(18):1772–82.
- [24] Sanders DP, Benet L, lucaferanti, Agarwal K, Richard B, Grawitter J, et al. *JuliaIntervals/IntervalArithmetic.jl: v0.20.3*. 2022.