

Consistency of Augmentation Graph and Network Approximability in Contrastive Learning

Chenghui Li

*Department of Statistics, University of Wisconsin-Madison
Wisconsin, USA*

CLI539@WISC.EDU

A. Martina Neuman

*Faculty of Mathematics, University of Vienna
Vienna, Austria*

NEUMANA53@UNIVIE.AC.AT

Editor: Brian Kulis

Abstract

Contrastive learning leverages data augmentation to develop feature representation without relying on large labeled data sets. However, despite its empirical success, the theoretical foundations of contrastive learning remain incomplete, with many essential guarantees left unaddressed, particularly the realizability assumption concerning neural approximability of an optimal spectral contrastive loss solution. In this work, we overcome these limitations by analyzing pointwise and spectral consistency of the augmentation graph Laplacian. We establish that, under specific conditions for data generation and graph connectivity, as the augmented data set size increases, the augmentation graph Laplacian converges to a weighted Laplace-Beltrami operator on the natural data manifold. These consistency results ensure that the graph Laplacian spectrum effectively captures the manifold geometry. Consequently, they give way to a robust framework for establishing neural approximability, directly resolving the realizability assumption in a current paradigm.

Key words: contrastive learning, manifold learning, augmentation graph, spectral convergence, pointwise consistency, approximation theory

1. Introduction

A major challenge in machine learning is the high cost of collecting large, labeled image data sets, which are essential for the success of many algorithms. To address this, data augmentation techniques have been widely utilized, as noted in (Mumuni and Mumuni, 2022; Shorten et al., 2021). In computer vision, for example, these techniques include transformations such as rotations, flips, and alterations in lighting or color. As a result, augmented data not only diversify the training set but also significantly improve the representational power of neural networks, enabling robust generalization from otherwise limited data sets.

In self-supervised learning, augmentation plays a pivotal role by generating multiple views of the same image, thus providing the essential self-supervision signals that drive effective learning. Recent advances have significantly benefited from contrastive learning, a paradigm that harnesses data augmentation to drive feature representation development without the need for extensive labeled data sets. In computer vision, contrastive methods aim to maximize the similarity between positive pairs, such as augmented versions

of the same image, while minimizing the similarity between negative pairs, such as randomly sampled images. In text-image representation learning, contrastive methods exploit semantically aligned text-image pairs, using an objective function to capture the underlying semantic relationships. While empirically successful, a robust theoretical foundation for contrastive learning remains incomplete. Existing theoretical frameworks, including the original work (Saunshi et al., 2019), and (Tosh et al., 2021a,b), often assume conditional independence between positive pairs given the class label. However, this assumption is frequently violated in practice, particularly when highly correlated positive pairs, such as image augmentations, are employed. As observed by (HaoChen et al., 2021), the two views of the same natural image exhibit strong correlations that cannot be fully mitigated through conditioning on the label. As a result, these frameworks tend to fall short of explaining the practical success of contrastive learning.

(HaoChen et al., 2021) proposed a framework for (spectral) contrastive learning that does not rely on the conditional independence of positive pairs. They introduced the notion of an *augmentation graph*, in which edges between augmented data points are weighted according to the likelihood that those points originate from the same underlying natural data point (e.g., through Gaussian noise augmentation). For an illustration, see Figure 1 below. The central insight of their framework is that the leading eigenvectors of this graph encode its latent geometric and semantic structure. Consequently, the goal is to learn neural representations that approximate a spectral embedding of the augmentation graph. To achieve this, (HaoChen et al., 2021) introduced a spectral contrastive loss whose population minimizers correspond to a spectral embedding of the augmentation graph (see, e.g. (3.20)). In this way, contrastive learning is reformulated as (1) performing spectral embedding on the population-level augmentation graph, and the resulting embedding task can then be expressed directly as (2) an optimization problem over neural network representations. Both are equivalent interpretations of the same learning objective.

Despite its utility in various applications, the theoretical framework introduced by (HaoChen et al., 2021) relies on several assumptions that may be restrictive. First, it implicitly assumes (e.g., (HaoChen et al., 2021, Example 3.8)) that different classes or semantic categories are supported on disconnected geometric structures. For example, images of dogs are assumed to lie on one manifold, while images of cats lie on another. We refer to this as the *separability assumption*. Such an assumption underlies the derivation of downstream generalization error bounds within their spectral contrastive framework. However, this perspective can be limiting in practice. For example, visually similar dog breeds such as Huskies and Alaskan Malamutes may lie on the same overarching dog manifold. In such situations, the distinction between these subclasses arises from finer geometric features, such as class boundaries, rather than from manifold-level separation.

Second, the framework requires a *recoverability assumption* (HaoChen et al., 2021, Assumption 3.6). Informally, this means that the spectral embedding of the augmentation graph reliably reveals the underlying clustering structure of the data, so that distinct clusters can be identified from the leading eigenvectors of the graph. Yet distinct clusters may naturally share a common manifold, which highlights the tension between the recoverability assumption and the (strict) separability assumption discussed above.

Finally, the framework relies on a *realizability assumption* (HaoChen et al., 2021, Assumption 3.7), which can be paraphrased as




Figure 1: Illustration of the augmentation graph constructed from augmented data points generated by applying Gaussian perturbations to natural data points, shown here using a conceptual dog-image manifold containing two dog breeds. **The left cloud** contains augmentation points corresponding to Huskies. **The right cloud** contains augmentation points corresponding to Beagles. **Orange edges** represent strongly connected positive pairs arising from nearby augmentations of similar dog breeds (between two Huskies or two Beagles). **Green edges** represent weakly connected negative pairs between augmentation points corresponding to different dog breeds (between a Husky and a Beagle). The figure is intended only as a conceptual illustration of the augmentation graph and its connectivity structure.

“There exists a deep neural network capable of achieving a global minimum of the spectral contrastive loss on population data.”

In other words, the solution to the resulting spectral embedding is assumed to be representable by the chosen neural network architecture. However, establishing such a result implicitly requires knowledge of the regularity of the target solution function. Since this regularity has not been rigorously established, it remains difficult to determine the network capacity necessary to realize the target embedding, thus limiting the scope and strength of the theoretical guarantees.

To address these limitations, we propose a detailed study of pointwise and spectral consistency of the augmentation graph Laplacian. We show that the augmentation graph Laplacian can be leveraged to unveil the geometry of the underlying natural data manifold. Specifically, within a range of data generation and graph connectivity parameters (see e.g.

(2.1), (2.4) below respectively), the graph Laplacian converges to a continuum (weighted) Laplace-Beltrami operator on the manifold—i.e. the natural generalization of the Euclidean Laplacian, see (3.8)—thereby revealing its underlying structure. We give a meaning to this “passage to the continuum” next. We assume that the original natural data set consists of i.i.d. N samples drawn from an m -dimensional smooth, connected, compact manifold embedded in $\mathbb{R}^d$. From these natural data points, an augmented data set of size $n \geq N$ is generated. Each pair of augmented data points is connected by an edge whose weight is determined by their average shared similarity, up to a graph connectivity parameter ε . As n tends to infinity and ε approaches zero sufficiently slowly (compared to n), the augmentation graph Laplacian, when applied to a smooth test function within a sufficiently small neighborhood of the manifold, converges with high probability to the continuum operator acting on the test function evaluated over the projections of the augmented data. Then, we demonstrate spectral consistency by showing that the graph Laplacian eigenvectors converge to the corresponding weighted Laplace-Beltrami eigenfunctions, measured in an L^2 -type distance.

The eigenfunctions of a weighted Laplace-Beltrami operator are known to be regular (see, e.g., (Gilbarg and Trudinger, 1977, Section 6.3)). Consequently, an immediate implication of our consistency results is the effective relaxation of the recoverability assumption. This is because weak (strong) energetic connections between distinct (identical) augmented subgroups can be mapped to observable variations in the eigenfunctions. More significantly, our results also facilitate a relaxation of the realizability assumption. The regularity of the eigenfunctions ensures their approximation by a neural network with bounded complexity, and their proximity to the graph Laplacian eigenvectors further enables the same network to approximate these eigenvectors with high accuracy. This offers a concrete estimate of the neural network capacity necessary to replicate an optimal spectral contrastive loss solution and thus concludes the validity of the realizability assumption. Finally, our analysis characterizes a precise regime of data generation and graph connectivity parameters in which the spectral proximity necessary for neural approximability holds with high probability. By relating the graph Laplacian to the Laplace-Beltrami operator, this specification also pinpoints the conditions that enable the graph structure to faithfully capture the geometry of the underlying data manifold. Such detailed guarantees are absent in (HaoChen et al., 2021), as we illustrate empirically in Section 1.1.

We now outline our novel theoretical contributions, which are twofold. The primary contribution lies in the field of manifold learning, while the secondary, an application of the first, pertains to approximation theory.

Manifold Learning. In one aspect, our work falls within the domain of manifold learning, where the underlying geometry of the graphical data set, and its subsequent impact on the performance of machine learning algorithms, are explored through the consistency analysis of the corresponding graph Laplacians in the non-asymptotic, large-sample regime. However, a central assumption in the classical setting is that graphical data are i.i.d. samples drawn from a distribution μ supported on a low-dimensional manifold (see a wealth of works such as (Burago et al., 2014; García Trillos et al., 2020; Calder and García Trillos, 2022; Dunson et al., 2021; Wormell and Reich, 2021; García Trillos et al., 2023; Calder et al., 2022)). In (García Trillos et al., 2025), the authors established minimax rates for


approximating eigenfunctions of the weighted Laplace-Beltrami operator under this data-generation model, and demonstrated that graph-Laplacian-based methods achieve nearly minimax-optimal convergence rates. In (Chaubet and Divol, 2025), under the additional assumption that the manifold is known, the authors considered a weaker metric for measuring the approximation of PDE eigenfunctions and obtained corresponding minimax rates. In contrast, augmented graphical data points may deviate from the natural data manifold (e.g. through a Gaussian perturbation), preventing the direct application of traditional manifold learning techniques to analyze the augmentation graph Laplacian. Moreover, the augmented data points may exhibit dependencies. Therefore, ensuring a ratio of augmented to natural data points ($N \leq n \leq CN$) that guarantees sufficient independence for the desired concentration effects is crucial to the validity of our analysis. Additionally, existing research on graph Laplacian consistency has predominantly focused on specific graph constructions, often relying on *primary similarity measures* (e.g., the ε -proximity graph). Our augmentation graph, however, is based on the theoretical principles in (HaoChen et al., 2021), which employs a concept akin to a *secondary similarity measure* (see, e.g., the notion of shared nearest neighbors graph in (Ayad and Kamel, 2003)). Namely, given a natural data point x , we assume that the augmented data points $\bar{x}$ follow a distribution $\mathbb{P}(\bar{x}|x)$. Two augmented data points are considered similar if both are likely, on average, to be generated from the same natural data point. Accordingly, the weight between two augmented data points is determined by $\mathbb{E}_{x \sim \mu} [\mathbb{P}(\bar{x}|x)\mathbb{P}(\bar{x}'|x)]$, which reflects a second-order (secondary) similarity derived from the first-order similarities $\mathbb{P}(\bar{x}|x)$ and $\mathbb{P}(\bar{x}'|x)$. Thus, in this regard, our work makes a significant theoretical contribution to manifold learning by introducing novel methods to address data augmentation graphs. These methods, to the best of our knowledge, have not been previously presented. Our consistency results achieve the most optimal convergence rates available in the current literature. Explicitly, if the graph connectivity parameter ε scales like¹

$$C \left(\frac{\log n}{n} \right)^{\frac{1}{m+4}} \leq \varepsilon \ll 1, \quad (1.1)$$

then with high probability, both the pointwise and spectral convergence rates scale linearly in the graph connectivity parameter ε . Moreover, if $\varepsilon = C \left(\frac{\log n}{n} \right)^{\frac{1}{m+4}}$, then we obtain the convergence rate of $\mathcal{O}(n^{-\frac{1}{m+4}})$, up to a log factor. Since $N \leq n \leq CN$, all these estimates can be expressed in terms of the natural data set size N in place of the augmented data set size n .

Approximation Theory. Another, more application-oriented contribution of our work is a network approximation result for the target classifier function. This connects to the emerging approximation theory of spectral neural networks (see (Li et al., 2026)), which seeks to apply manifold learning to graph analysis. Building on the framework in (HaoChen et al., 2021), contrastive learning can be interpreted as a parametric variant of the spectral embedding algorithm applied to augmented graphical data sets. Consequently, the eigenvectors of the augmentation graph Laplacian encapsulate the spectral embedding information of a spectral contrastive loss solution, which corresponds to the target classifier function.

1. The constant C in (1.1) may be chosen to depend only on the geometry of the natural data manifold and the natural data sampling density; see Theorems 11, 12.



(a) When choosing $\varepsilon_p, \varepsilon_w, \varepsilon_n$ according to Table 1 with $\tau = 2, \eta = 1$.

(b) When choosing $\varepsilon_p = \varepsilon_w = \varepsilon$.

Figure 2: (a) When choosing $\varepsilon_p, \varepsilon_w, \varepsilon_n$ according to Table 1 with $\tau = 2, \eta = 1$, the first non-trivial eigenvectors of the augmentation graph Laplacian detects the bottleneck. (b) When choosing $\varepsilon_p = \varepsilon_w = \varepsilon$, the first non-trivial eigenvectors of the augmentation graph Laplacian cannot detect the bottleneck.

This perspective enables us to deduce neural approximability of the classifier by leveraging the spectral convergence of the augmentation graph Laplacian to the Laplace-Beltrami operator on the natural data manifold and the regularity of the Laplace-Beltrami eigenfunctions. Specifically, the regularity of the eigenfunctions allows them to be approximated by a neural network with bounded complexity, and their proximity to the corresponding graph Laplacian eigenvectors further guarantees that the neural network can similarly approximate the latter with high fidelity. We demonstrate that, in the scale (1.1), with high probability (*only* due to the randomness of data) the required network architecture has a bounded complexity outlined in (4.7), with a depth of $\mathcal{O}(\log(\varepsilon^{-1}) + \log d)$ and a width of $\mathcal{O}(\varepsilon^{-m} + d)$. Thus, if $\varepsilon = C\left(\frac{\log n}{n}\right)^{\frac{1}{m+4}}$, then the network depth scales at most logarithmically with the data size n (or N), while the network width scales sublinearly. This result is different from standard approximation theory (e.g., (Chen et al., 2022)) and memorization frameworks (e.g., (Yun et al., 2019)), as it requires a network to replicate the eigenvectors while also generalizing to the true eigenfunctions of the (weighted) Laplace-Beltrami operator.

1.1 Simulation

In this simulation, the natural data x_i are uniformly sampled from the dumbbell. We expect the first nontrivial eigenvectors of the augmentation graph Laplacian to reveal the bottleneck, enabling geometric recovery via k -means clustering when projected onto the dumbbell. This expectation follows from the fact that the first nontrivial eigenfunctions of the limit Laplace-Beltrami operator of a dumbbell capture its shape (see (3.8) and Remark 9

afterward). Specifically, we first generate 4000 samples i.i.d. uniformly from a 2-dimensional dumbbell. Our generation scheme for augmentation graph requires three parameters: a graph connectivity parameter $\varepsilon = n^{-\frac{1}{2+4}}$, and different choices of, an edge weight parameter ε_w and a data generation parameter ε_p . A detailed introduction of the latter parameters is given in Section 2; see also as summary in Table 1. We note that the connectivity parameter $\varepsilon = n^{-\frac{1}{2+4}}$ corresponds to the optimal rate for the manifold learning problem discovered in (Calder and García Trillos, 2022)—which is $\varepsilon = n^{-\frac{1}{m+4}}$ with $m = 2$ —up to a logarithmic factor. When $\varepsilon_p = \varepsilon_w = \varepsilon$, the setting reduces to that in (HaoChen et al., 2021), causing the manifold’s geometry to become ambiguous; see Figure 2b. In contrast, selecting the parameters as suggested in Table 1, the manifold geometry is recovered; see Figure 2a.

1.2 Outline

Our paper is organized as follows. In Section 2, we provide an overview of data augmentation (Section 2.1) as well as essential preliminaries, conventions, and assumptions (Section 2.2) that will be consistently employed throughout the paper. Section 3 is dedicated to describing the framework for each of the main results. Section 3.1 focuses on establishing prerequisites for pointwise convergence, and Section 3.2 on the framework for spectral convergence. In Section 3.3, we provide the setup for spectral contrastive learning. The main results are formally stated in Section 4. Section 5 presents numerical experiments to validate the choice of parameters that work for contrastive learning. Section 6 is devoted to the proof of Theorem 11, with the main proof given in Section 6.3. Section 7 is devoted to the proof of Theorem 12, with the main proof given in Section 7.2. Finally, Section 8 presents the proof of Theorem 13, a consequence of Theorem 12. The Appendix contains all supplementary results, including relevant concentration results and helpful geometry, as well as detailed calculations to support the main text.

2. Preliminary

2.1 Augmented Data and Augmentation Graph Laplacian

Let $\mathcal{M} \subset \mathbb{R}^d$ be an m -dimensional smooth, closed (compact), connected, orientable manifold without boundary. Let μ be a distribution on $\mathcal{M}$, and let $x_1, \dots, x_N \sim \mu$ be i.i.d. samples, referred to as natural data points. Let $\mathcal{X} := \{x_1, \dots, x_N\}$. The specific ordering x_i is immaterial; thus we refer to a sample as $x \in \mathcal{X}$. From $\mathcal{X}$, we generate an augmented data set $\bar{\mathcal{X}}$, where each element $\bar{x} \in \bar{\mathcal{X}}$ is an augmented datum. For each $x \in \mathcal{M}$, let $\mathbb{P}(\cdot|x)$ denote the conditional distribution that $\bar{x} \in \bar{\mathcal{X}}$ follows. We specify that $\mathbb{P}(\cdot|x)$ takes the form

$$\mathbb{P}(\bar{x}|x) = \pi_{\varepsilon_p}(\bar{x}|x) := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \exp\left(-\frac{\|x - \bar{x}\|^2}{2\varepsilon_p^2}\right), \quad (2.1)$$

for some $\varepsilon_p > 0$ to be chosen later. Here, $\|\cdot\|$ denotes the usual Euclidean vector norm in $\mathbb{R}^d$. Intuitively, $\mathbb{P}(\bar{x}|x)$ gives the likelihood of $\bar{x}$ being generated from $x \in \mathcal{X}$. Note that,

$$\int_{\mathbb{R}^d} \mathbb{P}(y|x) dy = \int_{\mathbb{R}^d} \pi_{\varepsilon_p}(y|x) dy = 1.$$

Therefore, we can write, $\bar{\mathcal{X}} := \{\bar{x} : \bar{x} \text{ is sampled from the density } \pi_{\varepsilon_p}(\cdot|x) \text{ for some } x \in \mathcal{X}\}$. The law (2.1) indicates that the majority of the generated data is not far from $\mathcal{M}$. Precisely, there exists $0 < r < \infty$ such that the event

$$\text{dist}(\bar{x}, \mathcal{M}) < r \quad (2.2)$$

holds with a high probability, where $\text{dist}(\bar{x}, \mathcal{M}) := \inf\{\|\bar{x} - z\| : z \in \mathcal{M}\}$. This essential fact is the content of Lemma 2 below.

Let $\#\bar{\mathcal{X}} = n \geq N$. Note that $\bar{\mathcal{X}}$ itself is not an independent set of random variables but only partially so. To ensure sufficient independence from the augmented data cloud for our analysis, we impose the following assumption, which is reasonably aligned with practical settings.

Assumption 1 *Let $C_0 \in \mathbb{N}$ be predetermined. We assume either of the following holds:*

1. *for each augmented data point $\bar{x}$, there exist at most C_0 other points $\bar{x}'$ dependent on $\bar{x}$;*
2. *for any natural data point x , at most C_0 augmented data points $\bar{x}$ are generated from x .*

Note that point (2) above, which is easier to implement in practice, implies point (1). Under Assumption 1, $N \leq n \leq C_0 N^2$. We construct an augmentation graph $G_{\bar{\mathcal{X}}}$ from $\bar{\mathcal{X}}$ as follows. For $\bar{x} \neq \bar{x}' \in \bar{\mathcal{X}}$, let

$$\omega(\bar{x}, \bar{x}') := (\varepsilon_w \sqrt{2\pi})^{2d} \mathbb{E}_{x \sim \mu} \left[\pi_{\varepsilon_w}(\bar{x}|x) \pi_{\varepsilon_w}(\bar{x}'|x) \right], \quad (2.3)$$

for some ε_w related to ε_p in (2.1), to be specified shortly. Let $\varepsilon > 0$. With $\omega(\bar{x}, \bar{x}') \in (0, 1)$ representing the connection strength between $\bar{x}$ and $\bar{x}'$, we define the augmentation graph $G_{\bar{\mathcal{X}}}$ to consist of a vertex set in $\bar{\mathcal{X}}$ and the edge set

$$\{(\bar{x}, \bar{x}') \in \bar{\mathcal{X}} \times \bar{\mathcal{X}} : 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon\}.$$

To the graph $G_{\bar{\mathcal{X}}}$, we associate an *unnormalized* graph Laplacian operator $\mathcal{L}_{\text{aug}}$ as follows³. For $f : \bar{\mathcal{X}} \rightarrow \mathbb{R}$, we let

$$\mathcal{L}_{\text{aug}} f(\bar{x}) := \frac{1}{n \varepsilon_w^m \varepsilon^{m+2}} \sum_{\bar{x}' \in \bar{\mathcal{X}} : 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} \omega(\bar{x}, \bar{x}') (f(\bar{x}) - f(\bar{x}')). \quad (2.4)$$

For the derivation of the associated graph-theoretic formulation of the graph Laplacian, see Section 3.3. It is evident from (2.4) that ε acts as a graph connectivity parameter. We now dictate the relationship among the parameters $\varepsilon_p, \varepsilon_w$ and ε . Let $\tau > 0, \eta > 0$. Set

$$\varepsilon_w := \varepsilon^\tau \quad \text{and} \quad \varepsilon_p := \varepsilon^{\tau+1} \eta^{1/d}.$$

The choice of $\varepsilon_p \ll \varepsilon_w$ is made to ensure consistency for $\mathcal{L}_{\text{aug}}$, as we shall see. In addition, it directly supports the validity of the claim (2.2), whose proof occupies the remainder of this section.

2. The condition $n \geq N$ reflects practical considerations, while $n \leq C_0 N$ ensures sufficient independence in the augmented data cloud.
 3. The subscript “aug” stands for “augmentation.”

Lemma 2 *Let $\varepsilon > 0$ and $0 < r < \infty$. Then with probability at least $1 - 2dn \exp(-r^2/(2d\varepsilon_p^2))$, $\text{dist}(\bar{x}, \mathcal{M}) < r$, for all $\bar{x} \in \bar{\mathcal{X}}$. In particular, letting $\varepsilon_n := \varepsilon^{\tau+1}$, the event*

$$\text{dist}(\bar{x}, \mathcal{M}) < \varepsilon_n, \quad \forall \bar{x} \in \bar{\mathcal{X}}, \quad (2.5)$$

with probability at least $1 - 2dn \exp(-\eta^{-2/d}/(2d))$.

Proof [Proof of Lemma 2] Conditioning on $x \in \mathcal{M}$, we let $\bar{x} \sim \pi_{\varepsilon_p}(\cdot|x)$. Then it follows from (2.1) that $(\bar{x})_l \sim \mathcal{N}((x)_l, \varepsilon_p^2)$ i.i.d, for every $l = 1, \dots, d$. By Hoeffding's inequality (Hoeffding, 1994),

$$\mathbb{P}(|(\bar{x})_l - (x)_l|^2 \geq t^2) = \mathbb{P}(|(\bar{x})_l - (x)_l| \geq t) \leq 2 \exp\left(-\frac{t^2}{2\varepsilon_p^2}\right), \quad (2.6)$$

for every $t > 0$. If $r^2 \leq \|\bar{x} - x\|^2 = \sum_{l=1}^d |(\bar{x})_l - (x)_l|^2$, then there must exist one l such that (2.6) holds with $t = r/\sqrt{d}$. Consequently,

$$\mathbb{P}(\|\bar{x} - x\|^2 \geq r^2) = \mathbb{P}(\|\bar{x} - x\| \geq r) \leq 2d \exp\left(-\frac{r^2}{2d\varepsilon_p^2}\right).$$

Given that $\|\bar{x} - x\| \geq \text{dist}(\bar{x}, \mathcal{M})$, and the continuity of the distance functions, we conclude

$$\mathbb{P}(\text{dist}(\bar{x}, \mathcal{M}) \geq r) \leq 2d \exp\left(-\frac{r^2}{2d\varepsilon_p^2}\right).$$

By performing a union bound over n points of $\bar{x} \in \bar{\mathcal{X}}$, we obtain uniformly for all such $\bar{x}$, $\text{dist}(\bar{x}, \mathcal{M}) < r$ with probability at least $1 - 2dn \exp(-r^2/(2d\varepsilon_p^2))$. Then (2.5) automatically follows from the selection $r = \varepsilon_n = \varepsilon^{\tau+1}$. $\blacksquare$

Remark 3 *The edge weight $\omega(\bar{x}, \bar{x}') = (\varepsilon_w \sqrt{2\pi})^{2d} \mathbb{E}_{x \sim \mu} [\pi_{\varepsilon_w}(\bar{x}|x) \pi_{\varepsilon_w}(\bar{x}'|x)]$, is an idealized population-level quantity modeling the average likelihood that the augmented data points $\bar{x}, \bar{x}'$ arise from the same underlying natural data point. In practice, since the sampling distribution μ is unknown, the expectation may be approximated empirically using the observed i.i.d. natural data points $X_1, \dots, X_N \sim \mu$, namely through the Monte Carlo approximation*

$$\frac{1}{N} \sum_{i=1}^N \pi_{\varepsilon_w}(\bar{x}|X_i) \pi_{\varepsilon_w}(\bar{x}'|X_i).$$

Because π_{ε_w} is a Gaussian density with variance ε_w^2 , the dominant contributions to such averages arise from natural data points lying simultaneously near both $\bar{x}, \bar{x}'$, within distances on the order of ε_w . Thus, the construction may be viewed heuristically as measuring the degree of shared local natural data between the two augmented data points, in a manner reminiscent of shared nearest neighbor constructions, see e.g. (Neuman, 2023). From a statistical perspective, there are therefore two levels of approximation. First, standard concentration results for empirical averages of i.i.d. random variables (e.g. Hoeffding-type

inequalities (Boucheron et al., 2013, Theorem 2.8), (Vershynin, 2018, Theorem 2.6.2)) justify replacing the population expectation by its empirical Monte Carlo approximation when N is sufficiently large. Second, known quantitative Gaussian tail estimates yield sharp exponential decay of contributions from sufficiently distant natural data points and may be used to control the additional approximation error from Gaussian localization. A detailed quantitative analysis of such estimation errors, however, lies beyond the scope of the present theoretical work.

Remark 4 While in this paper we specifically focus on the Gaussian augmentation graph, the framework is intended to model local perturbative augmentations more broadly. In particular, sufficiently small transformations, including translations $x \mapsto x + v$ with $\|v\| \ll 1$ and rotations $x \mapsto Rx$ with $R \approx I$, may be viewed as local perturbations of the underlying data cloud of small magnitude in the ambient space. From the viewpoint of our analysis, isotropic Gaussian perturbations may therefore be interpreted as a mathematically tractable generic local perturbation model probing the local geometry of the manifold in all directions simultaneously. Thus, isotropic noise augmentation does not primarily enforce invariance to a particular global transformation, but rather promotes local geometric stability around the data manifold. In this sense, the role of isotropic augmentation is less about imposing a specific symmetry and more about enriching the data cloud without substantially distorting the intrinsic geometry relevant to downstream tasks.

2.2 Basic Assumptions and Convention

Basic Notation. We write $\mathbb{R}_{\geq 0}$ and $\mathbb{R}_+$ to denote the sets of nonnegative and positive real numbers, respectively. We write 0 to denote either the scalar zero, the zero vector, or the zero function, with the distinction made clear from context. For a vector x , we write x^i to refer to its i th component. We use $\|\cdot\|$ to represent either the Euclidean vector norm (with the dimension also inferred from context) or the Frobenius matrix norm, and $|\cdot|$ to represent the absolute value of a scalar. Additionally, we use $\|\cdot\|$ with a subscript to specify a functional norm, such as $\|\cdot\|_{L^\infty(\mathcal{M})}$. In the case of a vector or a matrix, we also use $\|\cdot\|_0$ to specify the *nuclear norm*, counting the number of nonzero entries, and $\|\cdot\|_\infty$ to specify the ℓ^∞ -norm. These last two notations are used exclusively in Section 3.3, the statement of Theorem 13, and its proof in Section 8.

Basic Manifold Geometry Assumptions. We endow $\mathcal{M}$ the Riemannian structure induced by the ambient space $\mathbb{R}^d$ and let $d\mathcal{V}$ be the volume form of $\mathcal{M}$ with respect to the induced metric tensor. We assume that all the sectional curvatures of the closed manifold $\mathcal{M}$ are bounded above in absolute value by some $K < \infty$, the manifold reach $R < \infty$, and further, the manifold injectivity radius is bounded below by some $i_0 > 0$. We regard K, R, i_0 to be *intrinsic*⁴ values of the manifold learning problem.

We use $B^m(0, r)$ and $B^d(0, r)$ to denote an m -dimensional and d -dimensional centered Euclidean ball, respectively, and $\mathcal{B}(x, r)$ to denote a geodesic ball centered at x on $\mathcal{M}$, all with radius r . We write $d(x, y)$ to indicate the geodesic distance on $\mathcal{M}$. Note that, from


4. unrelated to the notions of “intrinsic” and “extrinsic” in differential geometry

(García Trillos et al., 2020, Proposition 2),

$$\|x - y\| \leq d(x, y) \leq \|x - y\| + \frac{8}{R^2} \|x - y\|^3, \quad (2.7)$$

whenever $\|x - y\| \leq R/2$, with the first inequality holding universally.

At $x \in \mathcal{M}$, the affine *tangent plane* to $\mathcal{M}$ is a translated isomorphic copy of $\mathbb{R}^m$. We distinguish between such physical plane $\tilde{T}_x\mathcal{M}$, which sits in $\mathbb{R}^d$, and its isomorphic copy, *the tangent space* $T_x\mathcal{M}$, which can be embedded in $\mathbb{R}^m$. Similarly, we distinguish between the affine *normal plane* $\tilde{N}_x\mathcal{M}$ of $\mathcal{M}$ at x in $\mathbb{R}^d$ and its translated isomorphic copy, *the normal space* $N_x\mathcal{M} \cong \mathbb{R}^{d-m}$. See Figure 3 below for an illustration.



(a) A two-dimensional sphere in $\mathbb{R}^3$ and one of its tangent planes colored in pink

(b) The Euclidean space $\mathbb{R}^2$ as the tangent space

Figure 3: A two-dimensional sphere, one of its tangent planes, and an isomorphic copy of the latter as the Euclidean space $\mathbb{R}^2$

We will further elaborate on these relationships in Appendix B, as they will then be necessary for our analysis.

If $0 < r < \min\{K^{-1/2}, i_0\}$, and $x \in \mathcal{M}$, then the exponential map

$$\text{Exp}_x : T_x\mathcal{M} \supset B^m(0, r) \rightarrow \mathcal{B}(x, r) \subset \mathcal{M}, \quad (2.8)$$

where $T_x\mathcal{M}$ denotes the tangent plane of $\mathcal{M}$ at x , is a diffeomorphism. Likewise, its inverse, which is the logarithm map

$$\text{Log}_x : \mathcal{M} \supset \mathcal{B}(x, r) \rightarrow B^m(0, r) \subset T_x\mathcal{M}, \quad (2.9)$$

is also a diffeomorphism.

Let $J_x(v)$ denote the Jacobian of Exp_x at $v \in B^m(0, r) \subset T_x\mathcal{M}$. Then it follows as a consequence of the Rauch comparison theorem that (do Carmo, 1992, Chapter 10, Theorem 2.3) (see also (Burago et al., 2014, Equation (2.5)))

$$(1 + CmK\|v\|^2)^{-1} \leq J_x(v) \leq 1 + CmK\|v\|^2, \quad (2.10)$$

which implies

$$|\mathcal{V}(\mathcal{B}(x, r)) - \alpha_m r^m| \leq C m K r^{m+2}, \quad (2.11)$$

where $\mathcal{V}(S) := \int_S d\mathcal{V}(x)$, for $S \subset \mathcal{M}$, and α_m is the volume of the m -dimensional unit ball.

From Manifold to Tangent Plane. The exponential and logarithm maps, (2.8) and (2.9) respectively, offer a framework for translating calculus from the manifold $\mathcal{M}$ to its tangent planes. We introduce a system of changing variables and modified notations aligned with this translation, which will be employed throughout.

Let $x \in \mathcal{M}$, and let $y \in \mathcal{B}(x, r)$ for some $0 < r < \min\{K^{-1/2}, i_0\}$. Let $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$, and let $f : \mathcal{M} \rightarrow \mathbb{R}$ be measurable. Then according to (2.8), there exists a unique $v \in B^m(0, r)$ such that $y = \text{Exp}_x(v)$, whence

$$g(d(x, y)) = g(\|v\|).$$

Further, we write

$$\tilde{f}(v) := f \circ \text{Exp}_x(v). \quad (2.12)$$

In this case, $\tilde{f} : B^m(0, r) \rightarrow \mathbb{R}$, and particularly, $\tilde{f}(0) = f(x)$.

Smooth Functions on $\mathcal{M}$. Note that (2.12) allows a function f on $\mathcal{M}$ to be expressed locally as a function $\tilde{f}$ on $B^m(0, r)$. We say that $f \in \mathcal{C}^i(\mathcal{M})$ if at every $x \in \mathcal{M}$ its local version satisfies $\tilde{f} \in \mathcal{C}^i(B^m(0, r))$. We write $\|f\|_{\mathcal{C}^i(\mathcal{M})} := \sum_{j=0}^i \|D^{(j)}f\|_{L^\infty(\mathcal{M})}$, where $D^{(j)}$ is the j th-differentiation operator. Then by expressing the covariant derivatives $D^{(j)}f$ in normal coordinates around x (Hebey, 1996, Chapter 2), and applying standard expansions for the Christoffel symbols and metric tensor in these coordinates, we can establish that

$$c\|\tilde{f}\|_{\mathcal{C}^i(B^m(0, r))} \leq \|f\|_{\mathcal{C}^i(\mathcal{B}(x, r))} \leq C\|\tilde{f}\|_{\mathcal{C}^i(B^m(0, r))}, \quad (2.13)$$

for some universal constants $C \geq c > 0$.

Basic Density Assumption. We have assumed at the beginning of Section 2.1 that the natural data $x_i \sim \mu$ i.i.d. for some distribution μ supported on $\mathcal{M}$. We posit further that $d\mu = \rho d\mathcal{V}$, where $d\mathcal{V}$ denotes the Riemannian volume form associated with $\mathcal{M}$, and $\rho : \mathcal{M} \rightarrow \mathbb{R}_+$ is a density function such that $\rho \in \mathcal{C}^2(\mathcal{M})$ and that for every $x \in \mathcal{M}$,

$$0 < \rho_{\min} \leq \rho(x) \leq \rho_{\max} < \infty. \quad (2.14)$$

Basic Assumptions on ε . We assume the following regarding the graph connectivity parameter ε in (2.4):

$$\varepsilon \leq \min \left\{ 1, K^{-1/2}, i_0, \frac{R}{2} \right\}. \quad (2.15)$$

The requirement $\varepsilon \leq \min\{K^{-1/2}, i_0\}$ is justified by the considerations (2.8), (2.9). Further, that $\varepsilon \leq \frac{R}{2}$ implies (2.7), along with

$$\mathbb{P} \left(\bigcup_{\bar{x} \in \bar{\mathcal{X}}} \{\text{dist}(\bar{x}, \mathcal{M}) \leq R\} \right) \geq \mathbb{P} \left(\bigcup_{\bar{x} \in \bar{\mathcal{X}}} \{\text{dist}(\bar{x}, \mathcal{M}) \leq \varepsilon_n\} \right) \geq 1 - 2dn \exp(-\eta^{-2/d}/(2d)),$$

according to Lemma 2. Throughout, we maintain the option to rescale ε to $C\varepsilon$, for some universal constant $C > 0$, and assume that $C\varepsilon$ continues to satisfy (2.15) in place of ε .

Analytic Constants and Their Dependence. We adhere to the convention that analytic majorant constants, denoted by C or c , may vary between occurrences, with $C \geq 1$ and $0 < c \leq 1$, implicitly. All constants are permitted to depend on at most ρ and the manifold $\mathcal{M}$. The former includes dependencies on $\rho_{\min}$, $\rho_{\max}$, $\|\rho\|_{\mathcal{C}^2(\mathcal{M})}$, for instances, while the latter involves the intrinsic values K, R, i_0 , the dimension m , as well as the quantity

$$\mathfrak{D} := \sup_{x, x' \in \mathcal{M}} \|x - x'\|. \quad (2.16)$$

For clarity, the dependencies of constants are always stated in result statements. When necessary, the constants are further distinguished in these statements. Specific dependencies, such as on η or a function, will be declared by a subscript. For instance, for the latter, $C_f > 0$ denotes a constant depending on f , in addition to $\mathcal{M}$, ρ at most, and the nature of such dependence will be announced. While most parametric dependencies can be made explicit, we do not pursue this.

The Notation $\mathcal{O}$. Let $a > 0$. We write $A = \mathcal{O}(a)$ to mean $|A| \leq Ca$ or $\|A\| \leq Ca$, depending on the nature of A , where $C > 0$ is either a universal constant or dependent on $\mathcal{M}$, ρ at most. For instance, if A, B are scalars, then $A = B + \mathcal{O}(a)$ means $|A - B| \leq Ca$. Since the non-asymptotic settings we consider involve small $\varepsilon > 0$, this notation convention does not introduce significant ambiguity. Further, a subscript will indicate a special dependence. For instance, $A = \mathcal{O}_f(a)$ means $|A| \leq C_f a$, where $C_f > 0$ depends on f .

3. Setup

In the following discussions, the basic assumptions from Section 2.2 concerning the geometry of $\mathcal{M}$, the regularity of the density ρ of the natural data, and, in particular, the microscopic range (2.15) of the parameter ε , will be understood throughout.

3.1 Setting for Pointwise Consistency

Recall that $\varepsilon_n = \varepsilon^{\tau+1}$, with $\tau > 0$. Recall further from Section 2.2 that $\tilde{N}_x \mathcal{M}$ denotes the affine normal plane of $\mathcal{M}$ at x in $\mathbb{R}^d$. We define

$$\mathcal{N}_{\varepsilon_n}(\mathcal{M}) := \{x + h : x \in \mathcal{M}, h \in \tilde{N}_x \mathcal{M} - x, \text{ and } \|h\| < \varepsilon_n\}$$

to be an ε_n -*tubular neighborhood* of $\mathcal{M}$. For $\bar{x} \in \bar{\mathcal{X}} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})$, we denote by $Q_{\bar{x}}$ the closest point on $\mathcal{M}$ to $\bar{x}$. I.e., $Q_{\bar{x}}$ is the orthogonal projection point of $\bar{x}$ onto $\mathcal{M}$, and $\|\bar{x} - Q_{\bar{x}}\| = \text{dist}(\bar{x}, \mathcal{M})$. By (2.15) and the definition of manifold reach, $Q_{\bar{x}}$ exists uniquely for every such $\bar{x}$. Further, recall from Lemma 2 that $\bar{\mathcal{X}} = \mathcal{X} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})$ with probability at least $1 - 2dn \exp(-\eta^{-2/d}/(2d))$. Hence, from this point onward, we assume the following.

Assumption 5 $\bar{\mathcal{X}} \subset \mathcal{N}_{\varepsilon_n}(\mathcal{M})$.

Moreover, we demonstrate that there exists a one-to-one correspondence between a point $\bar{x}$ and its projection $Q_{\bar{x}}$.

Lemma 6 *Let $\bar{x}, \bar{x}' \in \bar{\mathcal{X}}$ be such that $\bar{x} \neq \bar{x}'$. Let $Q_{\bar{x}}, Q_{\bar{x}'}$ denote the orthogonal projections of $\bar{x}, \bar{x}'$ onto $\mathcal{M}$, respectively. Then with probability 1, $Q_{\bar{x}} \neq Q_{\bar{x}'}$.*

Proof [Proof of Lemma 6] The event E_1 that $Q_{\bar{x}} = Q_{\bar{x}'}$ is equivalent to the event E_2 that both $\bar{x}, \bar{x}'$ belong to $\tilde{N}_x\mathcal{M}$ for some $x \in \mathcal{M}$. However, as $\tilde{N}_x\mathcal{M}$ is of dimension $d - m$, and not of full dimension d , $\mathbb{P}(E_2) = 0$. Integrating this probability over $x \sim \mu$, we obtain that $\mathbb{P}(E_1) = 0$ as well. $\blacksquare$

A summary of the relevant parameters and hyperparameters is given in Table 1.

Symbol	Meaning
η, τ	hyperparameters
ε	graph connectivity parameter
$\varepsilon_w = \varepsilon^\tau$	edge weight parameter
$\varepsilon_p = \eta^{1/d} \varepsilon^{\tau+1}$	variance parameter of the data-generating probability
$\varepsilon_n = \varepsilon^{\tau+1}$	neighborhood parameter of the augmented data

Table 1: Parameters/hyperparameters and their meanings. Note that $\tau, \eta > 0$ are treated as hyperparameters, i.e. given a priori and not subject to variation.

Under Assumption 5, an augmented point $\bar{x} \in \bar{\mathcal{X}}$ can be expressed as $z + h$, where $z \in \mathcal{M}$ and $h \in \tilde{N}_z\mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M}) - z$. Here, z is unique and represents the projected point $Q_{\bar{x}}$. Such points z are identically distributed according to a measure ν , which is supported on $\mathcal{M}$ and has a density

$$q(z) := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z\mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh. \quad (3.1)$$

Furthermore, the regularity of q can be established for a suitable range of ε .

Proposition 7 *Let $\tau > 0$, $\varepsilon > 0$, and $\varepsilon_n = \varepsilon^{\tau+1}$. Let $\varepsilon > 0$ be further sufficiently small, so that*

$$\exp\left(-\frac{c}{\varepsilon_n^{2/3}}\right) \leq \varepsilon_n^{m+2}. \quad (3.2)$$

Let q be given in (3.1). Then under Assumption 5, for any $z \in \mathcal{M}$, we have

$$|q(z) - \rho(z)| \leq C_\eta \varepsilon_n^{2/3}. \quad (3.3)$$

Moreover, we have $q \in \mathcal{C}^2(\mathcal{M})$, which means, specifically, for any $z \in \mathcal{M}$,

$$0 < c_\eta \leq q(z) \leq C_\eta < \infty, \quad (3.4)$$

and for any unit vectors $\tilde{u}, \tilde{w} \in \tilde{T}_z\mathcal{M}$,

$$|\partial_{\tilde{u}} q(z)| \leq C_\eta, \quad (3.5)$$

$$|\partial_{\tilde{w}} \partial_{\tilde{u}} q(z)| \leq C_\eta, \quad (3.6)$$

and finally, all the relevant partial derivatives are continuous in a neighborhood of z in $\mathcal{M}$. The constant $C_\eta > 0$ in (3.3) depends on η and $\mathcal{M}$. Then constants $C_\eta, c_\eta > 0$ in (3.4), (3.5), (3.6) depend on η and at most $\mathcal{M}, \rho$.

A detailed proof of Proposition 7 can be found in Appendix D. We briefly remark that (3.2) serves as a crucial condition to ensure the boundedness of the first and second derivatives of q in our proof, which relies heavily on a divide-and-conquer technique. However, in order to further guarantee the pointwise consistency of $\mathcal{L}_{\text{aug}}$, it will be necessary for us to enforce an additional condition, that is

$$\exp\left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}}\right) \leq \varepsilon_w^m \varepsilon^2. \quad (3.7)$$

It is also evident that $\tau > 2$ for (3.7) to hold. We summarize the key requirements for τ , η , ε as follows.

Assumption 8 (Parameter Assumption) *Let $\tau > 2$, $\eta > 0$, and let $\varepsilon > 0$ be sufficiently small as to satisfy (2.15), (3.2), (3.7).*

We are now ready to formulate the setting for our first main result.

Setting I *Let $\eta, \tau > 0$ and $\varepsilon > 0$. Let $\varepsilon_w, \varepsilon_p, \varepsilon_n$ be given in Table 1. Let $\mathcal{M} \subset \mathbb{R}^d$ be a closed manifold of dimension m , satisfying the fundamental geometry outlined in Section 2.2. Let $\mathcal{X} = \{x_1, \dots, x_N\}$ be the set of natural data, where $x_i \sim \mu$, whose density $\rho \in \mathcal{C}^2(\mathcal{M})$ satisfies (2.14). Let $\bar{\mathcal{X}} = \{\bar{x}_1, \dots, \bar{x}_n\}$ be the set of augmented data generated from $\mathcal{X}$ according to the law (2.1). Let Assumption 1 hold. Let $G_{\bar{\mathcal{X}}}$ be the augmentation graph on $\bar{\mathcal{X}}$ as described in Section 2.1, and let the associated augmentation graph Laplacian $\mathcal{L}_{\text{aug}}$ be given in (2.4). Let Assumptions 5, 8 hold. Let $q \in \mathcal{C}^2(\mathcal{M})$ be the density (3.1) of the measure ν to which the projections of the augmented data are distributed on $\mathcal{M}$. Finally, let Δ_{aug} be a Laplace-Beltrami operator defined as follows:*

$$\Delta_{\text{aug}} f(x) := -\frac{1}{2} \text{div}(q^2 \nabla f)(x), \quad (3.8)$$

for every $f \in \mathcal{C}^2(\mathcal{M})$ and $x \in \mathcal{M}$.

Remark 9 *It can be shown, by extending the techniques from the proof of Proposition 7, that q and ρ are also close in $\mathcal{C}^2(\mathcal{M})$ sense, i.e.*

$$\|q - \rho\|_{\mathcal{C}^2(\mathcal{M})} = \mathcal{O}_\eta(\varepsilon). \quad (3.9)$$

We choose to omit this due to the already extensive technical details involved. However, once (3.9) holds, then we can further rewrite the limit in (3.8) to $-\frac{1}{2} \text{div}(\rho^2 \nabla f)(x)$, making the natural data sampling density explicit in the expression.

3.2 Setting for Spectral Consistency

In this section, we demonstrate that the eigenvalues and eigenvectors of $\mathcal{L}_{\text{aug}}$ converge to the corresponding eigenvalues and eigenfunctions of the continuum operator Δ_{aug} . To lay the groundwork, we review the fundamental properties of a graph Laplacian and Laplace-Beltrami operator, with a specific focus on the respective examples of $\mathcal{L}_{\text{aug}}$ and Δ_{aug} .

Define on $\bar{\mathcal{X}} = \{\bar{x}_1, \dots, \bar{x}_n\}$ a discrete probability measure ϑ_n , that is

$$\vartheta_n := \frac{1}{n} \sum_{i=1}^n \delta_{\bar{x}_i}, \quad (3.10)$$

where δ denotes the Dirac delta measure. We then endow $L^2_{\vartheta_n}(\bar{\mathcal{X}})$ with an inner product

$$\langle f, g \rangle_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} := \frac{1}{n} \sum_{i=1}^n f(\bar{x}_i)g(\bar{x}_i).$$

It is known that $\mathcal{L}_{\text{aug}}$ is a positive semi-definite and self-adjoint operator on $L^2_{\vartheta_n}(\bar{\mathcal{X}})$ (von Luxburg, 2007). Particularly, we can list its eigenvalues in non-decreasing order (repeated according to multiplicity) as follows:

$$0 = \hat{\lambda}_1^{\text{aug}} \leq \hat{\lambda}_2^{\text{aug}} \leq \dots \leq \hat{\lambda}_n^{\text{aug}}. \quad (3.11)$$

An orthonormal basis for $L^2_{\vartheta_n}(\bar{\mathcal{X}})$ can be constructed from their associated eigenvectors, denoted $\hat{f}_1^{\text{aug}}, \dots, \hat{f}_n^{\text{aug}}$, where $\mathcal{L}_{\text{aug}} \hat{f}_l^{\text{aug}} = \hat{\lambda}_l^{\text{aug}} \hat{f}_l^{\text{aug}}$. Furthermore, the eigenvalues in (3.11) can also be characterized variationally according to the Courant-Fisher min-max principle:

$$\hat{\lambda}_l^{\text{aug}} = \min_{S \in \mathbb{G}_{\bar{\mathcal{X}}}^l} \max_{f \in S \setminus \{0\}} \frac{E_{\bar{\mathcal{X}}}(f)}{\|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2}. \quad (3.12)$$

Here, $\mathbb{G}_{\bar{\mathcal{X}}}^l$ denotes the Grassmannian manifold of all linear subspaces of $L^2_{\vartheta_n}(\bar{\mathcal{X}})$ of dimension l , and $E_{\bar{\mathcal{X}}}$ represents a *graph-based Dirichlet energy*, defined such that for $f \in L^2_{\vartheta_n}(\bar{\mathcal{X}})$:

$$E_{\bar{\mathcal{X}}}(f) := \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i,j=1}^n \omega(\bar{x}_i, \bar{x}_j) \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} (f(\bar{x}_i) - f(\bar{x}_j))^2 = 2 \langle \mathcal{L}_{\text{aug}} f, f \rangle_{\ell^2(\bar{\mathcal{X}})}, \quad (3.13)$$

with $\mathbb{1}$ being the indicator function. Thus, for each orthonormal eigenvector $\hat{f}_l$, its graph-based Dirichlet energy $E_{\bar{\mathcal{X}}}(\hat{f}_l)$ is given by

$$E_{\bar{\mathcal{X}}}(\hat{f}_l) = \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i,j=1}^n \omega(\bar{x}_i, \bar{x}_j) \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} (\hat{f}_l(\bar{x}_i) - \hat{f}_l(\bar{x}_j))^2.$$

Continuing, we endow $L^2(\mathcal{M})$ with an inner product

$$\langle f, g \rangle_{L^2(\mathcal{M})} := \int_{\mathcal{M}} f(x)g(x) d\mathcal{V}(x).$$

Next, recall the Laplace-Beltrami operator Δ_{aug} defined in (3.8). Then the operator Δ_{aug} is a positive semi-definite and self-adjoint operator on $L^2(\mathcal{M})$. Consequently, it has a well-defined point spectrum, allowing its eigenvalues to be listed in increasing order (repeated according to multiplicity) as

$$0 \leq \lambda_1 \leq \lambda_2 \leq \dots, \quad (3.14)$$

such that $\lim_{n \rightarrow \infty} \lambda_n = \infty$. Moreover, we can build an orthonormal basis for $L^2(\mathcal{M})$ comprising eigenfunctions f_l of Δ_{aug} , where $\Delta_{\text{aug}} f_l = \lambda_l f_l$. Denote by $H^1(\mathcal{M})$ the Sobolev space of functions possessing a weak gradient in $L^2(\mathcal{M})$. Alternatively, $H^1(\mathcal{M})$ is a subspace of $L^2(\mathcal{M})$ consisting of functions with a finite *continuum Dirichlet energy*, which is expressed for $f \in L^2(\mathcal{M})$, as

$$E_{\mathcal{M}}(f) := \begin{cases} \int_{\mathcal{M}} |\nabla f(x)|^2 q^2(x) d\mathcal{V}(x) = 2\langle \Delta_{\text{aug}} f, f \rangle_{L^2(\mathcal{M})} & \text{if } f \in H^1(\mathcal{M}), \\ +\infty & \text{otherwise.} \end{cases} \quad (3.15)$$

Similarly as before, the eigenvalues in (3.14) of Δ_{aug} can be written variationally in terms of Dirichlet energy; that is

$$\lambda_l = \min_{S \in \mathbb{G}_{\mathcal{M}}^l} \max_{f \in S \setminus \{0\}} \frac{E_{\mathcal{M}}(f)}{\|f\|_{L^2(\mathcal{M})}^2}, \quad (3.16)$$

where $\mathbb{G}_{\mathcal{M}}^l$ denotes the Grassmannian manifold of all linear subspaces of $L^2(\mathcal{M})$ of dimension l .

Inspired by (Calder and García Trillos, 2022), our spectral convergence result hinges on a crucial interpolation between the graph-based and continuum energies. In our context, this interpolation is achieved through an optimal transport that, roughly speaking, connects the sampling probability measure μ on $\mathcal{M}$ with the discrete probability measure ϑ_n on $\tilde{\mathcal{X}}$, mediated by a suitable discrete probability measure v_n on $\mathcal{M}$, as will be specified later in Proposition 20. To facilitate the construction of such a map, and to assess the error in our subsequent analysis of discrete-to-continuum approximation, we introduce two additional key parameters: ι_1, ι_2 . The first measures the proximity of v_n to ϑ_n in an optimal transport sense, and the second reflects the uniform closeness of v_n to ρ . Furthermore, a stricter condition on ε must be imposed, requiring it to lie within a narrower range than specified in (3.7), that is

$$\exp\left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}}\right) \leq \varepsilon_w^m \varepsilon^3. \quad (3.17)$$

The following summarizes the technical smallness assumptions on all the involved parameters and hyperparameters, ensuring we operate in the appropriate regime for the validity of our theorems.

Assumption 10 (Parameter Assumption) *Let $\tau > 2$, $\eta > 0$, and let $\varepsilon > 0$ be sufficiently small as to satisfy (2.15), (3.2), (3.17). Further, let $\iota_1, \iota_2 > 0$ satisfy:*

1. $\max\{n^{-1/m}, C\varepsilon^{3/2}\} < \iota_2 \leq \frac{\varepsilon}{4}$, where $C > 0$ depends on $\mathcal{M}$;
2. $\iota_1 + \iota_2 \leq c$, where $c > 0$ depends on $\mathcal{M}$, ρ .

Observe that ι_1, ι_2 respectively correspond to $\tilde{\delta}, \theta$ in (Calder and García Trillos, 2022). The additional lower bound $C\varepsilon^{3/2} \leq \iota_2$ in Assumption 10, compared to (Calder and García Trillos, 2022, Assumption 2.3), reflects an extra step necessary in our proof of Theorem 12, specifically in Lemma 24.

We present the setting for our second main result.

Setting II Let $\eta, \tau > 0$ and $\varepsilon, \iota_1, \iota_2 > 0$. Let $\varepsilon_w, \varepsilon_p, \varepsilon_n$ be given in Table 1. Let $\mathcal{M} \subset \mathbb{R}^d$ be a closed manifold of dimension m , satisfying the fundamental geometry outlined in Section 2.2. Let $\mathcal{X} = \{x_1, \dots, x_N\}$ be the set of natural data, where $x_i \sim \mu$, whose density $\rho \in \mathcal{C}^2(\mathcal{M})$ satisfies (2.14). Let $\bar{\mathcal{X}} = \{\bar{x}_1, \dots, \bar{x}_n\}$ be the set of augmented data generated from $\mathcal{X}$ according to the law (2.1). Let Assumption 1 hold. Let $G_{\bar{\mathcal{X}}}$ be the augmentation graph on $\bar{\mathcal{X}}$ as described in Section 2.1. Let the augmentation graph Laplacian $\mathcal{L}_{\text{aug}}$ be given in (2.4), possessing the spectrum (3.11) with the corresponding eigenvectors $\hat{f}_l^{\text{aug}}$, $l = 1, \dots, n$. Let Assumptions 5, 10 hold. Finally, let Δ_{aug} be the Laplace-Beltrami operator defined in (3.8), possessing the spectrum (3.14) with the corresponding eigenfunctions $\hat{f}_l$, $l \in \mathbb{N}$.

3.3 Setting for Realizability in Contrastive Learning

We begin by providing a matrix form for $\mathcal{L}_{\text{aug}}$ (2.4). Let $\mathbf{A}_{\text{aug}} \in \mathbb{R}^{n \times n}$ denote the symmetric adjacency matrix of the graph $G_{\bar{\mathcal{X}}}$ having the entries

$$(\mathbf{A}_{\text{aug}})_{ij} := \frac{1}{\varepsilon_w^m \varepsilon^{m+2}} \omega(\bar{x}_i, \bar{x}_j) \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} \quad \text{for } i, j = 1, \dots, n. \quad (3.18)$$

Let $\mathbf{D}_{\text{aug}} \in \mathbb{R}^{n \times n}$ be the corresponding diagonal degree matrix. Then $(\mathbf{D}_{\text{aug}})_{ij} = 0$ if $i \neq j$, and

$$(\mathbf{D}_{\text{aug}})_{ii} := \sum_{j=1}^n (\mathbf{A}_{\text{aug}})_{ij} = \frac{1}{\varepsilon_w^m \varepsilon^{m+2}} \sum_{j=1}^n \omega(\bar{x}_i, \bar{x}_j) \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}}. \quad (3.19)$$

For $f : \bar{\mathcal{X}} \rightarrow \mathbb{R}$, we identify f with the column vector $f = (f(\bar{x}_1), \dots, f(\bar{x}_n))^{\top} \in \mathbb{R}^n$. Then under such identification, we verify that

$$\mathcal{L}_{\text{aug}} f = \frac{1}{n} [\mathbf{D}_{\text{aug}} - \mathbf{A}_{\text{aug}}] f.$$

Indeed, adapted from definition (2.4), at $\bar{x}_i$, we have

$$\begin{aligned} \mathcal{L}_{\text{aug}} f(\bar{x}_i) &= \frac{1}{n \varepsilon_w^m \varepsilon^{m+2}} \sum_{\bar{x}_j \in \bar{\mathcal{X}}: 0 < \|\bar{x}_j - \bar{x}_i\| \leq \varepsilon} \omega(\bar{x}_i, \bar{x}_j) (f(\bar{x}_i) - f(\bar{x}_j)) \\ &= \frac{1}{n} \left(f(\bar{x}_i) (\mathbf{D}_{\text{aug}})_{ii} - \sum_{j=1}^n (\mathbf{A}_{\text{aug}})_{ij} f(\bar{x}_j) \right) \\ &= \frac{1}{n} \left(\sum_{j=1}^n \delta_{ij} (\mathbf{D}_{\text{aug}})_{ij} f(\bar{x}_j) - \sum_{j=1}^n (\mathbf{A}_{\text{aug}})_{ij} f(\bar{x}_j) \right) \\ &= \frac{1}{n} \sum_{j=1}^n ((\mathbf{D}_{\text{aug}})_{ij} - (\mathbf{A}_{\text{aug}})_{ij}) f(\bar{x}_j) = \frac{1}{n} [\mathbf{D}_{\text{aug}} - \mathbf{A}_{\text{aug}}] f, \end{aligned}$$

where the second equality follows from definitions (3.18), (3.19), and the fourth equality uses the fact that $\mathbf{D}_{\text{aug}}$ is diagonal. Therefore, we will refer to the graph Laplacian operator

$\mathcal{L}_{\text{aug}}$ as the matrix $\frac{1}{n}[\mathbf{D}_{\text{aug}} - \mathbf{A}_{\text{aug}}] \in \mathbb{R}^{n \times n}$ whenever appropriate⁵. Recall from Section 2.2 that $\|\cdot\|$ denotes the Euclidean vector norm or the Frobenius matrix norm. Let $k \in \mathbb{N}$, which acts as a parameter determining the spectral embedding dimension. Let $a \geq \max\{(\mathbf{D}_{\text{aug}})_{ii}\}$. It follows that $a\mathbf{I}_{n \times n} - \mathcal{L}_{\text{aug}}$, where $\mathbf{I}_{n \times n}$ is the $n \times n$ identity matrix, is positive semidefinite. Then by Eckart-Young-Minsky theorem (Eckart and Young, 1936), any minimizer $Y^* \in \mathbb{R}^{n \times k}$, of the optimization objective⁶

$$\min_{Y \in \mathbb{R}^{n \times k}} \|YY^\top - a\mathbf{I}_{n \times n} + \mathcal{L}_{\text{aug}}\|^2, \quad (3.20)$$

contains, as its columns, the scaled versions of the last k normalized (in the Euclidean norm) eigenvectors of $a\mathbf{I}_{n \times n} - \mathcal{L}_{\text{aug}}$ (corresponding to the largest k eigenvalues), up to multiplication on the right by an $k \times k$ orthogonal matrix. These eigenvectors are in turn, the first k normalized (in the Euclidean norm) eigenvectors of $\mathcal{L}_{\text{aug}}$ (corresponding to the smallest k eigenvalues). In other words, solving (3.20) is equivalent to performing spectral embedding of $G_{\mathcal{X}}$ using $\mathcal{L}_{\text{aug}}$. We aim to learn a minimizer of (3.20) over a hypothesis set of neural networks (NNs) with parameterizable architectures. Specifically, we solve the following optimization problem:

$$\min_{\theta \in \Theta} \|Y_\theta Y_\theta^\top - a\mathbf{I}_{n \times n} + \mathcal{L}_{\text{aug}}\|^2, \quad (3.21)$$

where θ denotes the parameter vector of a NN $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^k$, and Θ denotes the space of all admissible parameter vectors. Moreover, $Y_\theta \in \mathbb{R}^{n \times k}$ contains rows that are the outputs $f_\theta(\bar{x}_i)$, representing the embeddings of $\bar{x}_i \in \mathbb{R}^d$ into $\mathbb{R}^k$. Thus, the mapping realized by the NN⁷ f_θ can be interpreted as a *feature* or *representation map* for the input data cloud $\bar{\mathcal{X}}$. We specify the NNs f_θ as multi-layer ReLU neural networks (ReLU NNs), with the architectures given as follows. Let $\mathbf{M}, \mathbf{L}, \mathbf{p} \in \mathbb{N}$ and $\mathbf{m} > 0$. The NNs f_θ are functions of the form:

$$\mathbb{R}^d \ni x \mapsto W_{\mathbf{L}}(\text{ReLU}(W_{\mathbf{L}-1} \cdots \text{ReLU}(W_1 x + b_1) \cdots + b_{\mathbf{L}-1})) + b_{\mathbf{L}} \in \mathbb{R}^k. \quad (3.22)$$

Here, $W_i \in \mathbb{R}^{\mathbf{p} \times \mathbf{p}}, b_i \in \mathbb{R}^{\mathbf{p}}$ for $i = 2, \dots, \mathbf{L} - 1$, $W_1 \in \mathbb{R}^{\mathbf{p} \times d}, b_1 \in \mathbb{R}^{\mathbf{p}}$, and $W_{\mathbf{L}} \in \mathbb{R}^{k \times \mathbf{p}}, b_{\mathbf{L}} \in \mathbb{R}^k$. It follows that each network architecture described in (3.22) is characterized by a parameter vector

$$\theta := (W_1, b_1, \dots, W_{\mathbf{L}}, b_{\mathbf{L}}). \quad (3.23)$$

It can be seen from (3.22) that $\mathbf{L}$ corresponds to the *network depth* $\mathbf{L}(f_\theta)$ of f_θ , and $\mathbf{p}$ corresponds to the *network width* $\mathbf{p}(f_\theta)$. We say a parameter θ is admissible if, additionally, the following two conditions on *network complexity* hold:

$$\|W_i\|_\infty \leq \mathbf{m} \quad \text{and} \quad \|b_i\|_\infty \leq \mathbf{m} \quad \text{for} \quad i = 1, \dots, \mathbf{L}, \quad (3.24)$$

-
5. The choice to keep the factor $1/n$ explicit is intended to emphasize the underlying empirical Monte-Carlo averaging process, which is natural from the manifold-learning perspective. If instead one rescales $\mathbf{A} \leftarrow \mathbf{A}/n$ and $\mathbf{D} \leftarrow \mathbf{D}/n$ then one recovers the standard graph-theoretic definition of the unnormalized graph Laplacian $\mathcal{L}_{\text{aug}} = \mathbf{D} - \mathbf{A}$.
 6. Note that (3.20) can be rewritten as $\min_{Y \in \mathbb{R}^{n \times k}} \ell(Y)$ in terms of the loss function $\ell(Y) := \|YY^\top - a\mathbf{I}_{n \times n} + \mathcal{L}_{\text{aug}}\|^2$. Then by the reasoning given in (HaoChen et al., 2021), this optimization problem, involving matrix factorization, corresponds to an optimization problem based on their population spectral contrastive loss.
 7. In this paper, we will not differentiate between a network architecture and its realization.

and

$$\sum_{i=1}^L \|W_i\|_0 + \|b_i\|_0 \leq M. \quad (3.25)$$

Condition (3.24) imposes a bound on the magnitude $\mathbf{m}(f_\theta)$ of the network weights, whereas condition (3.25) restricts the number $M(f_\theta)$ of nonzero weights. Equivalently, the space Θ is defined as

$$\Theta := \{\theta \text{ given in (3.23)} : \theta \text{ satisfies (3.24), (3.25)}\}. \quad (3.26)$$

We conclude the subsection by introducing the setting for our third main result.

Setting III *Let Setting II hold. For $k \in \mathbb{N}$, consider the optimization problem (3.21), where*

$$Y_\theta = (f_\theta(\bar{x}_1), \dots, f_\theta(\bar{x}_n))^\top \in \mathbb{R}^{n \times k}, \quad (3.27)$$

and $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^k$ is a ReLU NN of the form (3.22) for some $L, p \in \mathbb{N}$. Moreover, θ given in (3.23) is a network parameter vector, and the space Θ of admissible parameters is defined in (3.26) for some $M \in \mathbb{N}$, $\mathbf{m} > 0$.

4. Main Results

Our main results are framed as probabilistic statements, expressed in terms of all relevant parameters and hyperparameters.

A Pointwise Consistency Result. Let

$$\alpha := \int_{\mathbb{R}^m} \exp(-\|y\|^2) dy \quad (4.1)$$

and

$$\beta := \int_{B^m(0,1)} |y^1|^2 dy \quad (4.2)$$

where,⁸ recall that y^1 denotes the first coordinate of $y \in \mathbb{R}^m$. Our first result demonstrates that, with high probability, $\mathcal{L}_{\text{aug}}$ behaves pointwise in close approximation to an $\alpha\beta\Delta_{\text{aug}}$. A proof is given in Section 6.3.

Theorem 11 *Assume Setting I. Let $f \in \mathcal{C}^3(\mathcal{M})$. Then for $\varepsilon \leq \delta \leq \varepsilon^{-1}$,*

$$\mathbb{P}\left(\max_{\bar{x} \in \mathcal{X}} \left| \mathcal{L}_{\text{aug}} f(\bar{x}) - \alpha\beta\Delta_{\text{aug}} f(Q_{\bar{x}}) \right| > C_\eta \delta \|f\|_{\mathcal{C}^3(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}\right) \leq 8n \exp(-c_\eta \delta^2 n \varepsilon^{m+2}),$$

for some $C_\eta, c_\eta > 0$ depending on η and $\mathcal{M}$, ρ at most.

We briefly remark that included in Setting I is Assumption 5. By Lemma 2, the event specified in the assumption holds with probability at least $1 - 2dn \exp(-\eta^{-2/d}/(2d))$. This probability can also be incorporated into the statement of Theorem 11.

8. A simple computation shows that $\beta = \sigma_m/(m+2)$, with σ_m being the volume of the m -dimensional Euclidean unit ball.

A Spectral Consistency Result. We introduce the concept of the eigengap. Let λ_l be the l th eigenvalue of Δ_{aug} . Let $\underline{\lambda}_{i_l}$ denote the i_l th *distinct* eigenvalue of Δ_{aug} , listed in an increasing order, that corresponds to λ_l , i.e. $\lambda_l = \underline{\lambda}_{i_l}$. Then the *eigengap* associated with λ_l is defined to be

$$\gamma_{\lambda_l} := \frac{1}{2} \min\{|\underline{\lambda}_{i_l} - \underline{\lambda}_{i_l-1}|, |\underline{\lambda}_{i_l} - \underline{\lambda}_{i_l+1}|\}. \quad (4.3)$$

Our second result establishes the spectral convergence of $\mathcal{L}_{\text{aug}}$ to $\alpha\beta\Delta_{\text{aug}}$ with high probability. A proof is given in Section 7.2.

Theorem 12 *Assume Setting II. Let $l \in \mathbb{N}$. Let $\lambda_l, \hat{\lambda}_l^{\text{aug}}$ denote the l th eigenvalue of $\Delta_{\text{aug}}, \mathcal{L}_{\text{aug}}$, respectively. Suppose it holds that*

$$\sqrt{\lambda_l}\varepsilon + C(\iota_1 + \iota_2) \leq \frac{1}{2l}, \quad (4.4)$$

and that

$$C(\varepsilon(\sqrt{\lambda_l} + 1) + \varepsilon + \iota_1)\lambda_l \leq \gamma_{\lambda_l}, \quad (4.5)$$

where γ_{λ_l} is given in (4.3).

Then with probability at least $1 - n \exp(-c\iota_1^2\iota_2^m) - 10n \exp(-c_\eta n \varepsilon^{m+4})$, for the l th normalized eigenvector $\hat{f}_l^{\text{aug}}$ (with respect to $L_{\vartheta_n}^2(\bar{\mathcal{X}})$) of $\mathcal{L}_{\text{aug}}$ associated with $\hat{\lambda}_l^{\text{aug}}$, there exists a normalized eigenfunction f_l (with respect to $L^2(\mathcal{M})$) of Δ_{aug} associated with λ_l , such that

$$\|\hat{f}_l^{\text{aug}}(\cdot) - f_l(Q_\cdot)\|_{L_{\vartheta_n}^2(\bar{\mathcal{X}})} \leq C_{\eta,l}\varepsilon, \quad (4.6)$$

where $Q_\cdot$ denotes the orthogonal projection of a point in $\bar{\mathcal{X}}$ onto $\mathcal{M}$. Here, all the constants $C, c > 0$ depend on $\mathcal{M}, \rho$ at most; the constant $c_\eta > 0$ additionally depends on η , and the constant $C_{\eta,l} > 0$ further depends on l .

The additional dependence on l of the constant $C_{\eta,l}$ can be explained via Theorem 11 as follows. Since $f_l \in \mathcal{C}^3(\mathcal{M})$, we can extend f_l to a function $\bar{f}_l$ in the neighborhood $\mathcal{N}_{\varepsilon_n}(\mathcal{M})$, ensuring that $\bar{f}_l(x) = \bar{f}_l(x+h)$ whenever $x+h \in \mathcal{N}_{\varepsilon_n}(\mathcal{M}) \cap \tilde{N}_x\mathcal{M}$ (as per the notation in Section 3.1). Note that this extension is well-defined, because a nonempty intersection between $\mathcal{N}_{\varepsilon_n}(\mathcal{M}) \cap \tilde{N}_x\mathcal{M}$ and $\mathcal{N}_{\varepsilon_n}(\mathcal{M}) \cap \tilde{N}_{x'}\mathcal{M}$ for $x \neq x'$ on $\mathcal{M}$ would imply $\varepsilon_n > R$, contradicting (2.15). Thus, the dependence on l manifests as a factor of $\|\bar{f}_l\|_{\mathcal{C}^3(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}$, which equals $\|f_l\|_{\mathcal{C}^3(\mathcal{M})}$.

Note that, under Assumption 10, if ε scales like $C\left(\frac{\log n}{n}\right)^{\frac{1}{m+4}} \leq \varepsilon \ll 1$, and we set $\iota_1 = C\varepsilon^2$ and $\iota_2 = c\varepsilon$ for some $C > 0$ and $c \in (0, 1/4]$, then with high probability the spectral convergence rate obtained from (4.6) scales linearly in ε . Particularly, if $\varepsilon = C\left(\frac{\log n}{n}\right)^{\frac{1}{m+4}}$, the convergence rate obtained is of $\mathcal{O}_\eta\left(\left(\frac{\log n}{n}\right)^{\frac{1}{m+4}}\right)$. For sufficiently large n , such parameter selection also ensures that (4.4), (4.5) are automatically satisfied.

A Realizability Result. Lastly, we provide an approximation guarantee for neural networks in spectral contrastive learning tasks. This result relaxes the realizability assumption (HaoChen et al., 2021, Assumption 3.7) in the spectral contrastive learning framework by asserting the existence of a neural network with bounded complexity capable of learning a minimizer of (3.20). A proof is given in Section 8.

Theorem 13 *Assume Setting III. Suppose (4.4), (4.5) hold for all $l = 1, \dots, k$. Then with probability at least $1 - kn \exp(-cn\iota_1^2\iota_2^m) - k10n \exp(-c_\eta n \varepsilon^{m+4})$, the following holds. For every $\tilde{\delta} \in (0, 1)$ and every minimizer $Y^* \in \mathbb{R}^{n \times k}$ of the problem (3.20), there exists a ReLU NN f_θ such that for Y_θ in (3.27),*

$$\|Y^* - Y_\theta\|_\infty \leq C_{\eta,k}(\varepsilon + \tilde{\delta}) \quad \text{and} \quad \|Y^* - Y_\theta\| \leq C_{\eta,k}(\varepsilon + \tilde{\delta}).$$

Moreover,

$$\mathcal{L}(f_\theta) \leq C \left(\log \left(\frac{1}{\tilde{\delta}} \right) + \log d \right) \quad \text{and} \quad \mathfrak{p}(f_\theta) \leq Ck(\tilde{\delta}^{-m} + d),$$

and $\theta \in \Theta$, for $M \in \mathbb{N}$, $m > 0$, such that

$$m \leq C \quad \text{and} \quad M \leq k \left((\tilde{\delta}^{-m} + d) \log \left(\frac{1}{\tilde{\delta}} \right) + d \log d \right). \quad (4.7)$$

Here, all the constants $C, c > 0$ depend only on $\mathcal{M}$; the constant $c_\eta > 0$ additionally depends on η , and the constants $C_{\eta,k} > 0$ further depend on k .

The dependence of the constants $C_{\eta,k}$ on k arises partly from a similar dependence in (4.6). Additionally, these constants include at most an extra factor of $\sqrt{k} \max_{l=1,\dots,k} \sqrt{\hat{\lambda}_l^{\text{aug}}}$.

Similarly as above, if ε scales as $C \left(\frac{\log n}{n} \right)^{\frac{1}{m+4}} \leq \varepsilon \ll 1$, and we set $\iota_1 = C_1 \varepsilon^2$, $\iota_2 = c_2 \varepsilon$, and $\tilde{\delta} = C_3 \varepsilon$ for some constants $C_1, C_3 > 0$ and $c_2 \in (0, 1/4]$, then with high probability, the required network depth scales logarithmically with respect to the inverse of ε , while the network width and the number of neurons scale sublinearly. Specifically, when $\varepsilon = C \left(\frac{\log n}{n} \right)^{\frac{1}{m+4}}$, the network depth is of $\mathcal{O}(\log(\frac{n}{\log n}))$, and the network width and the number of neurons are of $\mathcal{O}((\frac{n}{\log n})^{\frac{m}{m+4}})$, up to an extra log factor. Note that these rates depend solely on the embedded dimension m and not on the ambient dimension d , effectively mitigating the curse of dimensionality.

5. Numerical Experiments

In this section, we implement the spectral contrastive learning algorithm⁹ using parameter choices informed by our theoretical analysis¹⁰. We follow the same setup as in (HaoChen et al., 2021) and test on benchmark vision data sets. We minimize the empirical spectral contrastive loss with an encoder network f and generate augmentation in each iteration.

Encoder. The encoder function f contains a backbone network, a projection MLP, and a projection function. The backbone network is a standard ResNet architecture. The projection MLP is a fully connected network with the backbone network applied to each layer, and ReLU activation applied to each except for the last layer. The projection function takes a vector and projects it onto a sphere with radius $\sqrt{r}$, where $r > 0$ is a tunable hyperparameter. Additionally, we specify the relationship between the graph connectivity and data generation parameters according to our theory (see (2.4) and Table 1): we use

9. The original implementation is available at https://github.com/jhaochenz96/spectral_contrastive_learning.

10. Our implementation is in https://github.com/chl781/Augmented_spectral_contrastive_learning.

k NN to determine the connectivity parameter (for algorithm stability) and use quantile heuristic to compute ε which is used to generate Gaussian noise for the augmented data set.

Linear Evaluation Protocol. Given the pre-trained encoder network, we follow the standard linear evaluation protocol (Chen and He, 2021) and train a supervised linear classifier on frozen representations.

Results. We present the accuracy results on CIFAR-10/100 (Krizhevsky, 2009) in Table 2. Our empirical findings show that our spectral contrastive learning approach either outperforms or performs comparably to two widely used baseline algorithms, SimCLR (Chen et al., 2020b) and SimSiam (Chen et al., 2020a). Notably, our algorithm builds upon (HaoChen et al., 2021) by incorporating an explicit, theoretically motivated choice of the noise parameter. It therefore avoids exhaustive hyperparameter tuning, while exhibiting similar empirical behavior to that of (HaoChen et al., 2021).

Data sets	CIFAR-10			CIFAR-100		
Epochs	200	400	800	200	400	800
SimCLR	83.73	87.72	90.60	54.74	61.05	63.88
SimSiam	87.54	90.31	91.40	61.56	64.96	65.87
(HaoChen et al., 2021)	88.66	90.17	92.07	62.45	65.82	66.18
Ours	87.80	89.79	91.70	58.64	63.81	66.36

Table 2: Top-1 accuracy under linear evaluation protocol. SimCLR and SimSiam results are reported in (HaoChen et al., 2021).

6. Pointwise Consistency of Augmentation Graph Laplacian

6.1 From Augmentation Graph to Projected Graph

We maintain Assumption 5. For $\bar{x} \neq \bar{x}' \in \bar{\mathcal{X}}$ and each $x \in \mathcal{M}$, we define the random variable

$$X = X(x; \bar{x}, \bar{x}') := \mathbb{1}_{\{\|x - Q_{\bar{x}}\| < \varepsilon_n^{2/3}\} \wedge \{\|x - Q_{\bar{x}'}\| < \varepsilon_n^{2/3}\}}, \quad (6.1)$$

where $Q_{\bar{x}}, Q_{\bar{x}'}$ are the respective projections onto $\mathcal{M}$ of $\bar{x}, \bar{x}'$. For each such pair, we equip with a connection strength

$$\varsigma(Q_{\bar{x}}, Q_{\bar{x}'}):= \mathbb{E}_{x \sim \mu} \left[X \exp \left(\frac{-\|Q_{\bar{x}} - x\|^2 - \|Q_{\bar{x}'} - x\|^2}{2\varepsilon_w^2} \right) \right]. \quad (6.2)$$

Let $\mathcal{Q}_{\bar{\mathcal{X}}}$ denote the set of all $Q_{\bar{x}}$ from $\bar{x} \in \bar{\mathcal{X}}$. Let $G_{\mathcal{Q}_{\bar{\mathcal{X}}}}$ be a graph on the vertex set $\mathcal{Q}_{\bar{\mathcal{X}}}$, where each edge $(Q_{\bar{x}}, Q_{\bar{x}'})$ corresponds to an edge $(\bar{x}, \bar{x}')$ in $\bar{\mathcal{X}}$, i.e. $0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon$, and vice versa. Construct a projected graph Laplacian operator¹¹ $\mathcal{L}_{\text{proj}}$ such that for $f: \mathcal{Q}_{\bar{\mathcal{X}}} \rightarrow \mathbb{R}$,

$$\mathcal{L}_{\text{proj}} f(Q_{\bar{x}}) := \frac{1}{n\varepsilon_w^m \varepsilon^{m+2}} \sum_{Q_{\bar{x}'}: 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} \varsigma(Q_{\bar{x}}, Q_{\bar{x}'})(f(Q_{\bar{x}}) - f(Q_{\bar{x}'})). \quad (6.3)$$

11. The subscript “proj” stands for “projection.”

The augmentation graph Laplacian $\mathcal{L}_{\text{aug}}$ serves as an approximation to the graph Laplacian $\mathcal{L}_{\text{proj}}$. This is formally stated in the following proposition, whose proof can be found at the end of this subsection.

Proposition 14 *Let $\tau > 2$, $\varepsilon > 0$, and $\varepsilon_{\text{w}} = \varepsilon^\tau$, $\varepsilon_{\text{n}} = \varepsilon^{\tau+1}$. Let Assumption 5 hold. Let $\mathcal{L}_{\text{aug}}$ and $\mathcal{L}_{\text{proj}}$ be respectively defined as in (2.4) and in (6.3). Let $f \in \mathcal{C}^1(\mathcal{N}_{\varepsilon_{\text{n}}}(\mathcal{M}))$. Then for ε sufficiently small satisfying (3.7), and $\varepsilon^2 \leq \delta \leq 1$, we have*

$$\mathbb{P}\left(\max_{\bar{x} \in \bar{\mathcal{X}}} \left| \mathcal{L}_{\text{aug}} f(\bar{x}) - \mathcal{L}_{\text{proj}} f(Q_{\bar{x}}) \right| > C_\eta \varepsilon \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_{\text{n}}}(\mathcal{M}))} (1 + \delta)\right) \leq 6n \exp(-c\delta^2 n \varepsilon^m),$$

where $c > 0$ depends on $\mathcal{M}$, ρ at most, and $C_\eta > 0$ further depends on η .

In preparation, we introduce a key prerequisite result. Using (6.2), we define an intermediate connection between $\bar{x}, \bar{x}'$ that simultaneously approximates $\omega(\bar{x}, \bar{x}')$ and $\varsigma(Q_{\bar{x}}, Q_{\bar{x}'})$ up to a factor related to the distance of $\bar{\mathcal{X}}$ to the manifold $\mathcal{M}$; that is

$$\xi(\bar{x}, \bar{x}') := \exp\left(\frac{-\|\bar{x} - Q_{\bar{x}}\|^2 - \|\bar{x}' - Q_{\bar{x}'}\|^2}{2\varepsilon_{\text{w}}^2}\right) \varsigma(Q_{\bar{x}}, Q_{\bar{x}'}). \quad (6.4)$$

The relationship between ω , ξ is clarified by the following proposition, whose technical proof can be found in Appendix C.

Proposition 15 *Let $\tau > 0$, $\varepsilon > 0$. Let ω , ξ be as in (2.3), (6.4), respectively. Then under Assumption 5, the following holds,*

$$\exp(-C\varepsilon^2)\xi(\bar{x}, \bar{x}') \leq \omega(\bar{x}, \bar{x}') \leq \exp(C'\varepsilon^2)\xi(\bar{x}, \bar{x}') + \exp\left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}}\right), \quad (6.5)$$

for some universal $c > 0$, and $C, C' > 0$ depending on $\mathcal{M}$.

Proposition 15 provides an approximate Pythagorean decomposition. Indeed, this is encoded in the definition (6.4) of $\xi(\bar{x}, \bar{x}')$, which separates the contributions $\|x - Q_{\bar{x}}\|^2$ and $\|x - Q_{\bar{x}'}\|^2$ from the contributions $\|\bar{x} - Q_{\bar{x}}\|^2$ and $\|\bar{x}' - Q_{\bar{x}'}\|^2$ to the ambient Euclidean distances $\|x - \bar{x}\|^2$ and $\|x - \bar{x}'\|^2$, respectively, appearing in $\omega(\bar{x}, \bar{x}')$. Since the manifold is curved, such an approximate decomposition holds up to small curvature-induced additive error terms given in (6.5). These terms are controlled via the choice of range $\varepsilon_{\text{n}}^{2/3}$ in definition (6.1) of X . In turn, they ensure that the resulting error contributions remain sufficiently lower order in the subsequent estimates. Next, we present the final lemma before moving on to the proof of Proposition 14. This lemma explains the additional scaling factor ε_{w}^m in definition (2.4).

Lemma 16 *Let $\tau > 0$, $\varepsilon > 0$, and $\varepsilon_{\text{w}} = \varepsilon^\tau$. Then*

$$\max_{\bar{x}, \bar{x}' \in \bar{\mathcal{X}}} \varsigma(Q_{\bar{x}}, Q_{\bar{x}'}) \leq C\varepsilon_{\text{w}}^m \quad \text{and} \quad \max_{\bar{x}, \bar{x}' \in \bar{\mathcal{X}}} \xi(\bar{x}, \bar{x}') \leq C\varepsilon_{\text{w}}^m,$$

for some $C > 0$ depending on $\mathcal{M}$, ρ .

Proof [Proof of Lemma 16] From definition (6.4), it suffices to prove $\|\varsigma\|_\infty \leq C\varepsilon_w^m$. By utilizing (2.7) and definition (6.2), we get

$$\begin{aligned} \varsigma(Q_{\bar{x}}, Q_{\bar{x}'}) &\leq \int_{\{x: \|x - Q_{\bar{x}}\| < \varepsilon_n^{2/3}\} \cap \mathcal{M}} \exp\left(-\frac{\|Q_{\bar{x}} - x\|^2}{2\varepsilon_w^2}\right) \rho(x) d\mathcal{V}(x) \\ &\leq \int_{\{x: d(Q_{\bar{x}}, x) < c\varepsilon_n^{2/3}\} \cap \mathcal{M}} \exp\left(-\frac{cd(Q_{\bar{x}}, x)^2}{\varepsilon_w^2}\right) \rho(x) d\mathcal{V}(x), \end{aligned} \quad (6.6)$$

for some $c > 0$ depending on $\mathcal{M}$. Since ε is sufficiently small (2.15), the map $\text{Exp}_{Q_{\bar{x}}} : B^m(0, c\varepsilon_n^{2/3}) \rightarrow \mathcal{B}(Q_{\bar{x}}, c\varepsilon_n^{2/3})$ is a diffeomorphism. Hence, recalling the change of variables and notation system given in Section 2.2, for $x \in \mathcal{B}(Q_{\bar{x}}, c\varepsilon_n^{2/3})$, there exists a unique $y \in B^m(0, c\varepsilon_n^{2/3})$ such that $\text{Exp}_{Q_{\bar{x}}}(y) = x$. Writing $\tilde{\rho}(y) = \rho(\text{Exp}_{Q_{\bar{x}}}(y))$, we obtain from (2.14), (2.10), (6.6)

$$\begin{aligned} \varsigma(Q_{\bar{x}}, Q_{\bar{x}'}) &\leq \int_{B^m(0, c\varepsilon_n^{2/3}) \subset T_{Q_{\bar{x}}}\mathcal{M}} \exp\left(-\frac{c\|y\|^2}{\varepsilon_w^2}\right) \tilde{\rho}(y) J_{Q_{\bar{x}}}(y) dy \\ &\leq C\rho_{\max}\varepsilon_w^m \int_{B^m(0, c\varepsilon_n^{-1/3}\varepsilon^{2/3})} \exp(-c\|y\|^2) dy \\ &\leq C\varepsilon_w^m, \end{aligned}$$

and the desired conclusion follows. $\blacksquare$

Proof [Proof of Proposition 14] We define two auxiliary operators. For $f : \bar{\mathcal{X}} \rightarrow \mathbb{R}$, let

$$\mathcal{L}_{\bar{\mathcal{X}}}f(\bar{x}) := \frac{1}{n\varepsilon_w^m\varepsilon^{m+2}} \sum_{\bar{x}': 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} \xi(\bar{x}, \bar{x}')(f(\bar{x}) - f(\bar{x}')),$$

and for $f : \mathcal{Q}_{\bar{\mathcal{X}}} \rightarrow \mathbb{R}$, let

$$\mathcal{L}_{\mathcal{Q}_{\bar{\mathcal{X}}}}f(Q_{\bar{x}}) := \frac{1}{n\varepsilon_w^m\varepsilon^{m+2}} \sum_{\bar{x}': 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} \xi(\bar{x}, \bar{x}')(f(Q_{\bar{x}}) - f(Q_{\bar{x}'})).$$

Recall that $f \in \mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))$. We make the following claims:

$$\left| \mathcal{L}_{\text{aug}}f(\bar{x}) - \mathcal{L}_{\bar{\mathcal{X}}}f(\bar{x}) \right| \leq \frac{C\varepsilon\|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1, \quad (6.7)$$

$$\left| \mathcal{L}_{\bar{\mathcal{X}}}f(\bar{x}) - \mathcal{L}_{\mathcal{Q}_{\bar{\mathcal{X}}}}f(Q_{\bar{x}}) \right| \leq \frac{C\varepsilon^{\tau-1}\|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1, \quad (6.8)$$

$$\left| \mathcal{L}_{\mathcal{Q}_{\bar{\mathcal{X}}}}f(Q_{\bar{x}}) - \mathcal{L}_{\text{proj}}f(Q_{\bar{x}}) \right| \leq \frac{C\varepsilon\|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1. \quad (6.9)$$

We will validate (6.7), (6.8), (6.9) sequentially. First, note that if $\|\bar{x} - \bar{x}'\| \leq \varepsilon$, then

$$\|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq \|\bar{x} - \bar{x}'\| + \|Q_{\bar{x}} - \bar{x}\| + \|\bar{x}' - Q_{\bar{x}'}\| \leq c\varepsilon + 2\varepsilon_n \leq 2\varepsilon. \quad (6.10)$$

Proof of (6.7). Recall from Proposition 15 that

$$\left[\exp(-C\varepsilon^2) - 1 \right] \xi(\bar{x}, \bar{x}') \leq \omega(\bar{x}, \bar{x}') - \xi(\bar{x}, \bar{x}') \leq \left[\exp(C\varepsilon^2) - 1 \right] \xi(\bar{x}, \bar{x}') + \exp\left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}}\right), \quad (6.11)$$

which, together with Lemma 16 and that $\tau > 2$, delivers

$$|\omega(\bar{x}, \bar{x}') - \xi(\bar{x}, \bar{x}')| \leq C\varepsilon_w^m \varepsilon^2, \quad (6.12)$$

for sufficiently small ε satisfying (3.7). Using (6.10), (6.12), we conclude for $\bar{x} \in \bar{\mathcal{X}}$,

$$\begin{aligned} \left| \mathcal{L}_{\text{aug}} f(\bar{x}) - \mathcal{L}_{\bar{\mathcal{X}}} f(\bar{x}) \right| &\leq \frac{1}{n\varepsilon_w^m \varepsilon^{m+2}} \sum_{\bar{x}': 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} |\omega(\bar{x}, \bar{x}') - \xi(\bar{x}, \bar{x}')| |f(\bar{x}) - f(\bar{x}')| \\ &\leq \frac{C\varepsilon \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^m} \sum_{\bar{x}': 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} 1 \\ &\leq \frac{C\varepsilon \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1, \end{aligned}$$

which completes the proof for (6.7).

Proof of (6.8). We note that Assumption 5 implies

$$|f(\bar{x}) - f(Q_{\bar{x}})| \leq \varepsilon_n \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))} \quad \text{and} \quad |f(\bar{x}') - f(Q_{\bar{x}'})| \leq \varepsilon_n \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}. \quad (6.13)$$

Hence, from (6.10), (6.13) and Lemma 16 again, we can finish (6.8):

$$\begin{aligned} \left| \mathcal{L}_{\bar{\mathcal{X}}} f(\bar{x}) - \mathcal{L}_{Q_{\bar{\mathcal{X}}}} f(\bar{x}) \right| &\leq \frac{1}{n\varepsilon_w^m \varepsilon^{m+2}} \sum_{\bar{x}': 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} \xi(\bar{x}, \bar{x}') \left(|f(\bar{x}) - f(Q_{\bar{x}})| + |f(\bar{x}') - f(Q_{\bar{x}'})| \right) \\ &\leq \frac{2\varepsilon_n \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^{m+2}} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1 \\ &\leq \frac{2\varepsilon^{\tau-1} \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1. \end{aligned}$$

This completes the proof for (6.8).

Proof of (6.9). Assumption 5 gives

$$\exp\left(\frac{-\|\bar{x} - Q_{\bar{x}}\|^2 - \|\bar{x}' - Q_{\bar{x}'}\|^2}{2\varepsilon_w^2}\right) \geq \exp(-C\varepsilon^2),$$

whence, from definition (6.4), it follows that

$$|\xi(\bar{x}, \bar{x}') - \varsigma(Q_{\bar{x}}, Q_{\bar{x}'})| \leq \varsigma(Q_{\bar{x}}, Q_{\bar{x}'})(1 - \exp(-C\varepsilon^2)) \leq C\varepsilon_w^m \varepsilon^2. \quad (6.14)$$

Subsequently, from (6.14), and (6.10) once more,

$$\begin{aligned} \left| \mathcal{L}_{Q_{\bar{\mathcal{X}}}} f(\bar{x}) - \mathcal{L}_{\text{proj}} f(Q_{\bar{x}}) \right| &\leq \frac{1}{n\varepsilon_w^m \varepsilon^{m+2}} \sum_{\bar{x}': 0 < \|\bar{x} - \bar{x}'\| \leq \varepsilon} |\xi(\bar{x}, \bar{x}') - \varsigma(Q_{\bar{x}}, Q_{\bar{x}'})| |f(Q_{\bar{x}}) - f(Q_{\bar{x}'})| \\ &\leq \frac{C\varepsilon \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1. \end{aligned}$$

This completes the proof for (6.9).

We now estimate the right-hand sides of (6.7), (6.8), (6.9). Recall from Section 3.1 that the projection points $Q_{\bar{x}} \sim \nu$ on $\mathcal{M}$, with the probability density q . By Proposition 7, $q \in \mathcal{C}^2(\mathcal{M})$ and $q(x) \leq C_\eta$ for all $x \in \mathcal{M}$. By rescaling, we can assume 2ε satisfies (2.15) in place of ε . Let $\varepsilon^2 \leq \delta \leq 1$. We note that the points $Q_{\bar{x}}$ are not independent, as they are projections of the augmented data points. However, under Lemma 6 and Assumption 1, we can assert that with probability 1, each $Q_{\bar{x}}$ has at most C_0 other $Q_{\bar{x}'}$ dependent on $Q_{\bar{x}}$. Therefore, an application of Lemma 28 gives

$$\frac{1}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1 \leq \frac{1}{\varepsilon^m} \int_{\mathcal{B}(Q_{\bar{x}}, 2\varepsilon)} q(x) d\mathcal{V}(x) + C\delta \quad (6.15)$$

with probability at least $1 - 2\exp(-c\delta^2 n\varepsilon^m)$. Using the fact that $\text{Exp}_{Q_{\bar{x}}} : \mathcal{B}(Q_{\bar{x}}, 2\varepsilon) \rightarrow B^m(0, 2\varepsilon)$ is a diffeomorphism, we acquire

$$\int_{\mathcal{B}(Q_{\bar{x}}, 2\varepsilon)} q(x) d\mathcal{V}(x) = \int_{B^m(0, 2\varepsilon) \subset T_{Q_{\bar{x}}} \mathcal{M}} \tilde{q}(v) J_{Q_{\bar{x}}}(v) dv \leq C_\eta \varepsilon^m \quad (6.16)$$

where, as outlined in Section 2.2, we apply the change of variables $x = \text{Exp}_{Q_{\bar{x}}}(v)$ and write $\tilde{q}(v) = q(\text{Exp}_{Q_{\bar{x}}}(v))$. Putting (6.16) back in (6.15), we conclude

$$\frac{1}{n\varepsilon^m} \sum_{Q_{\bar{x}'}: 0 < \|Q_{\bar{x}} - Q_{\bar{x}'}\| \leq 2\varepsilon} 1 \leq C_\eta + C\delta, \quad (6.17)$$

with probability at least $1 - 2\exp(-c\delta^2 n\varepsilon^m)$. By combining (6.17) with (6.7), (6.8), (6.9), and recalling that $\tau > 2$, we arrive at

$$\begin{aligned} & |\mathcal{L}_{\text{aug}} f(\bar{x}) - \mathcal{L}_{\text{proj}} f(Q_{\bar{x}})| \\ & \leq |\mathcal{L}_{\text{aug}} f(\bar{x}) - \mathcal{L}_{\bar{\mathcal{X}}} f(\bar{x})| + |\mathcal{L}_{\bar{\mathcal{X}}} f(\bar{x}) - \mathcal{L}_{Q_{\bar{\mathcal{X}}}} f(\bar{x})| + |\mathcal{L}_{Q_{\bar{\mathcal{X}}}} f(\bar{x}) - \mathcal{L}_{\text{proj}} f(\bar{x})| \\ & \leq C_\eta \varepsilon \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))} (1 + \delta), \end{aligned}$$

with probability at least $1 - 6\exp(-c\delta^2 n\varepsilon^m)$. The proposition can now be concluded by another application of union bound over $\bar{x} \in \bar{\mathcal{X}}$. $\blacksquare$

6.2 From Discrete to Nonlocal

We extend definition (6.2) to a function ς on $\mathcal{M} \times \mathcal{M}$, that is

$$\varsigma(x, y) := \int_{\{\|z-x\| < \varepsilon_n^{2/3}\} \cap \{\|z-y\| < \varepsilon_n^{2/3}\} \cap \mathcal{M}} \exp\left(-\frac{\|z-x\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|z-y\|^2}{2\varepsilon_w^2}\right) \rho(z) d\mathcal{V}(z). \quad (6.18)$$

Observe that the conclusion of Lemma 16 remains valid for this extension; namely

$$\sup_{x, y \in \mathcal{M}} \varsigma(x, y) \leq C\varepsilon_w^m. \quad (6.19)$$

Define a *nonlocal* operator $\mathbf{L}_{\mathcal{M}}$ such that for $f \in L^2(\mathcal{M})$,

$$\mathbf{L}_{\mathcal{M}}f(x) := \frac{1}{\varepsilon_w^m \varepsilon^{m+2}} \int_{B(x, \varepsilon)} \varsigma(x, y)(f(x) - f(y))q(y)d\mathcal{V}(y). \quad (6.20)$$

The augmentation graph Laplacian $\mathcal{L}_{\text{aug}}$ provides, with high probability, a pointwise approximation of the nonlocal operator $\mathbf{L}_{\mathcal{M}}$, as detailed below.

Lemma 17 *Let $\tau > 2$, $\varepsilon > 0$, and $\varepsilon_n = \varepsilon^{\tau+1}$. Let $\mathcal{L}_{\text{aug}}$ and $\mathbf{L}_{\mathcal{M}}$ be respectively defined as in (2.4) and in (6.20). Let $f \in \mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))$. Then for $\varepsilon \leq \delta \leq \varepsilon^{-1}$,*

$$\mathbb{P}\left(\max_{\bar{x} \in \bar{\mathcal{X}}} \left| \mathcal{L}_{\text{aug}}f(\bar{x}) - \mathbf{L}_{\mathcal{M}}f(Q_{\bar{x}}) \right| > C_{\eta}\delta \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}\right) \leq 8n \exp(-c_{\eta}\delta^2 n \varepsilon^{m+2}),$$

for some $C_{\eta}, c_{\eta} > 0$ depending on η and $\mathcal{M}$, ρ at most.

Proof [Proof of Lemma 17] We resume from the conclusion of the proof of Proposition 14. Let $\bar{x} \in \bar{\mathcal{X}}$. At this stage, we require a concentration result provided by Lemma 27, a variant of Lemma 28. We apply the concentration result of Lemma 27 to $\mathcal{L}_{\text{proj}}f(Q_{\bar{x}})$, with the bounded open set Ω in the lemma corresponding to the orthogonal projection of $B^d(\bar{x}, \varepsilon) \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})$ onto $\mathcal{M}$. Let $y \in \mathcal{N}_{\varepsilon_n}(\mathcal{M})$ be such that $\|y - \bar{x}\| = \varepsilon$, and let Q_y denote the orthogonal projection of y onto $\mathcal{M}$. Then similar to the calculation in (6.10),

$$\|Q_{\bar{x}} - Q_y\| \leq \varepsilon + 2\varepsilon_n = \varepsilon(1 + 2\varepsilon_w), \quad (6.21)$$

and moreover,

$$\|Q_{\bar{x}} - Q_y\| \geq \|\bar{x} - y\| - \|Q_{\bar{x}} - \bar{x}\| - \|y - Q_y\| \geq \varepsilon - 2\varepsilon_n = \varepsilon(1 - 2\varepsilon_w). \quad (6.22)$$

It follows from (6.21), (6.22) that Ω , as the orthogonal projection of $B^d(\bar{x}, \varepsilon) \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})$ onto $\mathcal{M}$, satisfies

$$B^d(Q_{\bar{x}}, \varepsilon(1 - 2\varepsilon_w)) \cap \mathcal{M} \subset \Omega \subset B^d(Q_{\bar{x}}, \varepsilon(1 + 2\varepsilon_w)) \cap \mathcal{M}, \quad (6.23)$$

and thus, by considering (2.7), (2.11),

$$c\varepsilon^m \leq \mathcal{V}(\Omega) \leq C\varepsilon^m. \quad (6.24)$$

We now apply Lemma 27 to $\mathcal{L}_{\text{proj}}f(Q_{\bar{x}})$, taking into account Proposition 7 and (6.19), (6.24), which gives us, for $0 < \delta \leq 1$,

$$\left| \varepsilon^m \mathcal{L}_{\text{proj}}f(Q_{\bar{x}}) - \frac{1}{\varepsilon_w^m \varepsilon^2} \int_{\Omega \cap \mathcal{M}} \varsigma(Q_{\bar{x}}, y)(f(Q_{\bar{x}}) - f(y))q(y)d\mathcal{V}(y) \right| \leq C_{\eta} \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))} \varepsilon^m \delta \varepsilon^{-1} \quad (6.25)$$

with probability at least $1 - 2 \exp(-c_{\eta}\delta^2 n \varepsilon^m)$. Similarly, from (2.7), (2.10), as well as Proposition 7, (6.19), (6.23) again, we get

$$\begin{aligned} \frac{1}{\varepsilon_w^m \varepsilon^2} \left| \int_{\Omega \Delta (B^d(Q_{\bar{x}}, \varepsilon) \cap \mathcal{M})} \varsigma(Q_{\bar{x}}, y)(f(Q_{\bar{x}}) - f(y))q(y)d\mathcal{V}(y) \right| \\ \leq \frac{C_{\eta} \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{\varepsilon} \left| \int_{\Omega \Delta (B^d(Q_{\bar{x}}, \varepsilon) \cap \mathcal{M})} d\mathcal{V}(y) \right| \\ \leq \frac{C_{\eta} \varepsilon_w \varepsilon^m \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))}}{\varepsilon} \\ = C_{\eta} \|f\|_{\mathcal{C}^1(\mathcal{N}_{\varepsilon_n}(\mathcal{M}))} \varepsilon^{\tau-1} \varepsilon^m. \end{aligned} \quad (6.26)$$

Above, we use Δ to denote the set symmetric difference. Therefore, combining (6.25), (6.26), yields

$$\left| \mathcal{L}_{\text{proj}} f(Q_{\bar{x}}) - \frac{1}{\varepsilon_{\text{w}}^m \varepsilon^{m+2}} \int_{B^d(Q_{\bar{x}}, \varepsilon) \cap \mathcal{M}} \varsigma(Q_{\bar{x}}, y) (f(Q_{\bar{x}}) - f(y)) q(y) d\mathcal{V}(y) \right| \leq C_{\eta} \|f\|_{C^1(\mathcal{N}_{\varepsilon_{\text{n}}}(\mathcal{M}))} \max\{\delta \varepsilon^{-1}, \varepsilon^{\tau-1}\}, \quad (6.27)$$

with probability at least $1 - 2 \exp(-c_{\eta} \delta^2 n \varepsilon^m)$. Furthermore, as part of the conclusion of Lemma 28 (or Lemma 26) and (6.19),

$$\begin{aligned} \frac{1}{\varepsilon_{\text{w}}^m \varepsilon^{m+2}} \left| \int_{B^d(Q_{\bar{x}}, \varepsilon) \cap \mathcal{M}} \varsigma(Q_{\bar{x}}, y) (f(Q_{\bar{x}}) - f(y)) q(y) d\mathcal{V}(y) \right. \\ \left. - \int_{B(Q_{\bar{x}}, \varepsilon)} \varsigma(Q_{\bar{x}}, y) (f(Q_{\bar{x}}) - f(y)) q(y) d\mathcal{V}(y) \right| \leq C \|f\|_{C^1(\mathcal{N}_{\varepsilon_{\text{n}}}(\mathcal{M}))} \varepsilon. \end{aligned} \quad (6.28)$$

Now recall from the proof of Proposition 14 that for $\varepsilon^2 \leq \delta \leq 1$,

$$|\mathcal{L}_{\text{aug}} f(\bar{x}) - \mathcal{L}_{\text{proj}} f(Q_{\bar{x}})| \leq C \varepsilon \|f\|_{C^1(\mathcal{N}_{\varepsilon_{\text{n}}}(\mathcal{M}))} (1 + \delta) \quad (6.29)$$

with probability at least $1 - 6 \exp(-c \delta^2 n \varepsilon^m)$. Recall that $\tau > 2$. Then by setting $\delta' = \delta / \varepsilon$ for $\varepsilon^2 \leq \delta \leq 1$, and using (6.27), (6.28), (6.29), along with the definition (6.20), we obtain

$$|\mathcal{L}_{\text{aug}} f(\bar{x}) - \mathbf{L}_{\mathcal{M}} f(Q_{\bar{x}})| \leq C_{\eta} \|f\|_{C^1(\mathcal{N}_{\varepsilon_{\text{n}}}(\mathcal{M}))} \delta',$$

for $\varepsilon \leq \delta' \leq \varepsilon^{-1}$, with probability at least $1 - 8 \exp(-c_{\eta} (\delta')^2 n \varepsilon^{m+2})$. Finally, we conclude the proof with an application of union bound over $\bar{x} \in \mathcal{X}$. $\blacksquare$

6.3 Proof of Theorem 11

Proof Lemma 17 establishes that, with high probability $\mathcal{L}_{\text{aug}}$ behaves pointwise in close approximation of $\mathbf{L}_{\mathcal{M}}$. Therefore, to prove Theorem 11, it is enough to show that $\mathbf{L}_{\mathcal{M}}$ is approximately $\alpha \beta \Delta_{\text{aug}}$, with a small additive error. In the ensuing proof, we will employ the equivalence demonstrated in (2.13).

We define, for $x, y \in \mathcal{M}$,

$$\varsigma^*(x, y) := \int_{\{z \in \mathbb{R}^d : d(x, z) < \varepsilon_{\text{n}}^{2/3} \wedge d(y, z) < \varepsilon_{\text{n}}^{2/3}\} \cap \mathcal{M}} \exp\left(-\frac{d(x, z)^2}{2\varepsilon_{\text{w}}^2}\right) \exp\left(-\frac{d(y, z)^2}{2\varepsilon_{\text{w}}^2}\right) \rho(z) d\mathcal{V}(z),$$

the version of $\varsigma(x, y)$ in (6.18) based exclusively on geodesic distances. The following lemma, whose proof is given in Appendix E, applies.

Lemma 18 *Let ς, ς^* be defined as above. Then for $x, y \in \mathcal{M}$,*

$$\varsigma(x, y) = \varsigma^*(x, y) + \mathcal{O}(\varepsilon_{\text{w}}^m \varepsilon^2).$$

Leveraging Lemma 18, we can substitute $\varsigma(x, y)$ with $\varsigma^*(x, y)$ in the expression (6.20) for $\mathbf{L}_{\mathcal{M}}f(x)$. More specifically, since

$$\begin{aligned} \frac{1}{\varepsilon_{\mathbf{w}}^m \varepsilon^{m+2}} \left| \int_{\mathcal{B}(x, \varepsilon)} (\varsigma(x, y) - \varsigma^*(x, y))(f(x) - f(y))q(y) d\mathcal{V}(y) \right| \\ \leq \frac{C \varepsilon_{\mathbf{w}}^m \varepsilon^2}{\varepsilon_{\mathbf{w}}^m \varepsilon^{m+2}} \int_{\mathcal{B}(x, \varepsilon)} |f(x) - f(y)| q(y) d\mathcal{V}(y) \\ \leq \frac{C_{\eta} \|f\|_{C^1(\mathcal{B}(x, \varepsilon))} \varepsilon}{\varepsilon^m} \int_{\mathcal{B}(x, \varepsilon)} d\mathcal{V}(y) \\ \leq C_{\eta} \|f\|_{C^1(\mathcal{B}(x, \varepsilon))} \varepsilon, \end{aligned} \quad (6.30)$$

we get

$$\mathbf{L}_{\mathcal{M}}f(x) = \frac{1}{\varepsilon_{\mathbf{w}}^m \varepsilon^{m+2}} \int_{\mathcal{B}(x, \varepsilon)} \varsigma^*(x, y)(f(x) - f(y))q(y) d\mathcal{V}(y) + \mathcal{O}_{f, \eta}(\varepsilon).$$

For $y \in \mathcal{B}(x, \varepsilon)$, we write $\varsigma_x^*(y) := \varsigma^*(x, y)$. Following the change of variables and notation modification introduced in Section 2.2, we let $y \in \mathcal{B}(x, \varepsilon)$ correspond uniquely to $v \in B^m(0, \varepsilon)$ such that $\text{Exp}_x(v) = y$, and so $\text{Exp}_x(0) = x$. We further write $\tilde{\varsigma}_0^*(v) = \varsigma_x^*(y)$, $\tilde{f}(v) = f(y)$, $\tilde{q}(v) = q(y)$. Thus

$$\mathbf{L}_{\mathcal{M}}\tilde{f}(0) = -\frac{1}{\varepsilon_{\mathbf{w}}^m \varepsilon^{m+2}} \int_{B^m(0, \varepsilon) \subset T_x \mathcal{M}} \tilde{\varsigma}_0^*(v)(\tilde{f}(v) - \tilde{f}(0))\tilde{q}(v)J_x(v) dv + \mathcal{O}_{f, \eta}(\varepsilon). \quad (6.31)$$

From Taylor's expansion, we obtain the following result, with a proof provided in Appendix F.

Lemma 19 *Let $v \in B^m(0, \varepsilon)$. Then*

$$\tilde{\varsigma}_0^*(0) = \varepsilon_{\mathbf{w}}^m (\alpha \rho(x) + \mathcal{O}(\varepsilon)). \quad (6.32)$$

Moreover,

$$\nabla \tilde{\varsigma}_0^*(0) = 0 \quad \text{and} \quad \|D^2 \tilde{\varsigma}_0^*(0)\| = \mathcal{O}(\varepsilon_{\mathbf{w}}^m), \quad (6.33)$$

and in particular,

$$\tilde{\varsigma}_0^*(v) = \tilde{\varsigma}_0^*(0) + \mathcal{O}(\varepsilon_{\mathbf{w}}^m \varepsilon^2).$$

Utilizing Lemma 19, along with (2.13), we see that, as per the computations presented in (6.30),

$$\begin{aligned} \frac{1}{\varepsilon_{\mathbf{w}}^m \varepsilon^{m+2}} \left| \int_{B^m(0, \varepsilon) \subset T_x \mathcal{M}} (\tilde{\varsigma}_0^*(v) - \tilde{\varsigma}_0^*(0))(\tilde{f}(v) - \tilde{f}(0))\tilde{q}(v)J_x(v) dv \right| \\ \leq \frac{C}{\varepsilon^m} \int_{B^m(0, \varepsilon)} |\tilde{f}(v) - \tilde{f}(0)| dv \leq C \|f\|_{C^1(\mathcal{B}(x, \varepsilon))} \varepsilon. \end{aligned} \quad (6.34)$$

Consequently, from (6.31), (6.34),

$$\begin{aligned} \mathbf{L}_{\mathcal{M}}\tilde{f}(0) &= -\frac{\tilde{\varsigma}_0^*(0)}{\varepsilon_{\mathbf{w}}^m \varepsilon^{m+2}} \int_{B^m(0, \varepsilon) \subset T_x \mathcal{M}} (\tilde{f}(v) - \tilde{f}(0))\tilde{q}(v)J_x(v) dv + \mathcal{O}_{f, \eta}(\varepsilon) \\ &= -\frac{\tilde{\varsigma}_0^*(0)}{\varepsilon_{\mathbf{w}}^m \varepsilon^2} \int_{B^m(0, 1)} (\tilde{f}(\varepsilon v) - \tilde{f}(0))\tilde{q}(\varepsilon v)J_x(\varepsilon v) dv + \mathcal{O}_{f, \eta}(\varepsilon). \end{aligned} \quad (6.35)$$

Using Taylor's expansion once again, and considering Proposition 7, we acquire

$$\begin{aligned}\tilde{f}(\varepsilon v) &= \tilde{f}(0) + \varepsilon \langle \nabla \tilde{f}(0), v \rangle + \frac{1}{2} \varepsilon^2 \langle D^2 \tilde{f}(0) v, v \rangle + \mathcal{O}_{f,\eta}(\varepsilon^3) \\ \tilde{q}(\varepsilon v) &= \tilde{q}(0) + \varepsilon \langle \nabla \tilde{q}(0), v \rangle + \mathcal{O}_\eta(\varepsilon^2).\end{aligned}\tag{6.36}$$

Moreover, from (2.10),

$$J_x(\varepsilon v) = 1 + \mathcal{O}(\varepsilon^2).\tag{6.37}$$

Then putting (6.36) in (6.35), together with simple calculations, delivers

$$\mathbf{L}_\mathcal{M} \tilde{f}(v) = \frac{\tilde{\zeta}_0^*(0)}{\varepsilon_w^m} (\mathbf{L}_1 + \mathbf{L}_2 + \mathbf{L}_3) + \mathcal{O}_{f,\eta}(\varepsilon),\tag{6.38}$$

where

$$\begin{aligned}\mathbf{L}_1 &:= -\frac{\tilde{q}(0)}{\varepsilon} \int_{B^m(0,1)} \langle \nabla \tilde{f}(0), v \rangle J_x(\varepsilon v) dv \\ \mathbf{L}_2 &:= -\int_{B^m(0,1)} \langle \nabla \tilde{f}(0), v \rangle \langle \nabla \tilde{q}(0), v \rangle J_x(\varepsilon v) dv \\ \mathbf{L}_3 &:= -\frac{\tilde{q}(0)}{2} \int_{B^m(0,1)} \langle D^2 \tilde{f}(0) v, v \rangle J_x(\varepsilon v) dv.\end{aligned}$$

Observe, due to the symmetry of odd function,

$$\int_{B^m(0,1)} \langle \nabla \tilde{f}(0), v \rangle dv = 0,$$

which, combined with (6.37), subsequently gives us

$$\mathbf{L}_1 = \mathcal{O}_{f,\eta}(\varepsilon).\tag{6.39}$$

Regarding $\mathbf{L}_2$, $\mathbf{L}_3$, we also get from (6.37) and a standard computation,

$$\mathbf{L}_2 = -\int_{B^m(0,1)} \langle \nabla \tilde{f}(0), v \rangle \langle \nabla \tilde{q}(0), v \rangle J_x(\varepsilon v) dv = -\beta \langle \nabla \tilde{f}(0), \nabla \tilde{q}(0) \rangle + \mathcal{O}_f(\varepsilon^2),\tag{6.40}$$

and

$$\mathbf{L}_3 = -\frac{\tilde{q}(0)}{2} \int_{B^m(0,1)} \langle D^2 \tilde{f}(0) v, v \rangle J_x(\varepsilon v) dv = -\frac{\beta \tilde{q}(0)}{2} \Delta_\mathcal{M} \tilde{f}(0) + \mathcal{O}_f(\varepsilon^2).\tag{6.41}$$

By combining (6.38), (6.39), (6.40), (6.41), and applying Lemma 19, we arrive at

$$\begin{aligned}\mathbf{L}_\mathcal{M} \tilde{f}(0) &= -\frac{\tilde{\zeta}_0^*(0)\beta}{\varepsilon_w^m} (\langle \nabla \tilde{f}(0), \nabla \tilde{q}(0) \rangle + \frac{\tilde{q}(0)}{2} \Delta_\mathcal{M} \tilde{f}(0)) + \mathcal{O}_{f,\eta}(\varepsilon) \\ &= -\alpha\beta\rho(x) (\langle \nabla \tilde{f}(0), \nabla \tilde{q}(0) \rangle + \frac{\tilde{q}(0)}{2} \Delta_\mathcal{M} \tilde{f}(0)) + \mathcal{O}_{f,\eta}(\varepsilon),\end{aligned}$$

which simplifies to

$$\mathbf{L}_{\mathcal{M}}f(x) = \mathbf{L}_{\mathcal{M}}\tilde{f}(0) = -\frac{\alpha\beta\rho(x)}{2\tilde{q}(0)}\operatorname{div}(\tilde{q}^2\nabla\tilde{f})(0) + \mathcal{O}_{f,\eta}(\varepsilon),$$

or equivalently,

$$\mathbf{L}_{\mathcal{M}}f(x) = -\frac{\alpha\beta\rho(x)}{2q(x)}\operatorname{div}(q^2\nabla f)(x) + \mathcal{O}_{f,\eta}(\varepsilon). \quad (6.42)$$

Recall from Proposition 7 that $\|\rho - q\|_{L^\infty(\mathcal{M})} \leq C_\eta \varepsilon_n^{2/3}$. Thus, we have, subsequently from (6.42),

$$\mathbf{L}_{\mathcal{M}}f(x) = -\frac{\alpha\beta}{2}\operatorname{div}(q^2\nabla f)(x) + \mathcal{O}_{f,\eta}(\varepsilon) = -\alpha\beta\Delta_{\text{aug}}f(x) + \mathcal{O}_{f,\eta}(\varepsilon). \quad (6.43)$$

Finally, to complete the proof, it remains to integrate (6.43) with Lemma 17. ■

7. Spectral Consistency of Augmentation Graph Laplacian

7.1 Outline for the Proof of Theorem 12

With Theorem 11 established, the proof of Theorem 12 is primarily influenced by and aligned with the strategy in (Calder and García Trillos, 2022). Before proceeding with further discussion, however, it is necessary to delineate two key differences between our setting and that in the aforementioned work. For the first difference, the augmentation graph $G_{\bar{\mathcal{X}}}$ is not built from points sampled directly on the manifold $\mathcal{M}$, meaning that, for example, ϑ_n in (3.10) is not a measure on $\mathcal{M}$. Nevertheless, the groundwork in Section 3.1 ensures the construction of a unique projection point $Q_{\bar{x}} \in \mathcal{M}$ for each $\bar{x} \in \bar{\mathcal{X}}$. This bijective correspondence, guaranteed with probability 1 by Lemma 6, provides a matching between $\bar{\mathcal{X}} = \{\bar{x}_1, \dots, \bar{x}_n\}$ and the set $\mathcal{Q}_{\bar{\mathcal{X}}} = \{Q_{\bar{x}_1}, \dots, Q_{\bar{x}_n}\}$, thus enabling the transfer between functions on the graph $G_{\bar{\mathcal{X}}}$ and those on the graph $G_{\mathcal{Q}_{\bar{\mathcal{X}}}}$, as defined at the beginning of Section 6.1, and vice versa. In particular, Proposition 14 establishes, with high probability, the pointwise proximity between the augmentation graph Laplacian $\mathcal{L}_{\text{aug}}$ (2.4) and the projected graph Laplacian $\mathcal{L}_{\text{proj}}$ (6.3). For the second difference, the projection points $Q_{\bar{x}} \sim \nu$, which is not a probability measure on $\mathcal{M}$; see (3.1). However, as shown in Proposition 7, q closely approximates the sampling density ρ of μ , which is a probability measure. Therefore, with $\mathcal{L}_{\text{proj}}$ serving as a substitute for $\mathcal{L}_{\text{aug}}$, and ρ replacing q , the theory developed in (Calder and García Trillos, 2022) can be directly applied, with some additional considerations. We outline the application of this theory next.

Define a discrete probability measure $\check{\vartheta}_n$ to be

$$\check{\vartheta}_n := \frac{1}{n} \sum_{i=1}^n \delta_{Q_{\bar{x}_i}},$$

which is a measure on $\mathcal{M}$ with support on $\mathcal{Q}_{\bar{\mathcal{X}}}$. Further, we define a graph-based Dirichlet energy, on the graph $G_{\mathcal{Q}_{\bar{\mathcal{X}}}}$, such that for $f \in L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$,

$$\begin{aligned} E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f) &:= \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i,j=1}^n \varsigma(Q_{\bar{x}_i}, Q_{\bar{x}_j}) \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} (f(Q_{\bar{x}_i}) - f(Q_{\bar{x}_j}))^2 \\ &= 2 \langle \mathcal{L}_{\text{proj}} f, f \rangle_{L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})}, \end{aligned} \quad (7.1)$$

where

$$\langle f, g \rangle_{L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})} := \frac{1}{n} \sum_{i=1}^n f(Q_{\bar{x}_i}) g(Q_{\bar{x}_i})$$

is an inner product on $L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$. Then $\mathcal{L}_{\text{proj}}$ is a positive semi-definite and self-adjoint operator on $L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$, whose eigenvalues are given by

$$\hat{\lambda}_l^{\text{proj}} = \min_{S \in \mathbb{G}_{\mathcal{Q}_{\bar{\mathcal{X}}}}^l} \max_{f \in S \setminus \{0\}} \frac{E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f)}{\|f\|_{L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})}^2}, \quad (7.2)$$

where $\mathbb{G}_{\mathcal{Q}_{\bar{\mathcal{X}}}}^l$ denotes the Grassmannian manifold of all linear subspaces of $L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$ of dimension l , and $l = 1, \dots, n$. An orthonormal basis of $L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$ can be formed by the associated eigenvectors $\hat{f}_l^{\text{proj}}$ of $\mathcal{L}_{\text{proj}}$. Establishing a quantitative interpolation between the graph-based Dirichlet energy $E_{\mathcal{Q}_{\bar{\mathcal{X}}}}$ and the continuum energy $E_{\mathcal{M}}$ (3.15) is critical for our spectral convergence analysis. Such interpolation can be realized through two maps $P : L^2(\mathcal{M}) \rightarrow L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$ and $\mathcal{I} : L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}}) \rightarrow L^2(\mathcal{M})$, which act as near-isometries when restricted to functions with finite Dirichlet energy (in the discrete and continuum settings). Precisely, these maps further satisfy estimates of the form

$$\begin{aligned} E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(Pf) &\leq \alpha\beta(1 + \text{small error})E_{\mathcal{M}}(f), \quad \forall f \in L^2(\mathcal{M}), \\ \alpha\beta E_{\mathcal{M}}(\mathcal{I}f) &\leq (1 + \text{small error})E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f) \quad \forall f \in L^2_{\check{\vartheta}_n}(\mathcal{Q}_{\bar{\mathcal{X}}}); \end{aligned}$$

see Proposition 21 in the next subsection. When combined with the result of Theorem 11 and the variational identities (7.2), (3.16), these inequalities translate to proximity between the spectra of $\mathcal{L}_{\text{proj}}$ and $\alpha\beta\Delta_{\text{aug}}$, following the steps in (Calder and García Trillos, 2022). Then the final step of our analysis involves leveraging Proposition 14 and (3.12) to demonstrate the spectral alignment of $\mathcal{L}_{\text{aug}}$ and $\alpha\beta\Delta_{\text{aug}}$.

For the remainder of this subsection, we outline the construction of the maps $P, \mathcal{I}$, which are constructed using an ∞ -optimal transport (∞ -OT) map between $\check{\vartheta}_n$ and a suitably chosen probability measure on $\mathcal{M}$, provided below.

Proposition 20 *Let $\iota_1, \iota_2 > 0$ satisfy¹²:*

1. $n^{-1/m} < \iota_2 \leq 1$;

12. The proof of (Calder and García Trillos, 2022, Proposition 2.12) only requires that ι_2 be of small, fixed scale, greater than $n^{-1/m}$.

2. $\iota_1 + \iota_2 \leq c$, where $c > 0$ depends on $\mathcal{M}$, ρ .

Then with probability at least $1 - n \exp(-c' n \iota_1^2 \iota_2^m)$, there exist a probability measure v_n with density $p_n : \mathcal{M} \rightarrow \mathbb{R}$, such that

$$\min_{S_{\#} v_n = \check{v}_n} \sup_{x \in \mathcal{M}} d(x, S(x)) \leq \iota_2,$$

and that

$$\|p_n - \rho\|_{L^\infty(\mathcal{M})} \leq C(\iota_1 + \iota_2).$$

Here, the constant $C > 0$ depends on $\mathcal{M}$, ρ , and the constant $c' > 0$ depends solely on¹³ $\mathcal{M}$.

This proposition directly follows from (Calder and García Trillos, 2022, Proposition 2.12) and its proof.

We denote by T an ∞ -OT map between $\check{v}_n$ and v_n given in Proposition 20. Let $U_i := T^{-1}(\{Q_{\bar{x}_i}\})$. We define an *adaptive* contraction discretization map $P : L^2(\mathcal{M}) \rightarrow L^2_{\check{v}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$ by

$$Pf(Q_{\bar{x}_i}) := \frac{n \int_{U_i} f(x) p_n(x) d\mathcal{V}(x)}{\rho(Q_{\bar{x}_i})^{1/2}}, \quad (7.3)$$

as well as an adaptive extension map $P^* : L^2_{\check{v}_n}(\mathcal{Q}_{\bar{\mathcal{X}}}) \rightarrow L^2(\mathcal{M})$, by

$$P^*f(x) := \sum_{i=1}^n f(Q_{\bar{x}_i}) \rho(Q_{\bar{x}_i})^{1/2} \mathbb{1}_{U_i}(x). \quad (7.4)$$

Next, we describe the construction of an *interpolation* map $\mathcal{I} : L^2_{\check{v}_n}(\mathcal{Q}_{\bar{\mathcal{X}}}) \rightarrow H^1(\mathcal{M}) \subset L^2(\mathcal{M})$. It will be apparent below that such a map inherently possesses a smoothing property. Let $\phi : [0, \infty) \rightarrow [0, \infty)$ be a non-increasing function of choice with support on the interval $[0, 1]$, whose restriction to $[0, 1]$ is Lipschitz continuous and non-vanishing. Let $\psi : [0, \infty) \rightarrow [0, \infty)$ be a smooth function, given by

$$\psi(t) := \frac{1}{\beta} \int_t^\infty \phi(s) ds,$$

where we recall β is defined in (4.2).

Lastly, for a *bandwidth* $r > 0$, let $\Lambda_r : L^2(\mathcal{M}) \rightarrow H^1(\mathcal{M})$ be a convolution operator with kernel

$$K_r(x, y) := \frac{1}{r^m} \psi\left(\frac{d(x, y)}{r}\right),$$

such that

$$\Lambda_r f(x) := \left(\int_{\mathcal{M}} K_r(x, y) d\mu(y) \right)^{-1} \int_{\mathcal{M}} K_r(x, y) f(y) d\mu(y).$$

We then build an interpolation map $\mathcal{I} : L^2_{\check{v}_n}(\mathcal{Q}_{\bar{\mathcal{X}}}) \rightarrow H^1(\mathcal{M})$ with bandwidth $\varepsilon - 2\iota_2$, specifically

$$\mathcal{I}f(x) := \Lambda_{\varepsilon - 2\iota_2} P^*f(x). \quad (7.5)$$

13. specifically on its dimension m .

7.2 Proof of Theorem 12

Proof The beginning of the proof follows almost the same approach as that of (Calder and García Trillos, 2022, Theorem 2.4 (i)), and, subsequently, that of (Calder and García Trillos, 2022, Theorem 2.7 (ii)), with some modifications. To further clarify, we introduce Propositions 21, 22 below. They correspond to (Calder and García Trillos, 2022, Propositions 4.1, 4.2), respectively, which are the two main components for the proofs of (Calder and García Trillos, 2022, Theorems 2.4, 2.7). We recall the discretization map P defined in (7.3) and the interpolation map $\mathcal{I}$ defined in (7.5).

Proposition 21 *Let $\varepsilon > 0$ and $\iota_1, \iota_2 > 0$ satisfy Assumption 10. Then both of the following hold with probability at least $1 - n \exp(-cn\iota_1^2\iota_2^m)$,*

$$E_{\mathcal{Q}_{\bar{x}}}(Pf) \leq \alpha\beta \left(1 + C_\eta \left(\frac{\iota_2}{\varepsilon} + \varepsilon + \iota_1\right)\right) E_{\mathcal{M}}(f), \quad \forall f \in L^2(\mathcal{M}), \quad (7.6)$$

$$\alpha\beta E_{\mathcal{M}}(\mathcal{I}f) \leq \left(1 + C_\eta \left(\frac{\iota_2}{\varepsilon} + \varepsilon + \iota_1\right)\right) E_{\mathcal{Q}_{\bar{x}}}(f), \quad \forall f \in L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{x}}), \quad (7.7)$$

where α and β are defined in (4.1) and (4.2), respectively; the constants $C_\eta > 0$ depend on η and $\mathcal{M}$, ρ , and the constant $c > 0$ depends solely on $\mathcal{M}$.

Proposition 22 *Let $\varepsilon > 0$ and $\iota_1, \iota_2 > 0$ satisfy:*

1. $n^{-1/m} < \iota_2 \leq \frac{\varepsilon}{4}$;
2. $\iota_1 + \iota_2 \leq c$, where $c > 0$ depends on $\mathcal{M}$, ρ .

Then both of the following hold with probability at least $1 - n \exp(-c'n\iota_1^2\iota_2^m)$,

$$\left| \|f\|_{L^2(\mathcal{M})}^2 - \|Pf\|_{L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{x}})}^2 \right| \leq C\iota_2 \|f\|_{L^2(\mathcal{M})} \sqrt{E_{\mathcal{M}}(f)} + C(\iota_1 + \iota_2) \|f\|_{L^2(\mathcal{M})}^2, \quad \forall f \in L^2(\mathcal{M}), \quad (7.8)$$

$$\left| \|f\|_{L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{x}})}^2 - \|\mathcal{I}f\|_{L^2(\mathcal{M})}^2 \right| \leq C\varepsilon \|f\|_{L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{x}})} \sqrt{E_{\mathcal{Q}_{\bar{x}}}(f)} + C(\iota_1 + \iota_2) \|f\|_{L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{x}})}^2, \quad \forall f \in L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{x}}). \quad (7.9)$$

Here, all the constants $C > 0$ depend on $\mathcal{M}$, ρ , and the constant $c' > 0$ depends solely on $\mathcal{M}$.

Note that these are the near-isometric inequalities discussed at the beginning of Section 7.1.

The proof of Proposition 21 mirrors that of (Calder and García Trillos, 2022, Proposition 4.1), and the proof of Proposition 22 mirrors that of (Calder and García Trillos, 2022, Proposition 4.2). The former, given in Section 7.3, is more involved than the latter, outlined in Section 7.4. One modification in these proofs arises from our definition of discretization and extension maps, (7.3) and (7.4), which explicitly incorporate a local density factor. This definition is guided by the specific formulation (7.1) of the energy $E_{\mathcal{Q}_{\bar{x}}}$ and contrasts with the discretization and extension maps in (Calder and García Trillos, 2022). For instance, see the proof of Lemma 24 for how this difference is addressed. Another modification stems

from the fact that an edge $(Q_{\bar{x}}, Q_{\bar{x}'})$ in the graph $G_{\mathcal{Q}_{\bar{\mathcal{X}}}}$ is designed to correspond to the edge $(\bar{x}, \bar{x}')$ in the graph $G_{\bar{\mathcal{X}}}$, thus causing the Euclidean distance between $Q_{\bar{x}}, Q_{\bar{x}'}$ to vary from $\varepsilon(1 - 2\varepsilon_w)$ to $\varepsilon(1 + 2\varepsilon_w)$. This differs from the ε - and k NN graphs considered in (Calder and García Trillos, 2022), where all edges have a fixed length. To handle this, we will take advantage of the negligible difference between $\varepsilon(1 - 2\varepsilon_w) \approx \varepsilon \approx \varepsilon(1 + 2\varepsilon_w)$, and refer to the technical result in (García Trillos et al., 2020, Lemma 5) for further details.

Given Propositions 21, 22, we can follow the reasoning in (Calder and García Trillos, 2022, Theorem 2.4 (i)) (minmax principle) to conclude that, for the first step, if

$$\sqrt{\lambda_l} \iota_2 + C(\iota_1 + \iota_2) \leq \sqrt{\lambda_l} \varepsilon + C(\iota_1 + \iota_2) \leq \frac{1}{2l}, \quad (7.10)$$

then with probability at least $1 - n \exp(-cn \iota_1^2 \iota_2^m)$,

$$|\lambda_l - \alpha \beta \hat{\lambda}_l^{\text{proj}}| \leq C(\varepsilon(\sqrt{\lambda_l} + 1) + \varepsilon + \iota_1) \lambda_l. \quad (7.11)$$

Note that the probability derived here matches the probability that the conclusions of Propositions 21 and 22 hold, which, in turn, is guaranteed whenever the conclusion of Proposition 20 holds. For the second step, we suppose in addition that

$$C(\varepsilon(\sqrt{\lambda_l} + 1) + \varepsilon + \iota_1) \lambda_l \leq \gamma_{\lambda_l} \quad (7.12)$$

holds. We infer from the reasoning near the end of the proof in (Calder and García Trillos, 2022, Theorem 2.7(ii)) (Davis-Kahan-type expansion), and the proof of Lemma 17, the existence of an eigenfunction f_l of Δ_{aug} , with probability at least $1 - n \exp(-cn \iota_1^2 \iota_2^m) - 2n \exp(-c_\eta n \varepsilon^{m+4})$, such that

$$\|f_l - \hat{f}_l^{\text{proj}}\|_{L_{\vartheta_n}^2(\mathcal{Q}_{\bar{\mathcal{X}}})} \leq C_{\eta, l} \varepsilon. \quad (7.13)$$

The additional reduction in probability stems from the pointwise consistency probability, transitioning from $\mathcal{L}_{\text{proj}}$ to the nonlocal operator $\mathbf{L}_{\mathcal{M}}$ (6.20) (with $\delta = \varepsilon$) in the proof of Lemma 17, and ultimately to Δ_{aug} .

With (7.13), we complete the first part of our analysis, which establishes the spectral alignment between $\mathcal{L}_{\text{proj}}$ and Δ_{aug} , as discussed at the beginning of Section 7.1. In order to further extend (7.13) to a conclusion about the spectral proximity between $\mathcal{L}_{\text{aug}}$ and $\alpha \beta \Delta_{\text{aug}}$, we define a *matching* map M that sends $\bar{x} \in \bar{\mathcal{X}}$ to $Q_{\bar{x}} \in \mathcal{Q}_{\bar{\mathcal{X}}}$. By Lemma 6, this map M establishes a bijection between $\bar{\mathcal{X}}$ and $\mathcal{Q}_{\bar{\mathcal{X}}}$ with probability 1. Thus, if $f \in L_{\vartheta_n}^2(\mathcal{Q}_{\bar{\mathcal{X}}})$, then $f \circ M \in L_{\vartheta_n}^2(\bar{\mathcal{X}})$, and conversely, if $f \in L_{\vartheta_n}^2(\bar{\mathcal{X}})$, then $f \circ M^{-1} \in L_{\vartheta_n}^2(\mathcal{Q}_{\bar{\mathcal{X}}})$. Analogous to Propositions 21, 22, we have the following proposition.

Proposition 23 *Under Assumption 5, it holds with probability 1 that*

$$\|f\|_{L_{\vartheta_n}^2(\bar{\mathcal{X}})} = \|f \circ M^{-1}\|_{L_{\vartheta_n}^2(\mathcal{Q}_{\bar{\mathcal{X}}})}, \quad \forall f \in L_{\vartheta_n}^2(\bar{\mathcal{X}}), \quad (7.14)$$

$$\|f\|_{L_{\vartheta_n}^2(\mathcal{Q}_{\bar{\mathcal{X}}})} = \|f \circ M\|_{L_{\vartheta_n}^2(\bar{\mathcal{X}})}, \quad \forall f \in L_{\vartheta_n}^2(\mathcal{Q}_{\bar{\mathcal{X}}}). \quad (7.15)$$

Moreover, let $\varepsilon > 0$ be sufficiently small to satisfy (3.17). Then both of the following hold with probability at least $1 - 2n \exp(-c_\eta n \varepsilon^{m+4})$,

$$E_{\bar{\mathcal{X}}}(f) \leq (1 + C\varepsilon^2)E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f \circ \mathbf{M}^{-1}) + C_\eta \varepsilon \|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2, \quad \forall f \in L^2_{\vartheta_n}(\bar{\mathcal{X}}), \quad (7.16)$$

$$E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f) \leq (1 + C\varepsilon^2)E_{\bar{\mathcal{X}}}(f \circ \mathbf{M}), \quad \forall f \in L^2_{\vartheta_n}(\mathcal{Q}_{\bar{\mathcal{X}}}). \quad (7.17)$$

Here, the constants $C_\eta, c_\eta > 0$ depend on η and $\mathcal{M}$, ρ at most, while the constants $C > 0$ in (7.16), (7.17) depend on $\mathcal{M}$.

The proof of Proposition 23 is presented in Section 7.5. By combining (7.6), (7.16), we obtain, for $f \in L^2(\mathcal{M})$,

$$E_{\bar{\mathcal{X}}}(Pf \circ \mathbf{M}) \leq \alpha\beta(1 + C\varepsilon^2) \left(1 + C_\eta \left(\frac{\iota_2}{\varepsilon} + \varepsilon + \iota_1\right)\right) E_{\mathcal{M}}(f) + C_\eta \varepsilon \|Pf \circ \mathbf{M}\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2, \quad (7.18)$$

and by combining (7.7), (7.17), we get, for $f \in L^2_{\vartheta_n}(\bar{\mathcal{X}})$,

$$\alpha\beta E_{\mathcal{M}}(\mathcal{I}(f \circ \mathbf{M}^{-1})) \leq (1 + C\varepsilon^2) \left(1 + C_\eta \left(\frac{\iota_2}{\varepsilon} + \varepsilon + \iota_1\right)\right) E_{\bar{\mathcal{X}}}(f). \quad (7.19)$$

Similarly, by combining (7.8) with (7.15) and (7.9) with (7.14), (7.17), we respectively derive

$$\left| \|f\|_{L^2(\mathcal{M})}^2 - \|Pf \circ \mathbf{M}\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2 \right| \leq C\iota_2 \|f\|_{L^2(\mathcal{M})} \sqrt{E_{\mathcal{M}}(f)} + C(\iota_1 + \iota_2) \|f\|_{L^2(\mathcal{M})}^2, \quad \forall f \in L^2(\mathcal{M}), \quad (7.20)$$

$$\left| \|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2 - \|\mathcal{I}(f \circ \mathbf{M}^{-1})\|_{L^2(\mathcal{M})}^2 \right| \leq C\varepsilon \|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} \sqrt{E_{\bar{\mathcal{X}}}(f)} + C(\iota_1 + \iota_2) \|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2, \quad \forall f \in L^2_{\vartheta_n}(\bar{\mathcal{X}}). \quad (7.21)$$

With (7.18), (7.19), (7.20), (7.21) established, we proceed similarly to before, directly adapting the remainder of the proof to the variational argument in (Calder and García Trillos, 2022, Theorem 2.4 (i)), and subsequently to (Calder and García Trillos, 2022, Theorem 2.7 (ii)), line by line. In particular, if (7.10) holds, we derive that in place of (7.11),

$$|\lambda_l - \alpha\beta \hat{\lambda}_l^{\text{aug}}| \leq C \left(\varepsilon (\sqrt{\lambda_l} + 1) + \varepsilon + \iota_1 \right) \lambda_l,$$

with probability at least $1 - n \exp(-cn\iota_1^2\iota_2^m) - 2n \exp(-c_\eta n \varepsilon^{m+4})$. The additional reduction in probability, relative to what is obtained previously, arises from the probability that both (7.16) and (7.17) hold. Finally, if (7.12) is also met, we are guaranteed with probability at least $1 - n \exp(-cn\iota_1^2\iota_2^m) - 10n \exp(-c_\eta n \varepsilon^{m+4})$, the existence of an eigenfunction f_l of Δ_{aug} such that

$$\|\hat{f}_l^{\text{aug}}(\cdot) - f_l(Q \cdot)\|_{L^2_{\vartheta_n}(\mathcal{Q}_{\bar{\mathcal{X}}})} \leq C_{\eta, l} \varepsilon.$$

The further reduction in probability comes from leveraging Proposition 14. The proof is now concluded. $\blacksquare$

7.3 Proof of Proposition 21

Proof Let $r > 0$ be sufficiently small. We define two non-local energy forms for $f \in L^2(\mathcal{M})$,

$$\begin{aligned}\mathbf{E}_{\mathcal{M};r}(f) &:= \frac{\alpha}{r^{m+2}} \int_{\mathcal{M}} \int_{\mathcal{M}} \mathbb{1}_{\{d(x,y) \leq r\}} (f(x) - f(y))^2 p_n(x) p_n(y) d\mathcal{V}(x) d\mathcal{V}(y), \\ \mathbf{E}_{\mathcal{M};r}^*(f) &:= \frac{\alpha}{r^{m+2}} \int_{\mathcal{M}} \int_{\mathcal{M}} \mathbb{1}_{\{d(x,y) \leq r\}} (f(x) - f(y))^2 q(x) q(y) d\mathcal{V}(x) d\mathcal{V}(y).\end{aligned}$$

Moreover, we temporarily assume the validity of the following lemma, whose proof is given shortly after.

Lemma 24 *Let $\varepsilon > 0$ and $\iota_1, \iota_2 > 0$ satisfy Assumption 10. Let the discretization map P , the extension map P^* be given in (7.3), (7.4), respectively. Then*

$$E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(Pf) \leq \left(1 + \frac{C\iota_2}{\varepsilon}\right) \mathbf{E}_{\mathcal{M};\varepsilon+2\iota_2}(f), \quad \forall f \in L^2(\mathcal{M}), \quad (7.22)$$

$$\mathbf{E}_{\mathcal{M};\varepsilon-2\iota_2}(P^*f) \leq \left(1 + \frac{C\iota_2}{\varepsilon}\right) E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f), \quad \forall f \in L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{\mathcal{X}}}). \quad (7.23)$$

Here, the constants $C > 0$ depend on $\mathcal{M}, \rho$.

The remainder of the proof now follows closely the approach in the proof of (Calder and García Trillos, 2022, Proposition 4.1). For instance, given Lemma 24, we can apply Propositions 7, 20 to obtain

$$E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(Pf) \leq \left(1 + \frac{C\iota_2}{\varepsilon}\right) \left(1 + C(\iota_1 + \iota_2)\right) \left(1 + C_\eta \varepsilon_n^{2/3}\right) \mathbf{E}_{\mathcal{M};\varepsilon+2\iota_2}^*(f), \quad \forall f \in L^2(\mathcal{M}), \quad (7.24)$$

$$\mathbf{E}_{\mathcal{M};\varepsilon-2\iota_2}^*(P^*f) \leq \left(1 + \frac{C\iota_2}{\varepsilon}\right) \left(1 + C(\iota_1 + \iota_2)\right) \left(1 + C_\eta \varepsilon_n^{2/3}\right) E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f), \quad \forall f \in L^2_{\mathfrak{V}_n}(\mathcal{Q}_{\bar{\mathcal{X}}}), \quad (7.25)$$

where, both (7.24), (7.25) hold with probability at least $1 - n \exp(-cn\iota_1^2\iota_2^m)$. Then, applying the arguments from (García Trillos et al., 2020, Lemmas 5, 9), we deduce, respectively for $f \in L^2(\mathcal{M})$ and $r > 0$ sufficiently small,

$$\mathbf{E}_{\mathcal{M};r}^*(f) \leq \alpha\beta \left(1 + Cr + Cr^2\right) E_{\mathcal{M}}(f), \quad (7.26)$$

$$\alpha\beta E_{\mathcal{M}}(f)(\Lambda_r f) \leq \left(1 + Cr + Cr^2\right) \left(1 + Cr^2\right) \mathbf{E}_{\mathcal{M};r}^*(f). \quad (7.27)$$

With minor simplifications, combining (7.24), (7.26) for $r = \varepsilon + 2\iota_2$, gives us (7.6), while combining (7.25), (7.27) for $r = \varepsilon - 2\iota_2$, gives us (7.7). $\blacksquare$

It remains for us to prove Lemma 24.

Proof [Proof of Lemma 24] We recall from Lemma 19 and its proof, that if $d(x, y) \leq r$, for $r > 0$ sufficiently small, then

$$\frac{\varsigma(x, y)}{\varepsilon_w^m} = \alpha\rho(x)^{1/2}\rho(y)^{1/2} + \mathcal{O}(r^{1/2}). \quad (7.28)$$

Let $\varepsilon^* := \varepsilon(1 + 2\varepsilon_w)$, $\varepsilon_* := \varepsilon(1 - 2\varepsilon_w)$, and $\epsilon^* := \varepsilon^* + \frac{8}{R^2}(\varepsilon^*)^3$; all are small, positive quantities comparable to ε . We start with (7.22). The proof mirrors the approach outlined in (García Trillos et al., 2020, Lemma 13), albeit with a few modifications due to the adaptive definitions of discretization and extension maps, as well as the fluctuation of edge lengths in the graph $G_{\mathcal{Q}_{\bar{\mathcal{X}}}}$, as pointed out earlier. If $f \in L^2(\mathcal{M})$, then by definition $Pf \in L^2_{\mathcal{V}_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$. Thus from (6.21) and definition (7.1),

$$E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(Pf) \leq \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i,j=1}^n \varsigma(Q_{\bar{x}_i}, Q_{\bar{x}_j}) \mathbb{1}_{\{0 < \|Q_{\bar{x}_i} - Q_{\bar{x}_j}\| \leq \varepsilon^*\}} (Pf(Q_{\bar{x}_i}) - Pf(Q_{\bar{x}_j}))^2,$$

which, using (2.7), (7.28) with $r = \varepsilon^*$, simplifies to

$$E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(Pf) \leq \bar{E}_1(Pf) + C(\varepsilon^*)^{1/2} \bar{E}_2(Pf), \quad (7.29)$$

where

$$\begin{aligned} \bar{E}_1(Pf) &:= \frac{\alpha}{n^2 \varepsilon^{m+2}} \sum_{i,j=1}^n \rho(Q_{\bar{x}_i})^{1/2} \rho(Q_{\bar{x}_j})^{1/2} \mathbb{1}_{\{0 < \|Q_{\bar{x}_i} - Q_{\bar{x}_j}\| \leq \varepsilon^*\}} (Pf(Q_{\bar{x}_i}) - Pf(Q_{\bar{x}_j}))^2 \\ \bar{E}_2(Pf) &:= \frac{1}{n^2 \varepsilon^{m+2}} \sum_{i,j=1}^n \mathbb{1}_{\{0 < \|Q_{\bar{x}_i} - Q_{\bar{x}_j}\| \leq \varepsilon^*\}} (Pf(Q_{\bar{x}_i}) - Pf(Q_{\bar{x}_j}))^2. \end{aligned}$$

Considering $\bar{E}_1(Pf)$, we directly obtain from the definition (7.3) of the discretization map P that

$$\begin{aligned} \bar{E}_1(Pf) &\leq \frac{\alpha}{n \varepsilon^{m+2}} \sum_{j=1}^n \int_{U_j} \rho(Q_{\bar{x}_j})^{1/2} \mathbb{1}_{\{d(x, Q_{\bar{x}_j}) \leq \varepsilon^* + 2\iota_2\}} (f(x) - Pf(Q_{\bar{x}_j}))^2 p_n(x) d\mathcal{V}(x) \\ &= \frac{\alpha}{\varepsilon^{m+2}} \int_{\mathcal{M}} \int_{\mathcal{M}} \mathbb{1}_{\{d(x,y) \leq \varepsilon^* + 2\iota_2\}} (f(x) - f(y))^2 p_n(x) p_n(y) d\mathcal{V}(x) d\mathcal{V}(y). \end{aligned}$$

Then from this point on, we proceed identically to the proof of (García Trillos et al., 2020, Lemma 13) (and applying the technical (García Trillos et al., 2020, Lemma 5) to exchange $\varepsilon^* + 2\iota_2 \approx \varepsilon + 2\iota_2$) to get

$$\bar{E}_1(Pf) \leq \left(1 + \frac{C\iota_2}{\varepsilon}\right) \mathbf{E}_{\mathcal{M}; \varepsilon + 2\iota_2}^*(f). \quad (7.30)$$

Regarding $\bar{E}_2(Pf)$, we observe by (2.14) that

$$\begin{aligned} &\bar{E}_2(Pf) \\ &\leq \frac{C}{n^2 \varepsilon^{m+2}} \sum_{i,j=1}^n \rho(Q_{\bar{x}_i})^{1/2} \rho(Q_{\bar{x}_j})^{1/2} \mathbb{1}_{\{0 < \|Q_{\bar{x}_i} - Q_{\bar{x}_j}\| \leq \varepsilon^*\}} (Pf(Q_{\bar{x}_i}) - Pf(Q_{\bar{x}_j}))^2 = C \bar{E}_1(Pf); \end{aligned}$$

thus from (7.30)

$$(\varepsilon^*)^{1/2} \bar{E}_2(Pf) \leq C(\varepsilon^*)^{1/2} \left(1 + \frac{C\iota_2}{\varepsilon}\right) \mathbf{E}_{\mathcal{M}; \varepsilon + 2\iota_2}^*(f). \quad (7.31)$$

Since $C\varepsilon^{3/2} \leq \iota_2$ implies that $C(\varepsilon^*)^{1/2} \leq \iota_2/\varepsilon$, we combine (7.29), (7.30), (7.31) to conclude (7.22).

For (7.23), note that if $f \in L^2_{\vartheta_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$ then $P^*f \in L^2(\mathcal{M})$. We derive from (2.7), (6.22), (7.28) with $r = \varepsilon_*$, and definition (7.1) that

$$\begin{aligned} E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f) &\geq \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{Q_{\bar{x}_i}, Q_{\bar{x}_j} \in \mathcal{Q}_{\bar{\mathcal{X}}}} \varsigma(Q_{\bar{x}_i}, Q_{\bar{x}_j}) \mathbb{1}_{\{0 < \|Q_{\bar{x}_i} - Q_{\bar{x}_j}\| \leq \varepsilon_*\}} (f(Q_{\bar{x}_i}) - f(Q_{\bar{x}_j}))^2 \\ &\geq \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{Q_{\bar{x}_i}, Q_{\bar{x}_j} \in \mathcal{Q}_{\bar{\mathcal{X}}}} (\alpha \rho(Q_{\bar{x}_i})^{1/2} \rho(Q_{\bar{x}_j})^{1/2} - C\varepsilon_*^{1/2}) \mathbb{1}_{\{0 < \|Q_{\bar{x}_i} - Q_{\bar{x}_j}\| \leq \varepsilon_*\}} (f(Q_{\bar{x}_i}) - f(Q_{\bar{x}_j}))^2. \end{aligned}$$

The proof closely parallels that of (García Trillos et al., 2020, Lemma 14); thus adopting a similar approach to the one used earlier, we obtain

$$E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(f) \geq \left(1 - \frac{C\iota_2}{\varepsilon}\right) \mathbf{E}_{\mathcal{M}; \varepsilon - 2\iota_2}^*(P^*f),$$

which is (7.23). ■

7.4 Proof of Proposition 22

Proof The proof is nearly word-for-word identical to that of (Calder and García Trillos, 2022, Proposition 4.2), with modifications similar to those encountered in the proof of Lemma 24. One modification involves incorporating the adaptive definitions (7.3), (7.4), which yields

$$\|P^*Pf\|_{L^2_{v_n}(\mathcal{M})}^2 = \|(Pf)\rho^{1/2}\|_{L^2_{\vartheta_n}(\mathcal{Q}_{\bar{\mathcal{X}}})}^2,$$

with the extra density above being the distinguishing factor from the corresponding quantities in (Calder and García Trillos, 2022, Proposition 4.2). Another modification appears in the application of (García Trillos et al., 2020, Lemma 8) near the end. Specifically, ε needs to be replaced by $\varepsilon(1 + 2\varepsilon_w)$, and this small difference is again accounted for by the technical (García Trillos et al., 2020, Lemma 5). After these modifications, every other step in (Calder and García Trillos, 2022, Proposition 4.2) remains unchanged. Lastly, we mention that the probability $1 - n \exp(-cn\iota_1^2\iota_2^m)$ aligns with the probability of the existence of v_n in Proposition 20. Further details are omitted. ■

7.5 Proof of Proposition 23

Proof Since (7.14), (7.15) follow directly from the definition of the matching map M , and (7.17) is similar but less complex than (7.16), we will prove only (7.16) and briefly address (7.17) afterward. Let $f \in L^2_{\vartheta_n}(\bar{\mathcal{X}})$, and let $\bar{f} := f \circ M^{-1} \in L^2_{\vartheta_n}(\mathcal{Q}_{\bar{\mathcal{X}}})$. Then

$$E_{\bar{\mathcal{X}}}(f) = 2\langle \mathcal{L}_{\text{aug}}f, f \rangle_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} = 2\langle \mathcal{L}_{\text{aug}}f - (\mathcal{L}_{\text{proj}}\bar{f}) \circ M, f \rangle_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} + 2\langle (\mathcal{L}_{\text{proj}}\bar{f}) \circ M, f \rangle_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}. \quad (7.32)$$

Observe, since $f(\bar{x}) = \bar{f}(Q_{\bar{x}})$, we have

$$\begin{aligned} \langle (\mathcal{L}_{\text{proj}} \bar{f}) \circ M, f \rangle_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} &= \frac{1}{n} \sum_{i=1}^n ((\mathcal{L}_{\text{proj}} \bar{f}) \circ M)(\bar{x}_i) f(\bar{x}_i) \\ &= \frac{1}{n} \sum_{i=1}^n (\mathcal{L}_{\text{proj}} \bar{f})(Q_{\bar{x}_i}) \bar{f}(Q_{\bar{x}_i}) = \frac{1}{2} E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(\bar{f}), \end{aligned} \quad (7.33)$$

and moreover,

$$\begin{aligned} &\langle \mathcal{L}_{\text{aug}} f - (\mathcal{L}_{\text{proj}} \bar{f}) \circ M, f \rangle_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} \\ &= \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i=1}^n \sum_{j=1}^n (\omega(\bar{x}_i, \bar{x}_j) - \varsigma(Q_{\bar{x}_i}, Q_{\bar{x}_j})) \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} (\bar{f}(Q_{\bar{x}_i}) - \bar{f}(Q_{\bar{x}_j})) \bar{f}(Q_{\bar{x}_i}). \end{aligned} \quad (7.34)$$

Now, recalling the steps in the proof of Proposition 14, particularly (6.11), (6.14), we obtain

$$-C\varepsilon^2 \varsigma(Q_{\bar{x}_i}, Q_{\bar{x}_j}) \leq \omega(\bar{x}_i, \bar{x}_j) - \varsigma(Q_{\bar{x}_i}, Q_{\bar{x}_j}) \leq C\varepsilon^2 \varsigma(Q_{\bar{x}_i}, Q_{\bar{x}_j}) + \exp\left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}}\right). \quad (7.35)$$

By applying Lemma 27 (with $\delta = \varepsilon^2$), combined with a union bound, and adapting the argument leading to (6.24), we conclude that

$$\begin{aligned} \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} (\bar{f}(Q_{\bar{x}_i}))^2 &= \frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i=1}^n (\bar{f}(Q_{\bar{x}_i}))^2 \sum_{Q_{\bar{x}_j}: 0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon} 1 \\ &\leq \frac{C_\eta(1 + \varepsilon^2) \|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2}{\varepsilon_w^m \varepsilon^2} \end{aligned} \quad (7.36)$$

holds with probability at least $1 - 2n \exp(-c_\eta n \varepsilon^{m+4})$. Similarly,

$$\begin{aligned} &\frac{1}{n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} |\bar{f}(Q_{\bar{x}_i})| |\bar{f}(Q_{\bar{x}_j})| \\ &\leq \frac{1}{2n^2 \varepsilon_w^m \varepsilon^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \mathbb{1}_{\{0 < \|\bar{x}_i - \bar{x}_j\| \leq \varepsilon\}} (|\bar{f}(Q_{\bar{x}_i})|^2 + |\bar{f}(Q_{\bar{x}_j})|^2) \leq \frac{C_\eta(1 + \varepsilon^2) \|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2}{\varepsilon_w^m \varepsilon^2} \end{aligned} \quad (7.37)$$

holds with the same probability. We now substitute the second inequality in (7.35) into (7.34), integrating (7.36), (7.37), as well as invoking (3.17), to derive

$$\begin{aligned} \langle \mathcal{L}_{\text{aug}} f - (\mathcal{L}_{\text{proj}} \bar{f}) \circ M, f \rangle_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} &\leq C\varepsilon^2 E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(\bar{f}) + C_\eta \exp\left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}}\right) \frac{\|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2}{\varepsilon_w^m \varepsilon^2} \\ &\leq C\varepsilon^2 E_{\mathcal{Q}_{\bar{\mathcal{X}}}}(\bar{f}) + C_\eta \varepsilon \|f\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})}^2. \end{aligned} \quad (7.38)$$

Then, combining (7.38) with (7.32), (7.33) yields (7.16). It remains to see that the proof of (7.17) follows in the same manner, but utilizes the first inequality in (7.35). $\blacksquare$

8. Proof of Theorem 13

Proof Let $l = 1, \dots, k$. Let $\hat{f}_l^{\text{aug}}$ be the j th normalized eigenvector (with respect to $L^2_{\vartheta_n}(\bar{\mathcal{X}})$) of $\mathcal{L}_{\text{aug}}$, and, by Theorem 12, let f_l be the j th normalized eigenfunction (with respect to $L^2(\mathcal{M})$) of Δ_{aug} that satisfies (4.6). We have

$$\frac{1}{\sqrt{n}} \|\hat{f}_l^{\text{aug}}(\cdot) - f_l(Q \cdot)\|_{\ell^\infty(\bar{\mathcal{X}})} \leq \|\hat{f}_l^{\text{aug}}(\cdot) - f_l(Q \cdot)\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} \leq C_{\eta, l} \varepsilon, \quad (8.1)$$

and

$$\frac{1}{\sqrt{n}} \|\hat{f}_l^{\text{aug}}(\cdot) - f_l(Q \cdot)\| \leq \|\hat{f}_l^{\text{aug}}(\cdot) - f_l(Q \cdot)\|_{L^2_{\vartheta_n}(\bar{\mathcal{X}})} \leq C_{\eta, l} \varepsilon, \quad (8.2)$$

where $\|\cdot\|$ in (8.2) denotes the Euclidean vector norm. Based on the elliptic theory of PDEs (Gilbarg and Trudinger, 1977, Section 6.3) and Proposition 7, f_l is a Lipschitz function. Then as a consequence of (Chen et al., 2022, Theorem 3.1), there exists a ReLU NN $(f_l)_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$ of the form (3.22), where $\theta \in \Theta$ given in (3.26), such that for every $\tilde{\delta} \in (0, 1)$

$$\frac{1}{\sqrt{n}} \|f_l - (f_l)_\theta\|_{L^\infty(\mathcal{M})} \leq \|f_l - (f_l)_\theta\|_{L^\infty(\mathcal{M})} \leq \tilde{\delta}, \quad (8.3)$$

and moreover

$$\frac{1}{\sqrt{n}} \|f_l(Q \cdot) - (f_l)_\theta(Q \cdot)\| \leq \|f_l - (f_l)_\theta\|_{L^\infty(\mathcal{M})} \leq \tilde{\delta}. \quad (8.4)$$

Thus, we obtain, from (8.1), (8.3),

$$\frac{1}{\sqrt{n}} \|\hat{f}_l^{\text{aug}}(\cdot) - (f_l)_\theta(Q \cdot)\|_{\ell^\infty(\bar{\mathcal{X}})} \leq C_{\eta, l} \varepsilon + \tilde{\delta}, \quad (8.5)$$

and from (8.2), (8.4),

$$\frac{1}{\sqrt{n}} \|\hat{f}_l^{\text{aug}}(\cdot) - (f_l)_\theta(Q \cdot)\| \leq C_{\eta, l} \varepsilon + \tilde{\delta}, \quad (8.6)$$

Further, by Eckart-Young-Minsky theorem, a solution to (3.20) takes the form

$$Y^* = \frac{1}{\sqrt{n}} \left[\sqrt{\hat{\lambda}_l^{\text{aug}}} \hat{f}_l^{\text{aug}}(\bar{x}_i) \right]_{i=1, \dots, n; l=1, \dots, k} \in \mathbb{R}^{n \times k},$$

assuming the rotation matrix is the $k \times k$ identity matrix. Stack the ReLU NNs $\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta$, $l = 1, \dots, k$, to obtain a ReLU NN $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^k$, and let $Y_\theta = \frac{1}{\sqrt{n}}(f_\theta(\bar{x}_1), \dots, f_\theta(\bar{x}_n))^\top \in \mathbb{R}^{n \times k}$. We conclude respectively from (8.5), (8.6),

$$\|Y^* - Y_\theta\|_\infty \leq C_{\eta, k}(\varepsilon + \tilde{\delta}) \quad \text{and} \quad \|Y^* - Y_\theta\| \leq C'_{\eta, k}(\varepsilon + \tilde{\delta}), \quad (8.7)$$

as desired. Here, in (8.7), both $C_{\eta, k}$ and $C'_{\eta, k}$ involve a factor of $\max_{l=1, \dots, k} \sqrt{\hat{\lambda}_l^{\text{aug}}}$, with $C'_{\eta, k}$ further including a factor of $\sqrt{k}$. With $L(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta)$ denoting the network depth of $\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta$, $p(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta)$ its width, and $m(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta)$, $M(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta)$ the complexity

parameters from (3.24), (3.25) respectively, we obtain in addition from (Chen et al., 2022, Theorem 3.1),

$$\mathsf{L}\left(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta\right) \leq C\left(\log\left(\frac{1}{\tilde{\delta}}\right) + \log d\right) \quad \text{and} \quad \mathsf{p}\left(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta\right) \leq C(\tilde{\delta}^{-m} + d), \quad (8.8)$$

and

$$\mathsf{m}\left(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta\right) \leq C \max_{l=1,\dots,k} \sqrt{\hat{\lambda}_l^{\text{aug}}} \quad \text{and} \quad \mathsf{M}\left(\sqrt{\hat{\lambda}_l^{\text{aug}}}(f_l)_\theta\right) \leq \left((\tilde{\delta}^{-m} + d) \log\left(\frac{1}{\tilde{\delta}}\right) + d \log d\right), \quad (8.9)$$

where all the constants C depend only on $\mathcal{M}$. Therefore, it follows from (8.8), (8.9) respectively that

$$\mathsf{L}\left(\frac{1}{\sqrt{n}}f_\theta\right) \leq C\left(\log\left(\frac{1}{\tilde{\delta}}\right) + \log d\right) \quad \text{and} \quad \mathsf{p}\left(\frac{1}{\sqrt{n}}f_\theta\right) \leq Ck(\tilde{\delta}^{-m} + d),$$

and

$$\mathsf{m}\left(\frac{1}{\sqrt{n}}f_\theta\right) \leq \frac{C \max_{l=1,\dots,k} \sqrt{\hat{\lambda}_l^{\text{aug}}}}{\sqrt{n}} \leq C \quad \text{and} \quad \mathsf{M}\left(\frac{1}{\sqrt{n}}f_\theta\right) \leq k\left((\tilde{\delta}^{-m} + d) \log\left(\frac{1}{\tilde{\delta}}\right) + d \log d\right),$$

when n is sufficiently large. Finally, we perform a union bound over the probabilities of (4.6) occurring for all $l = 1, \dots, k$ to complete the proof. $\blacksquare$

Acknowledgments and Disclosure of Funding

CL gratefully acknowledges the support from the IFDS at UW-Madison and the NSF through the TRIPODS grant 2023239. AMN is supported by the Austrian Science Fund (FWF) Project P-37010.

Appendix A. Relevant Concentration Results

We present the following two known concentration lemmas.

Lemma 25 (Calder and García Trillos, 2022, Lemma 3.1 and its proof) (see also (Calder, 2019, Lemma 1)) *Let $\mathcal{M}$ satisfy the assumptions made in Section 2.2. Let $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ be bounded and Borel measurable with compact support in a bounded open set $\Omega \subset \mathbb{R}^d$ such that $\Omega \cap \mathcal{M} \neq \emptyset$. Let $x_1, \dots, x_n$ be a sample of i.i.d points on $\mathcal{M}$ following a distribution on $\mathcal{M}$ with a positive continuous density p , such that $p(x) \leq p_{\max} < \infty$ for all $x \in \mathcal{M}$. Let*

$$\Psi := \frac{1}{n} \sum_{i=1}^n \psi(x_i),$$

and let $a := \int_{\Omega \cap \mathcal{M}} \psi(z)p(z)d\mathcal{V}(z)$. Then for every $0 < \delta \leq 1$,

$$\mathbb{P}(|\Psi - a| > Cp_{\max}\|\psi\|_{L^\infty(\Omega)}\mathcal{V}(\Omega \cap \mathcal{M})\delta) \leq 2 \exp(-cp_{\max}n\mathcal{V}(\Omega \cap \mathcal{M})\delta^2),$$

where $C, c > 0$ depend only on $\mathcal{M}$.

When $\Omega = B^d(x, \varepsilon)$ for $x \in \mathcal{M}$, Lemma 28 below can be seen as a special case of Lemma 27.

Lemma 26 (*Calder and García Trillos, 2022, Lemma 3.1 and its proof*) *Let $\mathcal{M}$ satisfy the assumptions made in Section 2.2. Let $\psi : \mathcal{M} \rightarrow \mathbb{R}$ be bounded and Borel measurable. Let $x_1, \dots, x_n$ be a sample of i.i.d points on $\mathcal{M}$ following a distribution on $\mathcal{M}$ with a positive continuous density p , such that $p(x) \leq p_{\max} < \infty$ for all $x \in \mathcal{M}$. Let $\varepsilon > 0$. For $x \in \mathcal{M}$, we define $\Psi = \Psi(x)$ to be*

$$\Psi := \frac{1}{n} \sum_{i: \|x_i - x\| \leq \varepsilon} \psi(x_i).$$

Let

$$a := \int_{B^d(x, \varepsilon) \cap \mathcal{M}} \psi(z) p(z) d\mathcal{V}(z) \quad \text{and} \quad b := \int_{\mathcal{B}(x, \varepsilon)} \psi(z) p(z) d\mathcal{V}(z).$$

Then we have

$$|a - b| \leq Cp_{\max} \|\psi\|_{L^\infty(\mathcal{B}(x, 2\varepsilon))} \varepsilon^{m+2}.$$

Additionally, for any $\varepsilon^2 \leq \delta \leq 1$,

$$\mathbb{P}(|\Psi - b| \geq Cp_{\max} \|\psi\|_{L^\infty(\mathcal{B}(x, 2\varepsilon))} \delta \varepsilon^m) \leq 2 \exp(-cp_{\max} \delta^2 n \varepsilon^m),$$

where $C, c > 0$ depend only on $\mathcal{M}$.

Both Lemmas 25, 26 apply to i.i.d. random variables. However, in the context of our analysis, we require concentration results for the collection $\mathcal{Q}_{\bar{\mathcal{X}}}$ of projection points $Q_{\bar{x}}$. While these points can be viewed as samples from the measure ν with density q (3.1), they are not independent. Nevertheless, under Assumption 1, a similar partial independence condition holds for $\mathcal{Q}_{\bar{\mathcal{X}}}$, in that for every $Q_{\bar{x}}$, there exist at most $C_0 \in \mathbb{N}$ other $Q_{\bar{x}'}$ dependent on $Q_{\bar{x}}$. We present variants of both lemmas to address this scenario, as follows.

Lemma 27 *Let $\mathcal{M}$ satisfy the assumptions made in Section 2.2. Let $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ be bounded and Borel measurable with compact support in a bounded open set $\Omega \subset \mathbb{R}^d$ such that $\Omega \cap \mathcal{M} \neq \emptyset$. Let $x_1, \dots, x_n$ be a sample of i.i.d points on $\mathcal{M}$ following a distribution on $\mathcal{M}$ with a positive continuous density p , such that $p(x) \leq p_{\max} < \infty$ for all $x \in \mathcal{M}$. Suppose further that for each point x_i there are at most C_0 other points x_j dependent on x_i . Let*

$$\Psi := \frac{1}{n} \sum_{i=1}^n \psi(x_i),$$

and let $a := \int_{\Omega \cap \mathcal{M}} \psi(z) p(z) d\mathcal{V}(z)$. Then for every $0 < \delta \leq 1$,

$$\mathbb{P}(|\Psi - a| > Cp_{\max} \|\psi\|_{L^\infty(\Omega)} \mathcal{V}(\Omega \cap \mathcal{M}) \delta) \leq 2 \exp(-cp_{\max} n \mathcal{V}(\Omega \cap \mathcal{M}) \delta^2), \quad (\text{A.1})$$

where $C, c > 0$ depend only on $\mathcal{M}$.

Lemma 28 *Let $\mathcal{M}$ satisfy the assumptions made in Section 2.2. Let $\psi : \mathcal{M} \rightarrow \mathbb{R}$ be bounded and Borel measurable. Let $x_1, \dots, x_n$ be a sample of points on $\mathcal{M}$ following an identical distribution on $\mathcal{M}$ with a positive continuous density p , such that $p(x) \leq p_{\max} < \infty$*

for all $x \in \mathcal{M}$. Suppose further that for each point x_i there are at most C_0 other points x_j dependent on x_i . Let $\varepsilon > 0$. For $x \in \mathcal{M}$, we define $\Psi = \Psi(x)$ to be

$$\Psi := \frac{1}{n} \sum_{i: \|x_i - x\| \leq \varepsilon} \psi(x_i).$$

Let

$$a := \int_{B^d(x, \varepsilon) \cap \mathcal{M}} \psi(z) p(z) d\mathcal{V}(z) \quad \text{and} \quad b := \int_{\mathcal{B}(x, \varepsilon)} \psi(z) p(z) d\mathcal{V}(z).$$

Then we have

$$|a - b| \leq Cp_{\max} \|\psi\|_{L^\infty(\mathcal{B}(x, 2\varepsilon))} \varepsilon^{m+2}.$$

Additionally, for any $\varepsilon^2 \leq \delta \leq 1$,

$$\mathbb{P}(|\Psi - b| \geq Cp_{\max} \|\psi\|_{L^\infty(\mathcal{B}(x, 2\varepsilon))} \delta \varepsilon^m) \leq 2 \exp(-cp_{\max} \delta^2 n \varepsilon^m), \quad (\text{A.2})$$

where $C, c > 0$ depend only on $\mathcal{M}$.

The constant c in both probabilistic estimates (A.1), (A.2) implicitly includes a factor of C_0^{-1} , which we assume to be predetermined. The proofs of Lemma 27 and Lemma 28 can be easily adapted from those in (Calder and García Trillos, 2022, Lemma 3.1) and (Janson, 2004, Theorem 2.1). We skip further details.

Appendix B. Relevant Geometric Construction for Propositions 7, 15

Let $r > 0$ be sufficiently small. Then recall from Section 2.2 that the logarithm map

$$\text{Log}_x : \mathcal{M} \supset \mathcal{B}(x, r) \rightarrow B^m(0, r) \subset T_x \mathcal{M},$$

is a diffeomorphism. Let $\mathbf{l}_m : \mathbb{R}^m \rightarrow \mathbb{R}^d$, where $\mathbf{l}_m(y^1, \dots, y^m) = (y^1, \dots, y^m, 0, \dots, 0)$, be a *canonical immersion* map (see (Guillemin and Pollack, 2010, Chapter 3)). Let A_x denote a composition of a translation by x and an appropriate rotation R_x , such that A_x maps

$$\mathbb{R}^m \times \{0\} := \{(y^1, \dots, y^m, 0, \dots, 0) : y^i \in \mathbb{R}\}$$

onto $\tilde{T}_x \mathcal{M}$. See Figure 4 below for an illustration. Then $A_x \circ \mathbf{l}_m$ maps $T_x \mathcal{M} \cong \mathbb{R}^m$ bijectively and isometrically to $\tilde{T}_x \mathcal{M} \subset \mathbb{R}^d$. Expressively,

$$A_x \circ \mathbf{l}_m(T_x \mathcal{M}) = \tilde{T}_x \mathcal{M}. \quad (\text{B.1})$$


By letting $L_x := A_x \circ \mathbf{l}_m \circ \text{Log}_x$, we obtain for $x' \in \mathcal{B}(x, r)$ a unique image $L_x(x') \in \tilde{T}_x \mathcal{M}$, such that

$$\|L_x(x') - x\| = \|\text{Log}_x(x')\| = d(x, x'). \quad (\text{B.2})$$

Next, let $s > 0$. We consider the neighborhood $\mathcal{N}_s(\mathcal{B}(x, r))$ of $\mathcal{B}(x, r)$, defined to be

$$\mathcal{N}_s(\mathcal{B}(x, r)) := \{x' + h : x' \in \mathcal{B}(x, r) \text{ and } h \in \tilde{N}_{x'} \mathcal{M} \cap \mathcal{N}_s(\mathcal{M}) - x'\}.$$

Figure 4: A two-dimensional sphere and one of its tangent planes under rotation and translation



Note, $\tilde{N}_{x'}\mathcal{M} - x' \cong N_{x'}\mathcal{M} \cong \mathbb{R}^{d-m}$. Similarly, as before, such identification can be made explicit, for the rotation $R_{x'}$ also maps

$$\{0\} \times \mathbb{R}^{d-m} := \{(0, \dots, 0, y^1, \dots, y^{d-m}) : y^i \in \mathbb{R}\}$$

onto $\tilde{N}_{x'}\mathcal{M} - x'$. Let $J_{d-m} : \mathbb{R}^{d-m} \rightarrow \mathbb{R}^d$ be another immersion map, where

$$J_{d-m}(y^1, \dots, y^{d-m}) = (0, \dots, 0, y^1, \dots, y^{d-m}).$$

Then, through the bijective isometric map $(R_{x'} \circ J_{d-m})^{-1}$, $h \in \tilde{N}_{x'}\mathcal{M} \cap \mathcal{N}_s(\mathcal{M}) - x'$ is uniquely identified with $k = (R_{x'} \circ J_{d-m})^{-1}(h) \in B^{d-m}(0, s)$. Now using the diffeomorphic map

$$\text{Exp}_x : T_x\mathcal{M} \supset B^m(0, r) \rightarrow \mathcal{B}(x, r) \subset \mathcal{M},$$

we can construct another diffeomorphic map:

$$\begin{aligned} \overline{\text{Exp}}_x : B^m(0, r) \times B^{d-m}(0, s) &\rightarrow \mathcal{N}_s(\mathcal{B}(x, r)) \\ (v, k) &\mapsto x' + h, \end{aligned}$$

where $x' = \text{Exp}_x(v)$ and, as noted, $h = R_{x'} \circ J_{d-m}(k)$. It is readily checked that the volume form of $\overline{\text{Exp}}_x$ at (v, k) is $J_x(v)dv \times dk$.

Appendix C. Proof of Proposition 15

Proof We begin by establishing the necessary groundwork. Let $\tilde{T}_{Q_{\bar{x}}}\mathcal{M}$ and $\tilde{T}_{Q_{\bar{x}'}}\mathcal{M}$ be the tangent planes of $\mathcal{M}$ at $Q_{\bar{x}}$ and at $Q_{\bar{x}'}$, respectively. It follows that

$$Q_{\bar{x}} - \bar{x} \perp \tilde{T}_{Q_{\bar{x}}}\mathcal{M} \quad \text{and} \quad Q_{\bar{x}'} - \bar{x}' \perp \tilde{T}_{Q_{\bar{x}'}}\mathcal{M}, \quad (\text{C.1})$$




Figure 5: $Q_{\bar{x}}$ and $Q_{\bar{x}'}$ are the closest points to $\bar{x}$ and $\bar{x}'$ on $\mathcal{M}$, respectively. $\tilde{T}_{Q_{\bar{x}}}\mathcal{M}$ and $\tilde{T}_{Q_{\bar{x}'}}\mathcal{M}$ are the tangent planes of $\mathcal{M}$ on $Q_{\bar{x}}$ and $Q_{\bar{x}'}$, respectively.

where $\perp$ denotes the usual orthogonal notion in $\mathbb{R}^d$; see Figure 5. From (2.7), (2.15), we obtain¹⁴

$$\begin{aligned} d(x, Q_{\bar{x}}) &\leq \|x - Q_{\bar{x}}\| + \frac{8}{R^2}\|x - Q_{\bar{x}}\|^3 \leq C\varepsilon_n^{2/3}, \\ d(x, Q_{\bar{x}'}) &\leq \|x - Q_{\bar{x}'}\| + \frac{8}{R^2}\|x - Q_{\bar{x}'}\|^3 \leq C\varepsilon_n^{2/3}. \end{aligned}$$

Hence, by applying (B.2) to the locations $Q_{\bar{x}}$, $Q_{\bar{x}'}$, we acquire for x , respective unique images $L_{Q_{\bar{x}}}(x)$, $L_{Q_{\bar{x}'}}(x)$, such that

$$\begin{aligned} \|L_{Q_{\bar{x}}}(x) - Q_{\bar{x}}\| &= \|\text{Log}_{Q_{\bar{x}}}(x)\| = d(x, Q_{\bar{x}}) \leq C\varepsilon_n^{2/3}, \text{ and} \\ \|L_{Q_{\bar{x}'}}(x) - Q_{\bar{x}'}\| &= \|\text{Log}_{Q_{\bar{x}'}}(x)\| = d(x, Q_{\bar{x}'})) \leq C\varepsilon_n^{2/3}. \end{aligned} \quad (\text{C.2})$$

We now present the following two technical lemmas that will be indispensable in our analysis. Their proofs are provided in Appendices C.1, C.2, respectively in the order they appear.

Lemma 29 *Let x , $Q_{\bar{x}}$, $Q_{\bar{x}'}$, $L_{Q_{\bar{x}}}(x)$, $L_{Q_{\bar{x}'}}(x)$ be as above. Then it holds that*

$$\|x - L_{Q_{\bar{x}}}(x)\| = \mathcal{O}(\|L_{Q_{\bar{x}}}(x) - Q_{\bar{x}}\|^2) \quad \text{and} \quad \|x - L_{Q_{\bar{x}'}}(x)\| = \mathcal{O}(\|L_{Q_{\bar{x}'}}(x) - Q_{\bar{x}'}\|^2) \quad (\text{C.3})$$

where the majorant constants depend only on $\mathcal{M}$.

Lemma 30 *Let x , $Q_{\bar{x}}$, $Q_{\bar{x}'}$, $L_{Q_{\bar{x}}}(x)$, $L_{Q_{\bar{x}'}}(x)$ be as above. Then under Assumption 5, the following hold,*

$$\begin{aligned} \|x - \bar{x}\|^2 &\leq \|\bar{x} - Q_{\bar{x}}\|^2 + (1 + C\varepsilon_n^{2/3})\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2, \\ \|x - \bar{x}'\|^2 &\leq \|\bar{x}' - Q_{\bar{x}'}\|^2 + (1 + C\varepsilon_n^{2/3})\|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2, \end{aligned} \quad (\text{C.4})$$

14. Whether these inequalities are presented as strict or non-strict will not affect the subsequent analysis.

and moreover,

$$\begin{aligned} \|x - \bar{x}\|^2 &\geq \|\bar{x} - Q_{\bar{x}}\|^2 + (1 - C'\varepsilon_n^{2/3})\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2, \\ \|x - \bar{x}'\|^2 &\geq \|\bar{x}' - Q_{\bar{x}'}\|^2 + (1 - C'\varepsilon_n^{2/3})\|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2, \end{aligned} \quad (\text{C.5})$$

where $C, C' > 0$ dependent on $\mathcal{M}$.

Recalling the random variable

$$X = \mathbb{1}_{\{\|x - Q_{\bar{x}}\| < \varepsilon_n^{2/3}\} \wedge \{\|x - Q_{\bar{x}'}\| < \varepsilon_n^{2/3}\}}$$

from (6.1), we define another, related variable

$$Y := 1 - X = \mathbb{1}_{\{\|x - Q_{\bar{x}}\| \geq \varepsilon_n^{2/3}\} \vee \{\|x - Q_{\bar{x}'}\| \geq \varepsilon_n^{2/3}\}}. \quad (\text{C.6})$$

By expanding upon (2.3), we attain

$$\begin{aligned} \omega(\bar{x}, \bar{x}') &= \mathbb{E}_{x \sim \mu} \left[\exp \left(\frac{-\|x - \bar{x}\|^2 - \|x - \bar{x}'\|^2}{2\varepsilon_w^2} \right) \right] \\ &= \mathbb{E}_{x \sim \mu} \left[(X + Y) \exp \left(\frac{-\|x - \bar{x}\|^2 - \|x - \bar{x}'\|^2}{2\varepsilon_w^2} \right) \right]. \end{aligned} \quad (\text{C.7})$$

We will first prove the lower bound in (6.5). By virtue of Lemma 30, we have

$$\begin{aligned} &(X + Y) \exp \left(\frac{-\|x - \bar{x}\|^2 - \|x - \bar{x}'\|^2}{2\varepsilon_w^2} \right) \\ &\geq X \exp \left(\frac{-\|\bar{x} - Q_{\bar{x}}\|^2 - \|\bar{x}' - Q_{\bar{x}'}\|^2 - (1 + C\varepsilon_n^{2/3})(\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 + \|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2)}{2\varepsilon_w^2} \right). \end{aligned} \quad (\text{C.8})$$

Further, by Lemma 29 along with (C.2), it holds that

$$\begin{aligned} \|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 &\leq \|Q_{\bar{x}} - x\|^2 + 2\|Q_{\bar{x}} - x\|\|x - L_{Q_{\bar{x}}}(x)\| + \|x - L_{Q_{\bar{x}}}(x)\|^2 \leq \|Q_{\bar{x}} - x\|^2 + C\varepsilon_n^2 \\ \|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2 &\leq \|Q_{\bar{x}'} - x\|^2 + 2\|Q_{\bar{x}'} - x\|\|x - L_{Q_{\bar{x}'}}(x)\| + \|x - L_{Q_{\bar{x}'}}(x)\|^2 \leq \|Q_{\bar{x}'} - x\|^2 + C\varepsilon_n^2, \end{aligned} \quad (\text{C.9})$$

and that

$$\begin{aligned} C\varepsilon_n^{2/3}\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 &\leq C\varepsilon_n^2 \\ C\varepsilon_n^{2/3}\|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2 &\leq C\varepsilon_n^2. \end{aligned} \quad (\text{C.10})$$

Hence, combining (C.8), (C.9), (C.10) gives us the desired lower bound

$$\omega(\bar{x}, \bar{x}') \geq \exp \left(-\frac{C\varepsilon_n^2}{\varepsilon_w^2} \right) \xi(\bar{x}, \bar{x}') = \exp(-C\varepsilon^2) \xi(\bar{x}, \bar{x}').$$

Now, for the upper bound in (6.5), we first consider the case when

$$\|x - Q_{\bar{x}}\| \geq \varepsilon_n^{2/3} \quad \text{or} \quad \|x - Q_{\bar{x}'}\| \geq \varepsilon_n^{2/3},$$

i.e. when $Y = 1$ in (C.6). Suppose the first inequality holds. Then since $\varepsilon_n^{2/3} \gg \varepsilon_n$, $\|x - Q_{\bar{x}}\| - \|\bar{x} - Q_{\bar{x}}\| \geq c\varepsilon_n^{2/3} > 0$, and consequently,

$$\|\bar{x} - x\|^2 \geq (\|x - Q_{\bar{x}}\| - \|\bar{x} - Q_{\bar{x}}\|)^2 \geq c\varepsilon_n^{4/3}.$$

Therefore,

$$\mathbb{E}_{x \sim \mu} \left[Y \exp \left(\frac{-\|x - \bar{x}\|^2 - \|x - \bar{x}'\|^2}{2\varepsilon_w^2} \right) \right] \leq \exp \left(-\frac{c\varepsilon_n^{4/3}}{\varepsilon_w^2} \right) = \exp \left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}} \right). \quad (\text{C.11})$$

Next, suppose

$$\|x - Q_{\bar{x}}\| < \varepsilon_n^{2/3} \quad \text{and} \quad \|x - Q_{\bar{x}'}\| < \varepsilon_n^{2/3},$$

i.e. when $X = 1$ in (6.1). Then, due to

$$\begin{aligned} \|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 &\geq \|Q_{\bar{x}} - x\|^2 - 2\|Q_{\bar{x}} - x\|\|x - L_{Q_{\bar{x}}}(x)\| + \|x - L_{Q_{\bar{x}}}(x)\|^2 \\ \|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2 &\geq \|Q_{\bar{x}'} - x\|^2 - 2\|Q_{\bar{x}'} - x\|\|x - L_{Q_{\bar{x}'}}(x)\| + \|x - L_{Q_{\bar{x}'}}(x)\|^2, \end{aligned}$$

we get, from (C.2) and Lemma 29,

$$\begin{aligned} \|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 &\geq \|Q_{\bar{x}} - x\|^2 - C\varepsilon_n^2 \\ \|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2 &\geq \|Q_{\bar{x}'} - x\|^2 - C\varepsilon_n^2. \end{aligned} \quad (\text{C.12})$$

Further, by virtue of Lemma 30,

$$\begin{aligned} -\|\bar{x} - x\|^2 &\leq -\|\bar{x} - Q_{\bar{x}}\|^2 - (1 - C\varepsilon_n^{2/3})\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 \\ -\|\bar{x}' - x\|^2 &\leq -\|\bar{x}' - Q_{\bar{x}'}\|^2 - (1 - C\varepsilon_n^{2/3})\|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2, \end{aligned}$$

we get from another application of (C.2)

$$\begin{aligned} -\|\bar{x} - x\|^2 &\leq -\|\bar{x} - Q_{\bar{x}}\|^2 - \|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 + C\varepsilon_n^2 \\ -\|\bar{x}' - x\|^2 &\leq -\|\bar{x}' - Q_{\bar{x}'}\|^2 - \|Q_{\bar{x}'} - L_{Q_{\bar{x}'}}(x)\|^2 + C\varepsilon_n^2. \end{aligned} \quad (\text{C.13})$$

Altogether, (C.12), (C.13) demonstrate

$$\mathbb{E}_{x \sim \mu} \left[X \exp \left(\frac{-\|x - \bar{x}\|^2 - \|x - \bar{x}'\|^2}{2\varepsilon_w^2} \right) \right] \leq \exp \left(\frac{C\varepsilon_n^2}{\varepsilon_w^2} \right) \xi(\bar{x}, \bar{x}') = \exp(C\varepsilon^2) \xi(\bar{x}, \bar{x}'). \quad (\text{C.14})$$

Combining (C.7), (C.11), (C.14), we conclude

$$\omega(\bar{x}, \bar{x}') \leq \exp(C\varepsilon^2) \xi(\bar{x}, \bar{x}') + \exp \left(-\frac{c\varepsilon^{4/3}}{\varepsilon^{2\tau/3}} \right),$$

which is the desired upper bound. ■

C.1 Proof of Lemma 29

Proof We will only prove the first inequality in (C.3); the second inequality is handled similarly. Let $r > 0$. Let $y \in B^d(Q_{\bar{x}}, r) \cap \tilde{T}_{Q_{\bar{x}}}\mathcal{M}$, where $\tilde{T}_{Q_{\bar{x}}}\mathcal{M}$ denotes the affine plane tangent to $\mathcal{M}$ at $Q_{\bar{x}}$. Recalling further that $(A_{Q_{\bar{x}}} \circ \mathbf{l}_m)^{-1}(\tilde{T}_{Q_{\bar{x}}}\mathcal{M}) = T_{Q_{\bar{x}}}\mathcal{M}$ (see (B.1)), we define $E_{Q_{\bar{x}}} := \text{Exp}_{Q_{\bar{x}}} \circ (A_{Q_{\bar{x}}} \circ \mathbf{l}_m)^{-1}$ and write

$$E_{Q_{\bar{x}}}(y) = Q_{\bar{x}} + I_{Q_{\bar{x}}}(y) + \frac{1}{2}P_{Q_{\bar{x}}}(y) + \dots \quad (\text{C.15})$$

where $I_{Q_{\bar{x}}}, P_{Q_{\bar{x}}}$ are linear and quadratic forms on $\tilde{T}_{Q_{\bar{x}}}\mathcal{M}$, respectively, with values in $\mathbb{R}^d$. Consider the plane $\mathcal{P}$ determined by the points $x, Q_{\bar{x}}, L_{Q_{\bar{x}}}(x)$, and let γ be the geodesic curve connecting $Q_{\bar{x}}, x$ that lies in the intersection $\mathcal{P} \cap \mathcal{B}(Q_{\bar{x}}, r)$. Let v be the unit vector in the positive direction of $L_{Q_{\bar{x}}}(x) - Q_{\bar{x}}$. We can suppose that γ is parametrized by its arc length s , such that

$$\gamma(0) = Q_{\bar{x}} \quad \text{and} \quad \frac{d}{ds}\gamma(0) = v.$$

Now, setting $y = vs$ in (C.15) leads to

$$\gamma(s) = E_{Q_{\bar{x}}}(vs) = Q_{\bar{x}} + I_{Q_{\bar{x}}}(v)s + \frac{1}{2}P_{Q_{\bar{x}}}(v)s^2 + \dots \quad (\text{C.16})$$

where $I_{Q_{\bar{x}}}(v) = v = \frac{d}{ds}\gamma(0)$ and $P_{Q_{\bar{x}}}(v) = \frac{d^2}{ds^2}\gamma(0)$. Furthermore, since γ is a geodesic curve on $\mathcal{M}$, we have

$$\frac{d^2}{ds^2}\gamma(0) \perp \tilde{T}_{\gamma(0)}\mathcal{M} = \tilde{T}_{Q_{\bar{x}}}\mathcal{M},$$

and since γ is a plane curve on $\mathcal{P}$, we can write (Struik, 2012, Section 1-4)

$$\frac{d^2}{ds^2}\gamma(0) = \kappa \mathbf{n}$$

where $\mathbf{n}$ is the unit normal vector of γ in $\mathcal{P}$ at $s = 0$, and κ is the curvature of γ at $s = 0$. By the assumptions made on $\mathcal{M}$ in Section 2.2, $|\kappa| \leq K$. Hence, if we let $s = s^* > 0$ be such that $E_{Q_{\bar{x}}}(vs^*) = x$ in (C.16), we can conclude that

$$\|x - Q_{\bar{x}} - vs^*\| = \mathcal{O}((s^*)^2), \quad (\text{C.17})$$

and the majorant constant is bounded by a multiple of K . Since $Q_{\bar{x}} + vs^* = L_{Q_{\bar{x}}}(x)$, we obtain from (C.17),

$$\|x - L_{Q_{\bar{x}}}(x)\| = \mathcal{O}(\|L_{Q_{\bar{x}}}(x) - Q_{\bar{x}}\|^2),$$

and the proof is completed. ■

C.2 Proof of Lemma 30

Proof We will only prove the first inequalities in (C.4) and in (C.5). We prepare two facts. First, it follows from the orthogonality in (C.1) that

$$\|\bar{x} - L_{Q_{\bar{x}}}(x)\|^2 = \|\bar{x} - Q_{\bar{x}}\|^2 + \|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2, \quad (\text{C.18})$$

and second, from Lemma 29,

$$\|x - L_{Q_{\bar{x}}}(x)\| \leq C\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2. \quad (\text{C.19})$$

For the first inequality in (C.4), we have,

$$\begin{aligned} \|x - \bar{x}\|^2 &\leq (\|L_{Q_{\bar{x}}}(x) - \bar{x}\| + \|L_{Q_{\bar{x}}}(x) - x\|)^2 \\ &= \|L_{Q_{\bar{x}}}(x) - \bar{x}\|^2 + \|L_{Q_{\bar{x}}}(x) - x\|^2 + 2\|L_{Q_{\bar{x}}}(x) - \bar{x}\|\|L_{Q_{\bar{x}}}(x) - x\|. \end{aligned} \quad (\text{C.20})$$

Observe that

$$\begin{aligned} \|L_{Q_{\bar{x}}}(x) - \bar{x}\|\|L_{Q_{\bar{x}}}(x) - x\| &\leq C\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 \sqrt{\|\bar{x} - Q_{\bar{x}}\|^2 + \|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2} \\ &\leq C\varepsilon_n^{2/3}\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2, \end{aligned} \quad (\text{C.21})$$

where the first inequality is due to (C.18), (C.19) and the second to (C.2). Putting (C.21) back in (C.20), and invoking all (C.2), (C.18), (C.19) again, allows us to obtain

$$\begin{aligned} \|x - \bar{x}\|^2 &\leq \|Q_{\bar{x}} - \bar{x}\|^2 + \|L_{Q_{\bar{x}}}(x) - Q_{\bar{x}}\|^2 + \|L_{Q_{\bar{x}}}(x) - x\|^2 + C\varepsilon_n^{2/3}\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 \\ &\leq \|Q_{\bar{x}} - \bar{x}\|^2 + \|L_{Q_{\bar{x}}}(x) - Q_{\bar{x}}\|^2 + C\varepsilon_n^{2/3}\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 + C\varepsilon_n^{2/3}\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 \\ &\leq \|Q_{\bar{x}} - \bar{x}\|^2 + (1 + C\varepsilon_n^{2/3})\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2. \end{aligned}$$

Similarly, to handle the first inequality in (C.5), we note

$$\begin{aligned} \|x - \bar{x}\|^2 &\geq (\|L_{Q_{\bar{x}}}(x) - \bar{x}\| - \|L_{Q_{\bar{x}}}(x) - x\|)^2 \\ &= \|L_{Q_{\bar{x}}}(x) - \bar{x}\|^2 + \|L_{Q_{\bar{x}}}(x) - x\|^2 - 2\|L_{Q_{\bar{x}}}(x) - \bar{x}\|\|L_{Q_{\bar{x}}}(x) - x\| \\ &\geq \|L_{Q_{\bar{x}}}(x) - \bar{x}\|^2 - 2\|L_{Q_{\bar{x}}}(x) - \bar{x}\|\|L_{Q_{\bar{x}}}(x) - x\| \\ &\geq \|\bar{x} - Q_{\bar{x}}\|^2 + \|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 - C\varepsilon_n^{2/3}\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2 \\ &= \|\bar{x} - Q_{\bar{x}}\|^2 + (1 - C\varepsilon_n^{2/3})\|Q_{\bar{x}} - L_{Q_{\bar{x}}}(x)\|^2, \end{aligned}$$

where the last inequality is an application of (C.18), (C.21). ■

Appendix D. Proof of Proposition 7

Proof In Appendix D.1, we express the relevant partial derivatives of q as path derivatives and address their continuity. We prove (3.5), (3.6) in Appendices D.2, D.4, respectively. The proof of (3.4) and (3.3) relies on arguments similar to those in Appendix D.2 and is presented in Appendix D.3. ■

D.1 First and Second Partial Derivatives of q as Path Derivatives

Let $z \in \mathcal{M}$. Consider the diffeomorphic map

$$\begin{aligned} \overline{\text{Exp}}_z : B^m(0, \varepsilon_n) \times B^{d-m}(0, \varepsilon_n) &\rightarrow \mathcal{N}_{\varepsilon_n}(\mathcal{B}(z, \varepsilon_n)) \\ (v, k) &\mapsto z' + h \end{aligned}$$

constructed in Appendix B, with $r = s = \varepsilon_n$. Here, $z' = \text{Exp}_z(v)$ and $h = R_{z'} \circ J_{d-m}(k)$. Under this map, a straight line $\Gamma : [0, 1] \rightarrow B^m(0, \varepsilon_n)$, such that $\Gamma(t) = tu$, with $u \in \mathbb{S}^{m-1}$, is mapped to a geodesic curve $\gamma : [0, 1] \rightarrow \mathcal{B}(z, \varepsilon_n)$, such that $\gamma(0) = z$ and

$$\frac{d}{dt}\gamma(t)|_{t=0} = \tilde{u} = \langle \nabla \text{Exp}_z(0), u \rangle \in \tilde{T}_z \mathcal{M} - z$$

is a unit vector. Letting z' denote a generic point on γ , and following definition (3.1), we express $q(z')$ as

$$q(\gamma(t)) = \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{B^{d-m}(0, \varepsilon_n)} \int_{\mathcal{M}} \exp\left(-\frac{\|\text{Exp}_z(tu) + R_{\gamma(t)} \circ J_{d-m}(k) - x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dk.$$

Hence,

$$\begin{aligned} & \frac{d}{dt}q(\gamma(t))|_{t=0} \\ &= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{B^{d-m}(0, \varepsilon_n)} \int_{\mathcal{M}} \frac{d}{dt} \exp\left(-\frac{\|\text{Exp}_z(tu) + R_{\gamma(t)} \circ J_{d-m}(k) - x\|^2}{2\varepsilon_p^2}\right) \Big|_{t=0} \rho(x) d\mathcal{V}(x) dk \end{aligned} \quad (\text{D.1})$$

provides us with the directional derivative $\partial_{\tilde{u}}q(z)$. Observe, since $R_{\gamma(t)}$ is a rotation, we get

$$\left\langle \frac{d}{dt} R_{\gamma(t)} y, R_{\gamma(t)} y \right\rangle = 0, \quad (\text{D.2})$$

for all $y \in \mathbb{R}^d$. Then, from (D.1) and a straightforward calculation,

$$\partial_{\tilde{u}}q(z) = \frac{d}{dt}q(\gamma(t))|_{t=0} = E + Z, \quad (\text{D.3})$$

where, due to the orthogonality $h \perp \tilde{u}$,

$$\begin{aligned} E &:= -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \frac{\langle z + h - x, \tilde{u} \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z + h - x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ &= -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \frac{\langle z - x, \tilde{u} \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z + h - x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \end{aligned}$$

and due to (D.2),

$$Z := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \frac{\langle z - x, Ah \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z + h - x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh,$$

where $A := \frac{d}{dt} R_{\gamma(t)}|_{t=0}$.

Let $w \in \mathbb{S}^{m-1}$ and $\tilde{w} = \nabla \text{Exp}_z(0)w \in \tilde{T}_z \mathcal{M} - z$. Repeating this process for $\partial_{\tilde{u}}q$ in place of q , and $w, \tilde{w}$ in place of $u, \tilde{u}$, we interpret the second partial derivative $\partial_{\tilde{w}}\partial_{\tilde{u}}q$ as the time derivative of $\partial_{\tilde{u}}q(\zeta(t))$ at $t = 0$, where $\zeta : [0, 1] \rightarrow \mathcal{B}(z, \varepsilon_n)$ is a geodesic curve on $\mathcal{M}$ starting at z with velocity $\tilde{w}$. Thus, from (D.3)

$$\partial_{\tilde{w}}\partial_{\tilde{u}}q(z) = \frac{d}{dt}\partial_{\tilde{u}}q(\zeta(t))|_{t=0} = D + V + F + S \quad (\text{D.4})$$

where¹⁵

$$D := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \frac{\langle z-x, \tilde{u} + Ah \rangle \langle z-x, \tilde{w} + Ah \rangle}{\varepsilon_p^4} \\ \times \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh,$$

and

$$V := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \frac{\langle \tilde{w}, \tilde{u} \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ F := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \frac{\langle \tilde{w}, Ah \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ S := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M}} \frac{\langle z-x, Bh \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh,$$

where $B := \frac{d^2}{dt^2} R_{\gamma(t)}|_{t=0}$.

We conclude by observing from (D.3), (D.4) that $\partial_{\tilde{u}} q$, $\partial_{\tilde{w}} \partial_{\tilde{u}} q$ are both continuous in terms of z .

D.2 Proof of (3.5) in Proposition 7

Proof From (D.3), let $E = E_1 + E_2$, where

$$E_1 := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \tilde{u} \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \\ E_2 := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \tilde{u} \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh.$$

Starting with E_2 , we note from (2.7) that, $d(z, x) \geq \varepsilon_n^{2/3}$ implies $\|x - z\| \geq c\varepsilon_n^{2/3}$. Further, from the fact that $\|h\| \leq \varepsilon_n$, we also acquire

$$c\varepsilon_n^{2/3} \leq cd(z, x) \leq \|z - x\| \leq c'(\|z - x\| - \|h\|) \leq c'\|z + h - x\|. \quad (\text{D.5})$$

Recall the quantity $\mathfrak{D} = \sup\{\|x - x'\| : x, x' \in \mathcal{M}\}$ in (2.16). Since $\mathcal{M}$ is a closed manifold, $\mathfrak{D} < \infty$. Then gathering all in (D.5), we deduce,

$$|E_2| \leq \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\|z-x\|}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ \leq \frac{\mathfrak{D}}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{1}{\varepsilon_p^2} \exp\left(-\frac{c}{\eta^{2/d} \varepsilon_n^{2/3}}\right) \rho(x) d\mathcal{V}(x) dh. \quad (\text{D.6})$$

By taking ε sufficiently small satisfying (3.2), we acquire from (D.6),

$$|E_2| \leq \frac{C_\eta \varepsilon_n^m}{\eta^{2/d} (\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} dh \leq C_\eta. \quad (\text{D.7})$$

15. The letters E, Z, D, V, F, S are chosen correspondingly to the German words for “first, second, third, fourth, fifth, sixth” which are “Erster, Zweiter, Dritter, Vierter, Fünfter, Sechster” respectively.

Moving on to E_1 , recall that we can associate for each $x \in \mathcal{B}(z, \varepsilon_n^{2/3})$, a unique $L_z(x) - z \in \tilde{T}_z \mathcal{M} - z$ such that Lemma 29 holds. To simplify the notation, we will write $v_x := L_z(x) - z$ from here onward. Then

$$\|v_x\| = d(z, x) \quad \text{and} \quad \|x - (v_x + z)\| = \|x - L_z(x)\| \leq C\|L_z(x) - z\|^2 = Cd(z, x)^2, \quad (\text{D.8})$$

and hence, due to the orthogonality $v_x \perp h$,

$$\|z + h - x\|^2 = \|z - x\|^2 + \|h\|^2 + 2\langle z - x, h \rangle = \|z - x\|^2 + \|h\|^2 + 2\langle z + v_x - x, h \rangle.$$

The following straightforward lemma holds.

Lemma 31 *For $x \in \mathcal{B}(z, \varepsilon_n^{2/3})$, it holds that*

$$\exp\left(-\frac{\|x - z\|^2}{2\varepsilon_p^2} + \frac{d(z, x)^2}{2\varepsilon_p^2}\right) \leq \exp\left(\frac{C\|x - z\|^4}{\varepsilon_p^2}\right) \leq 1 + \frac{C\|x - z\|^4}{\varepsilon_p^2}, \quad (\text{D.9})$$

and that

$$\exp\left(-\frac{Cd(z, x)^2\|h\|}{\varepsilon_p^2}\right) \leq \exp\left(-\frac{\langle z + v_x - x, h \rangle}{\varepsilon_p^2}\right) \leq \exp\left(\frac{Cd(z, x)^2\|h\|}{\varepsilon_p^2}\right). \quad (\text{D.10})$$

Furthermore, if $h + z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})$, then

$$\begin{aligned} \exp\left(-\frac{\langle z + v_x - x, h \rangle}{\varepsilon_p^2}\right) &= 1 - \frac{\langle z + v_x - x, h \rangle}{\varepsilon_p^2} + \frac{\langle z + v_x - x, h \rangle^2}{2\varepsilon_p^4} - \frac{\langle z + v_x - x, h \rangle^3}{3!\varepsilon_p^6} \\ &\quad + \frac{\langle z + v_x - x, h \rangle^4}{4!\varepsilon_p^8} - \frac{\langle z + v_x - x, h \rangle^5}{5!\varepsilon_p^{10}} + \mathcal{O}_\eta(\varepsilon_n^2). \end{aligned} \quad (\text{D.11})$$

Proof [Proof of Lemma 31] Observe that (D.9) follows from a combination of (2.7) and Taylor's expansion, (D.10) from (D.8), and (D.11) from Taylor's expansion again. $\blacksquare$

Coming back to E_1 , we split $E_1 = E_{11} + E_{12}$, where

$$\begin{aligned} E_{11} &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle \tilde{u}, z + v_x - x \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z + h - x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \\ E_{12} &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} -\frac{\langle \tilde{u}, v_x \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z + h - x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh. \end{aligned}$$

Guided by Lemma 31, we let

$$\begin{aligned} E_{11}^* &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle \tilde{u}, z + v_x - x \rangle}{\varepsilon_p^2} \\ &\quad \times \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{\|x - z\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \end{aligned}$$

and

$$E_{12}^* := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} -\frac{\langle \tilde{u}, v_x \rangle}{\varepsilon_p^2} \\ \times \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{\|x-z\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh.$$

Then another lemma applies.

Lemma 32 *Let E_{1i}^* be defined as above, for $i = 1, 2$. Then $|E_{1i}^*| \leq C$.*

Proof [Proof of Lemma 32] It is straightforward from (2.7), (D.8) and a familiar calculation that

$$|E_{11}^*| \leq \frac{C}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{d(z, x)^2}{\varepsilon_p^2} \\ \times \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{cd(z, x)^2}{\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ = \frac{C\varepsilon_p^{d-m}}{(\varepsilon_p \sqrt{2\pi})^d} \int_{[-\eta^{-1/d}, \eta^{-1/d}]^{d-m}} \exp\left(-\frac{\|h\|^2}{2}\right) dh \int_{B^m(0, \varepsilon_n^{2/3}) \subset T_z \mathcal{M}} \frac{\|y\|^2}{\varepsilon_p^2} \\ \times \exp\left(-\frac{c\|y\|^2}{\varepsilon_p^2}\right) \tilde{\rho}(y) J_z(y) dy \\ \leq C.$$

We next focus on a version of E_{12}^* , which is

$$E_{12}^{**} := -\frac{\rho(z)}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle \tilde{u}, v_x \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{d(z, x)^2}{2\varepsilon_p^2}\right) d\mathcal{V}(x) dh.$$

Note, for a fixed $h + z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})$, if $x, x' \in \mathcal{M}$ are such that $v_x = -v_{x'}$, then $d(z, x) = d(z, x')$, and conversely. Therefore

$$-\int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle \tilde{u}, v_x \rangle}{\varepsilon_p^2} \exp\left(-\frac{d(z, x)^2}{2\varepsilon_p^2}\right) d\mathcal{V}(x) = 0,$$

and we conclude that $E_{12}^{**} = 0$. Thus, by virtue of Lemma 31, particularly (D.9), the Lipschitz continuity of ρ , and (D.8),

$$|E_{12}^* - E_{12}^{**}| \leq \frac{C}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{d(z, x)^2}{\varepsilon_p^2} \\ \times \exp\left(-\frac{cd(z, x)^2}{\varepsilon_p^2}\right) d\mathcal{V}(x) dh \\ \leq C.$$

Henceforth, $|E_{12}^*| \leq C$. ■

In view of Lemma 31, particularly (D.11),

$$|E_{11} - E_{11}^*| \leq \frac{C\varepsilon_n}{(\varepsilon_p\sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|y\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{d(z, x)^4}{\varepsilon_p^4} \exp\left(-\frac{cd(z, x)^2}{\varepsilon_p^2}\right) d\mathcal{V}(x) dh \leq C\varepsilon_n, \quad (\text{D.12})$$

and

$$|E_{12} - E_{12}^*| \leq \frac{C}{(\eta^{1/d}\varepsilon_p\sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|y\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{d(z, x)^3}{\varepsilon_p^3} \times \exp\left(-\frac{cd(z, x)^2}{\varepsilon_p^2}\right) d\mathcal{V}(x) dh \leq C_\eta.$$

Thus, in view of Lemma 32, $|E_{1i}| \leq C_\eta$. Consequently,

$$|E_1| = |E_{11} + E_{12}| \leq C_\eta. \quad (\text{D.13})$$

Combining (D.7), (D.13), we arrive at

$$|E| = |E_1 + E_2| \leq C_\eta. \quad (\text{D.14})$$

Finally, regarding Z in (D.3), we note that Z can also be split into

$$\begin{aligned} Z_1 &:= -\frac{1}{(\varepsilon_p\sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \mathbf{A}h \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ &= -\frac{1}{(\varepsilon_p\sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \|Ah\| \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, Ah/\|Ah\| \rangle}{\varepsilon_p^2} \\ &\quad \times \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \end{aligned}$$

and

$$Z_2 := -\frac{1}{(\varepsilon_p\sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \mathbf{A}h \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh.$$

Since $|\langle z-x, \mathbf{A}h \rangle| \leq C\|z-x\|\varepsilon_n$, following a similar calculation to that in (D.6), (D.7), we get, for ε sufficiently small

$$|Z_2| \leq \frac{C_\eta \mathfrak{D}}{(\varepsilon_p\sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{1}{\varepsilon_p} \exp\left(-\frac{c}{\eta^{2/d}\varepsilon_n^{2/3}}\right) \rho(x) d\mathcal{V}(x) dh \leq C_\eta. \quad (\text{D.15})$$

For Z_1 , we draw an observation from the expression of Z_1 and the treatment of E_1 , particularly from the proof of Lemma 32 that

$$|Z_1| \leq C_\eta \varepsilon_n. \quad (\text{D.16})$$

Then (D.15), (D.16) imply

$$|Z| = |Z_1 + Z_2| \leq C_\eta. \quad (\text{D.17})$$

Altogether, (D.3), (D.14), (D.17) conclude $|\partial_{\tilde{u}} q(z)| \leq |E| + |Z| \leq C_\eta$, as desired. $\blacksquare$

D.3 Proof of (3.3) and (3.4) in Proposition 7

Proof To establish the boundedness of q in (3.1), we reuse the arguments presented in Appendix D.2. Specifically, we observe from the integral expression (3.1) that $q(z)$ can be approximated by

$$A := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \exp\left(-\frac{d(z, x)^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh. \quad (\text{D.18})$$

Indeed, using reasoning akin to that employed for estimating E_1 in Appendix D.2, we find that $|q(z) - A|$ to be bounded by a universal constant multiple of $B_1 + B_2$, where

$$\begin{aligned} B_1 &:= \frac{\varepsilon_n}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|y\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{d(z, x)^2}{\varepsilon_p^2} \exp\left(-\frac{d(z, x)^2}{2\varepsilon_p^2}\right) d\mathcal{V}(x) dh \\ B_2 &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh. \end{aligned}$$

To derive B_1 , we use (D.9), (D.11), and a similar approach to the one yielding (D.12). We skip further details. Then, $B_1 \leq \varepsilon_n$, and by referencing the calculations used to bound $|E_2|$ in (D.6), albeit with a simpler argument, we can show

$$B_2 = \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \leq C_\eta \varepsilon_n^2.$$

Back to A in (D.18), we write

$$\begin{aligned} A &= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \exp\left(-\frac{d(z, x)^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) (\rho(x) - \rho(z)) d\mathcal{V}(x) dh \\ &\quad + \frac{\rho(z)}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \exp\left(-\frac{d(z, x)^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) d\mathcal{V}(x) dh, \end{aligned} \quad (\text{D.19})$$

where J_z is the Jacobian matrix. Using a now-standard argument and the Lipschitz continuity of ρ , we find that the absolute value of the first integral in (D.19) is bounded above by $C_\eta \varepsilon_n^{2/3}$, while the second simplifies to

$$\begin{aligned} &\frac{\rho(z)}{(\varepsilon_p \sqrt{2\pi})^d} \int_{[-\varepsilon_n, \varepsilon_n]^{d-m}} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) dh \int_{B^m(0, \varepsilon_n^{2/3})} \exp\left(-\frac{\|x\|^2}{2\varepsilon_p^2}\right) J_z(x) dx \\ &= \frac{\rho(z)}{(2\pi)^{d/2}} \left(\int_{[-\eta^{-1/d}, \eta^{-1/d}]^{d-m}} \exp\left(-\frac{\|h\|^2}{2}\right) dh \int_{B^m(0, \eta^{-1/d} \varepsilon_n^{-1/3})} \exp\left(-\frac{\|x\|^2}{2}\right) dx \right) (1 + \mathcal{O}(\varepsilon_n^{4/3})) \\ &= \rho(z) + \mathcal{O}_\eta(\varepsilon_n^{4/3}). \end{aligned}$$

Combining all the derivations, we conclude that $q(z) = \rho(z) + \mathcal{O}_\eta(\varepsilon_n^{2/3})$, which is (3.3). Then (3.4) follows as a consequence of (3.3). $\blacksquare$

D.4 Proof of (3.6) in Proposition 7

Proof Recall the expression (D.4) in Appendix D.1 that describes $\partial_{\tilde{u}}\partial_{\tilde{w}}q(z)$. We proceed to establish that all quantities D, V, F, S in (D.4) are bounded. As usual, we split each quantity into two parts, a local part, and a tail part. Specifically, we let

$$D_1 := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \tilde{u}+Ah \rangle \langle z-x, \tilde{w}+Ah \rangle}{\varepsilon_p^4} \\ \times \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh,$$

and

$$V_1 := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle \tilde{w}, \tilde{u} \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ F_1 := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle \tilde{w}, Ah \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ S_1 := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, Bh \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh,$$

and $D_2 := D - D_1$, $V_2 := V - V_1$, $F_2 := F - F_1$, $S_2 := S - S_1$.

Bounding F . Let

$$F_1^* := -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \left(\int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \frac{\langle \tilde{w}, Ah \rangle}{\varepsilon_p^2} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) dh \right) \\ \times \left(\int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \exp\left(-\frac{\|x-z\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) \right).$$

Due to symmetry, we have $F_1^* = 0$. Then applying (D.8) and Lemma 31, we obtain

$$|F_1 - F_1^*| \leq \frac{C_\eta}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{d(z, x)^2}{\varepsilon_p} \\ \times \exp\left(-\frac{cd(z, x)^2}{\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \leq C_\eta \varepsilon_n^{2/3}.$$

Thus, $|F_1| \leq C_\eta \varepsilon_n^{2/3}$. Moreover, by employing a similar argument as in (D.6), (D.7)

$$|F_2| \leq \frac{C}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \frac{\|h\|}{\eta^{2/d} \varepsilon_p} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{1}{\varepsilon_p} \exp\left(-\frac{c}{\varepsilon_n^{2/3}}\right) \rho(x) d\mathcal{V}(x) dh \leq C_\eta.$$

Gathering all this, we conclude $|F| \leq C_\eta$.

Bounding S . We note that the expression of S is analogous to that of Z in (D.3). Therefore, by applying a similar argument to deduce (D.17) previously, we obtain $|S| \leq C_\eta$.

Bounding D and V . We write $D_1 = D_{11} + D_{12} + D_{13} + D_{14}$, where

$$\begin{aligned} D_{11} &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \tilde{u} \rangle}{\varepsilon_p^2} \frac{\langle z-x, \tilde{w} \rangle}{\varepsilon_p^2} \\ &\quad \times \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \\ D_{12} &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \tilde{u} \rangle}{\varepsilon_p^2} \frac{\langle z-x, Ah \rangle}{\varepsilon_p^2} \\ &\quad \times \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \\ D_{13} &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, Ah \rangle}{\varepsilon_p^2} \frac{\langle z-x, \tilde{w} \rangle}{\varepsilon_p^2} \\ &\quad \times \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \\ D_{14} &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, Ah \rangle}{\varepsilon_p^2} \frac{\langle z-x, Ah \rangle}{\varepsilon_p^2} \\ &\quad \times \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh. \end{aligned}$$

We will now bound D_{12} , D_{13} , D_{14} , $D_{11} + V_1$, sequentially. Starting with D_{12} , we consider

$$\begin{aligned} D_{12}^* &:= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \tilde{u} \rangle}{\varepsilon_p^2} \frac{\langle z-x, Ah \rangle}{\varepsilon_p^2} \\ &\quad \times \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{\|z-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh. \end{aligned}$$

Then by integrating in terms of h first, we get $D_{12}^* = 0$. Moreover, by Lemma 31

$$\begin{aligned} |D_{12} - D_{12}^*| &\leq \frac{C_\eta}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{d(z, x)^4}{\varepsilon_p^4} \\ &\quad \times \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \exp\left(-\frac{cd(z, x)^2}{\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \leq C_\eta, \end{aligned}$$

and we get $|D_{12}| \leq C_\eta$. Due to the analogous expressions, it follows that $|D_{13}| \leq C_\eta$ as well. As for D_{14} , we see simply that

$$\begin{aligned} |D_{14}| &\leq \frac{C}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\|z-x\|^2 \|h\|^2}{\varepsilon_p^4} \exp\left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \\ &\leq C_\eta. \end{aligned}$$

To facilitate the upcoming calculations for $D_{11} + V_1$, we use a translation by $-z$ to temporarily set $z = 0$, and use the rotation A_0^{-1} , as discussed in Appendix B, to align $\tilde{T}_0 \mathcal{M}$

with $\mathbb{R}^m \times \{0\} \subset \mathbb{R}^d$. Then $D_{11} + V_1$ becomes

$$\begin{aligned} & \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h \in \tilde{N}_0 \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(0, \varepsilon_n^{2/3})} \left(\frac{\langle x, \tilde{u} \rangle \langle x, \tilde{w} \rangle}{\varepsilon_p^2} - \frac{\langle \tilde{u}, \tilde{w} \rangle}{\varepsilon_p^2} \right) \\ & \quad \times \exp\left(-\frac{\|x\|^2}{2\varepsilon_p^2}\right) \mathfrak{f}(x, h) \rho(x) d\mathcal{V}(x) dh, \end{aligned}$$

where, from Lemma 31,

$$\mathfrak{f}(x, h) := \exp\left(-\frac{\langle v_x - x, h \rangle}{\varepsilon_p^2}\right) = 1 + \sum_{i=1}^5 \mathfrak{f}_i(x, h) + \mathcal{O}_\eta(\varepsilon_n^2) := 1 + \sum_{i=1}^5 \frac{(-1)^i \langle v_x - x, h \rangle^i}{i! \varepsilon_p^{2i}} + \mathcal{O}_\eta(\varepsilon_n^2). \quad (\text{D.20})$$

Note that, if i is odd, then $\mathfrak{f}_i(x, h)$ is odd in terms of h , and if i is even, then

$$|\mathfrak{f}_i(x, h)| \leq \frac{C \|x\|^{2i} \|h\|^i}{\varepsilon_p^{2i}}, \quad (\text{D.21})$$

due to (D.8). Let

$$\begin{aligned} I := & \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h \in \tilde{N}_0 \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(0, \varepsilon_n^{2/3})} \left(\frac{\langle x, \tilde{u} \rangle \langle x, \tilde{w} \rangle}{\varepsilon_p^2} - \frac{\langle \tilde{u}, \tilde{w} \rangle}{\varepsilon_p^2} \right) \\ & \quad \times \exp\left(-\frac{\|x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh. \end{aligned}$$

Then using the given observation, and particularly (D.20), (D.21), we can dominate $|D_{11} + V_1 - I|$ with a C_η multiple of $I_1 + I_2 + I_3$, where

$$\begin{aligned} I_i := & \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h \in \tilde{N}_0 \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(0, \varepsilon_n^{2/3})} \frac{\|x\|^{4i}}{\varepsilon_p^{4i}} \left(\frac{\|x\|^2}{\varepsilon_p^2} + 1 \right) \\ & \quad \times \exp\left(-\frac{\|x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh, \end{aligned}$$

for $i = 1, 2$, and

$$I_3 := \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h \in \tilde{N}_0 \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(0, \varepsilon_n^{2/3})} \left(\frac{\|x\|^2}{\varepsilon_p^2} + 1 \right) \exp\left(-\frac{\|x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh.$$

It is straightforward to see that $I_i \leq C_\eta$. To calculate I , we note that $v_x \in \mathbb{R}^m \times \{0\}$ is $\text{Log}_0(x) \in \mathbb{R}^m$ canonically embedded in $\mathbb{R}^d$. Moreover, since $\tilde{u}, \tilde{w} \in \mathbb{R}^m \times \{0\}$, they are canonically embedded images of some $\mathbf{u}, \mathbf{w} \in \mathbb{R}^m$. Therefore, using Lemma 31, particularly (D.9), we can bound $|I|$ by a sum of the integral

$$\begin{aligned} & \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h \in \tilde{N}_0 \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{\mathcal{B}(0, \varepsilon_n^{2/3})} \frac{\|x\|^2}{\varepsilon_p^2} \left(\frac{\|x\|^4}{\varepsilon_p^4} + \frac{\|x\|^2}{\varepsilon_p^2} \right) \\ & \quad \times \exp\left(-\frac{\|x\|^2}{2\varepsilon_p^2}\right) \rho(x) d\mathcal{V}(x) dh \leq C_\eta, \end{aligned}$$

and another integral, expressed in local coordinates as

$$\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h \in \tilde{N}_0 \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \exp\left(-\frac{\|h\|^2}{2\varepsilon_p^2}\right) \int_{B^m(0, \varepsilon_n^{2/3})} \left(\frac{\langle y, u \rangle}{\varepsilon_p^2} \frac{\langle y, w \rangle}{\varepsilon_p^2} - \frac{\langle u, w \rangle}{\varepsilon_p^2}\right) \\ \times \exp\left(-\frac{\|y\|^2}{2\varepsilon_p^2}\right) \tilde{\rho}(y) J_0(y) dy dh.$$

Let I^* be the inner integral above; that is,

$$I^* := \int_{B^m(0, \varepsilon_n^{2/3})} \left(\frac{\langle y, u \rangle}{\varepsilon_p^2} \frac{\langle y, w \rangle}{\varepsilon_p^2} - \frac{\langle u, w \rangle}{\varepsilon_p^2}\right) \exp\left(-\frac{\|y\|^2}{2\varepsilon_p^2}\right) \tilde{\rho}(y) J_0(y) dy.$$

For $y \in B^m(0, \varepsilon_n^{2/3})$, we rely on Taylor's expansion to get

$$\tilde{\rho}(y) = \tilde{\rho}(0) + \langle \nabla \tilde{\rho}(0), y \rangle + \frac{1}{2} \langle D^2 \tilde{\rho}(0) y, y \rangle + \mathcal{O}_\eta(\varepsilon_p^2), \quad (\text{D.22})$$

and on (2.10) to get

$$|J_0(y) - 1| \leq C \|y\|^2. \quad (\text{D.23})$$

Applying (D.22), (D.23), and a standard argument that uses linear symmetry for cancellation, we obtain

$$I^* = \varepsilon_p^m \int_{B^m(0, \eta^{-1/d} \varepsilon_n^{-1/3})} \left(\frac{\langle y, u \rangle}{\varepsilon_p^2} \frac{\langle y, w \rangle}{\varepsilon_p^2} - \frac{\langle u, w \rangle}{\varepsilon_p^2}\right) \exp\left(-\frac{\|y\|^2}{2}\right) \tilde{\rho}(0) dy + \mathcal{O}_\eta(\varepsilon_p^m). \quad (\text{D.24})$$

We now introduce a basic estimate for the m -dimensional Gaussian distribution on Euclidean spaces.

Lemma 33 *Let $\sigma > 0$. Let $X \in \mathbb{R}^m$ be a centered Gaussian distributed random variable such that*

$$f_X(x) = \frac{1}{(\sigma \sqrt{2\pi})^m} \exp\left(-\frac{\|x\|^2}{2\sigma^2}\right).$$

Then for any $v, w \in \mathbb{R}^m$ we have

$$\int_{\mathbb{R}^m} \left(\frac{\langle u, x \rangle}{\sigma^2} \frac{\langle w, x \rangle}{\sigma^2} - \frac{\langle u, w \rangle}{\sigma^2}\right) f_X(x) dx = 0.$$

Proof [Proof of Lemma 33] We first show the result for the case $m = 1$.

$$\int_{\mathbb{R}^m} \left(\frac{\langle u, x \rangle}{\sigma^2} \frac{\langle w, x \rangle}{\sigma^2} - \frac{\langle u, w \rangle}{\sigma^2}\right) f_X(x) dx = \int_{\mathbb{R}^m} \left(\frac{x^2 u w}{\sigma^4} - \frac{u w}{\sigma^2}\right) f_X(x) dx = \frac{u w}{\sigma^2} - \frac{u w}{\sigma^2} = 0,$$

as it becomes a computation of the second moment of Gaussian distribution. For the case $m > 1$, since we can decompose the high-dimensional Gaussian distribution X into $(X^1, \dots, X^m)$ where X^i are i.i.d. one-dimensional centered Gaussian distributed random variable, and the problem reduces to one-dimensional problem. The proof is now complete. $\blacksquare$

Applying Lemma 33 (with $\sigma = 1$) to (D.24) gives us

$$I^* = -\varepsilon_p^m \int_{\mathbb{R}^m - B^m(0, \eta^{-1/d} \varepsilon_n^{-1/3})} \left(\frac{\langle y, \mathbf{u} \rangle \langle y, \mathbf{w} \rangle}{\varepsilon_p^2} - \frac{\langle \mathbf{u}, \mathbf{w} \rangle}{\varepsilon_p^2} \right) \exp \left(-\frac{\|y\|^2}{2} \right) \tilde{\rho}(0) dy + \mathcal{O}_\eta(\varepsilon_p^m),$$

which, due to the rapid Gaussian tail decay, implies that $|I^*| \leq C_\eta \varepsilon_p^m$. Thus, we conclude that $|I| \leq C_\eta$. In addition, by recalling $|D_{11} + V_1 - I| \leq |I_1| + |I_2| + |I_3| \leq C_\eta$, we conclude $|D_{11} + V_1| \leq C_\eta$. Finally, since

$$\begin{aligned} D_2 &= D - D_1 \\ &= \frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle z-x, \tilde{u} + \mathbf{A}h \rangle \langle z-x, \tilde{w} + \mathbf{A}h \rangle}{\varepsilon_p^4} \\ &\quad \times \exp \left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2} \right) \rho(x) d\mathcal{V}(x) dh, \end{aligned}$$

and

$$\begin{aligned} V_2 &= V - V_1 \\ &= -\frac{1}{(\varepsilon_p \sqrt{2\pi})^d} \int_{h+z \in \tilde{N}_z \mathcal{M} \cap \mathcal{N}_{\varepsilon_n}(\mathcal{M})} \int_{\mathcal{M} \setminus \mathcal{B}(z, \varepsilon_n^{2/3})} \frac{\langle \tilde{w}, \tilde{u} \rangle}{\varepsilon_p^2} \exp \left(-\frac{\|z+h-x\|^2}{2\varepsilon_p^2} \right) \rho(x) d\mathcal{V}(x) dh, \end{aligned}$$

employing the rapid Gaussian tail decay, as in (D.6), (D.7), yields, $|D_2|, |V_2| \leq C_\eta$. Putting altogether, we get $|D + V| \leq |D_1 + V_1| + |D_2| + |V_2| \leq C_\eta$, and moreover,

$$|\partial_{\tilde{w}} \partial_{\tilde{u}} q(z)| \leq |D + V + F + S| \leq |D + V| + |F| + |S| \leq C_\eta,$$

as desired. ■

Appendix E. Proof of Lemma 18

Proof Define

$$\begin{aligned} \varsigma^1(x, y) &:= \int_{\{z \in \mathbb{R}^d : d(x, z) < \varepsilon_n^{2/3} \wedge \|y-z\| < \varepsilon_n^{2/3}\} \cap \mathcal{M}} \exp \left(-\frac{\|x-z\|^2}{2\varepsilon_w^2} \right) \exp \left(-\frac{\|y-z\|^2}{2\varepsilon_w^2} \right) \rho(z) d\mathcal{V}(z) \\ \varsigma^2(x, y) &:= \int_{\{z \in \mathbb{R}^d : d(x, z) < \varepsilon_n^{2/3} \wedge d(y, z) < \varepsilon_n^{2/3}\} \cap \mathcal{M}} \exp \left(-\frac{\|x-z\|^2}{2\varepsilon_w^2} \right) \exp \left(-\frac{\|y-z\|^2}{2\varepsilon_w^2} \right) \rho(z) d\mathcal{V}(z). \end{aligned}$$

Then, we observe

$$\begin{aligned} |\varsigma(x, y) - \varsigma^1(x, y)| &\leq \int_{(B^d(x, \varepsilon_n^{2/3}) \setminus \mathcal{B}(x, \varepsilon_n^{2/3})) \cap \mathcal{M}} \exp \left(-\frac{\|x-z\|^2}{2\varepsilon_w^2} \right) \rho(z) d\mathcal{V}(z) \\ &\leq \int_{\mathcal{B}(x, 2\varepsilon_n^{2/3}) \setminus \mathcal{B}(x, \varepsilon_n^{2/3})} \exp \left(-\frac{cd(x, z)^2}{\varepsilon_w^2} \right) \rho(z) d\mathcal{V}(z), \end{aligned}$$

as a consequence of (2.7). Therefore,

$$\begin{aligned}
 |\varsigma(x, y) - \varsigma^1(x, y)| &\leq \int_{(B^m(0, 2\varepsilon_n^{2/3}) \setminus B^m(0, \varepsilon_n^{2/3}))} \exp\left(-\frac{c\|u\|^2}{\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du \\
 &\leq C\varepsilon_w^m \int_{\mathbb{R}^m \setminus B^m(0, \varepsilon_w^{-1/3} \varepsilon^{2/3})} \exp(-c\|u\|^2) du \\
 &\leq C\varepsilon_w^m \varepsilon^2,
 \end{aligned} \tag{E.1}$$

once ε is sufficiently small. Applying the same approach for $|\varsigma^1(x, y) - \varsigma^2(x, y)|$, we also have

$$|\varsigma^1(x, y) - \varsigma^2(x, y)| \leq C\varepsilon_w^m \varepsilon^2. \tag{E.2}$$

Together, (E.1), (E.2) yield

$$\varsigma(x, y) = \varsigma^2(x, y) + \mathcal{O}(\varepsilon_w^m \varepsilon^2). \tag{E.3}$$

Continuing, we define

$$\begin{aligned}
 \varsigma^3(x, y) &:= \int_{\{z \in \mathbb{R}^d : d(x, z) < \varepsilon_n^{2/3} \wedge d(y, z) < \varepsilon_n^{2/3}\} \cap \mathcal{M}} \exp\left(-\frac{d(x, z)^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|y - z\|^2}{2\varepsilon_w^2}\right) \rho(z) d\mathcal{V}(z) \\
 \varsigma^4(x, y) &:= \int_{\{z \in \mathbb{R}^d : d(x, z) < \varepsilon_n^{2/3} \wedge d(y, z) < \varepsilon_n^{2/3}\} \cap \mathcal{M}} \exp\left(-\frac{d(x, z)^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{d(y, z)^2}{2\varepsilon_w^2}\right) \rho(z) d\mathcal{V}(z).
 \end{aligned}$$

Note that $\varsigma^*(x, y) = \varsigma^4(x, y)$. Then it can be readily checked from (2.7) that $\varsigma^3(x, y) \leq \varsigma^2(x, y)$. Moreover, for $\|x - z\| < \varepsilon_n^{2/3}$,

$$\exp\left(-\frac{\|x - z\|^2}{2\varepsilon_w^2}\right) \leq \exp\left(-\frac{d(x, z)^2}{2\varepsilon_w^2}\right) \exp\left(\frac{C\|x - z\|^4}{\varepsilon_w^2}\right) \leq \exp\left(-\frac{d(x, z)^2}{2\varepsilon_w^2}\right) \left(1 + C\varepsilon_w^{2/3} \varepsilon^{8/3}\right), \tag{E.4}$$

where Taylor expansion has been applied to the last inequality, akin to the derivation of (D.9). Hence, from (E.4), together with (6.19),

$$\varsigma^2(x, y) \leq \varsigma^3(x, y) + C\varepsilon_w^m (\varepsilon_w^{2/3} \varepsilon^{8/3}). \tag{E.5}$$

By a similar argument, we can show that

$$\varsigma^3(x, y) = \varsigma^4(x, y) + \mathcal{O}(\varepsilon_w^m (\varepsilon_w^{2/3} \varepsilon^{8/3})). \tag{E.6}$$

Combining (E.3), (E.5), (E.6), we get

$$\varsigma(x, y) = \varsigma^4(x, y) + \mathcal{O}(\varepsilon_w^m \varepsilon^2) = \varsigma^*(x, y) + \mathcal{O}(\varepsilon_w^m \varepsilon^2),$$

which is what we want. ■

Appendix F. Proof of Lemma 19

Proof We first demonstrate (6.32). Express ς_x^* , via a change of variables using the exponential map Exp_x (2.8), as

$$\varsigma_0^*(v) = \int_{\tilde{\mathcal{U}}_{0,v}} \exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|u-v\|^2}{2\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du, \quad (\text{F.1})$$

where $\tilde{\mathcal{U}}_{0,v} := \{u : \|u\| < \varepsilon_n^{2/3} \wedge \|u-v\| < \varepsilon_n^{2/3}\} \cap \mathcal{M}$, and $\text{Exp}_x(v) = y$. For instance, $\tilde{\mathcal{U}}_{0,0} = B^m(0, \varepsilon_n^{2/3})$. Then by an application of Taylor's expansion and a change in variables,

$$\begin{aligned} \varsigma_0^*(0) &= \int_{B^m(0, \varepsilon_n^{2/3})} \exp\left(-\frac{\|u\|^2}{\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du \\ &= \varepsilon_w^m \int_{B^m(0, \varepsilon_w^{-1/3} \varepsilon_n^{2/3})} \exp(-\|u\|^2) (\tilde{\rho}(0) + \mathcal{O}(\varepsilon_n^{2/3})) (1 + \mathcal{O}(\varepsilon_n^{2/3})) du. \end{aligned} \quad (\text{F.2})$$

It is readily apparent that

$$\begin{aligned} &\int_{B^m(0, \varepsilon_w^{-1/3} \varepsilon_n^{2/3})} \exp(-\|u\|^2) \tilde{\rho}(0) du \\ &= \tilde{\rho}(0) \int_{\mathbb{R}^m} \exp(-\|u\|^2) du - \tilde{\rho}(0) \int_{\mathbb{R}^m \setminus B^m(0, \varepsilon_w^{-1/3} \varepsilon_n^{2/3})} \exp(-\|u\|^2) du \\ &= \alpha \rho(x) + \mathcal{O}(\varepsilon) \end{aligned} \quad (\text{F.3})$$

due to the rapid Gaussian decay, once ε is sufficiently small. The claim (6.32) now follows from a combination of (F.2), (F.3).

We demonstrate (6.33) next. For $v \in B^m(0, \varepsilon_n^{2/3})$, we write $v = \sum_{i=1}^d v^i e_i$, where $v^i \in \mathbb{R}$, and e_i is the standard i th basis vector. Let ∂_i denote the partial derivative along the i th axis. By the Reynold's transport equation (Evans, 2010, Appendix C, Theorem 6), the partial derivative $\partial_i \tilde{\varsigma}_0^*$ of $\tilde{\varsigma}_0^*$ (F.1) at v is

$$\partial_i \tilde{\varsigma}_0^*(v) = \frac{d}{ds} \left(\int_{\tilde{\mathcal{U}}_{0, v+se_i}} \exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|u-v-se_i\|^2}{2\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du \right) \Big|_{s=0} = L_1 + L_2, \quad (\text{F.4})$$

where

$$\begin{aligned} L_1 &:= \int_{\tilde{\mathcal{U}}_{0,v}} \frac{u^i - v^i}{\varepsilon_w^2} \exp\left(-\frac{\|u-v\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du, \\ L_2 &:= \int_{\partial \tilde{\mathcal{U}}_{0,v}} \exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|u-v\|^2}{2\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) \langle e_i, \vec{n} \rangle dS. \end{aligned} \quad (\text{F.5})$$

Here, $\vec{n}$ denotes the outward pointing unit normal, while $\partial \tilde{\mathcal{U}}_{0,v}$ denote the surface of $\tilde{\mathcal{U}}_{0,v}$. Moreover, it can be seen from (F.5) that both L_1, L_2 are continuous in terms of v . Thus, we conclude that $\nabla \tilde{\varsigma}_0^*(0)$ exists. To show that

$$\nabla \tilde{\varsigma}_0^*(0) = 0, \quad (\text{F.6})$$

we make an observation from (F.1) that $\tilde{\zeta}_0^*(v) \leq \tilde{\zeta}_0^*(0)$, for any $v \in B^m(0, \varepsilon_n^{2/3})$. Hence in particular 0 is a global maximum of $\tilde{\zeta}_0^*$ along the i th axis. Therefore, $\partial_i \tilde{\zeta}_0^*(0) = 0$, and (F.6) must hold.

Continuing, we prove that

$$|D^2 \tilde{\zeta}_0^*(0)| = \mathcal{O}(\varepsilon_w^m). \quad (\text{F.7})$$

From (F.4), (F.5), and the Gauss-Green theorem (Evans, 2010, Appendix C, Theorem 1),

$$\begin{aligned} \partial_i \tilde{\zeta}_0^*(v) &= \int_{\tilde{\mathcal{U}}_{0,v}} \frac{u^i - v^i}{\varepsilon_w^2} \exp\left(-\frac{\|u - v\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du \\ &\quad + \int_{\tilde{\mathcal{U}}_{0,v}} \partial_i \left(\exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|u - v\|^2}{2\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) \right) du. \end{aligned} \quad (\text{F.8})$$

Then, with a small simplification, repeating the previous process with $v + te_j$ replacing v in (F.8), we obtain

$$\partial_j \partial_i \tilde{\zeta}_0^*(v) = \frac{d}{dt} \left(\int_{\tilde{\mathcal{U}}_{0,v+te_j}} \exp\left(-\frac{\|u - v - te_j\|^2}{2\varepsilon_w^2}\right) \partial_i \left(\exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) \right) du \right) \Big|_{t=0}. \quad (\text{F.9})$$

The existence of $\partial_j \partial_i \tilde{\zeta}_0^*(v)$ now follows from the Reynolds transport equation, and its continuous dependence on v is straightforward. It remains to calculate individual $\partial_j \partial_i \tilde{\zeta}_0^*(0)$, which is, from (F.9),

$$\partial_j \partial_i \tilde{\zeta}_0^*(0) = K_1 + K_2 + K_3 \quad (\text{F.10})$$

where

$$\begin{aligned} K_1 &:= \int_{B^m(0, \varepsilon_n^{2/3})} \frac{u^j}{\varepsilon_w^2} \frac{u^i}{\varepsilon_w^2} \exp\left(-\frac{\|u\|^2}{\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du - \int_{B^m(0, \varepsilon_n^{2/3})} \frac{\delta_{ij}}{\varepsilon_w^2} \exp\left(-\frac{\|u\|^2}{\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du \\ K_2 &:= - \int_{B^m(0, \varepsilon_n^{2/3})} \frac{u^i}{\varepsilon_w^2} \exp\left(-\frac{\|u\|^2}{\varepsilon_w^2}\right) \partial_j (\tilde{\rho}(u) J_x(u)) du \end{aligned}$$

and

$$K_3 := \frac{d}{dt} \left(\int_{\tilde{\mathcal{U}}_{0,v+te_j}} \exp\left(-\frac{\|u - v - te_j\|^2}{2\varepsilon_w^2}\right) \exp\left(-\frac{\|u\|^2}{2\varepsilon_w^2}\right) \partial_i (\tilde{\rho}(u) J_x(u)) du \right) \Big|_{t=0}.$$

Using a nearly identical argument as in Appendix D.4 to address $|D_{11} + V_1|$, particularly Lemma 33, we get

$$\begin{aligned} |K_1| &= \left| \int_{B^m(0, \varepsilon_n^{2/3})} \frac{u^i}{\varepsilon_w^2} \frac{u^j}{\varepsilon_w^2} \exp\left(-\frac{\|u\|^2}{\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du - \int_{B^m(0, \varepsilon_n^{2/3})} \frac{\delta_{ij}}{\varepsilon_w^2} \exp\left(-\frac{\|u\|^2}{\varepsilon_w^2}\right) \tilde{\rho}(u) J_x(u) du \right| \\ &= \mathcal{O}(\varepsilon_w^m). \end{aligned}$$

Another similarly straightforward argument, combined with Taylor's expansion, also leads to

$$|K_2| = \left| \int_{B^m(0, \varepsilon_n^{2/3})} \frac{u^i}{\varepsilon_w^2} \exp\left(-\frac{\|u\|^2}{\varepsilon_w^2}\right) \partial_j (\tilde{\rho}(u) J_x(u)) du \right| = \mathcal{O}(\varepsilon_w^m).$$

Lastly, we observe that the expression for K_3 mirrors that in (F.4). Therefore, by the same reasoning used to derive the boundedness of the first derivative, we can express K_3 as a sum of two terms analogous to L_1 , L_2 , and, through Taylor’s expansion, deduce that $|K_3| = \mathcal{O}(\varepsilon_w^m)$. Then putting all the obtained facts that $|K_i| = \mathcal{O}(\varepsilon_w^m)$ in (F.10), we arrive at (F.7), as wanted. ■

References

- Hanan Ayad and Mohamed Kamel. Refined shared nearest neighbors graph for combining multiple data clusterings. In *Advances in Intelligent Data Analysis V*, pages 307–318, 2003.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press, 2013.
- Dmitri Burago, Sergei Ivanov, and Yaroslav Kurylev. A graph discretization of the Laplace-Beltrami operator. *Journal of Spectral Theory*, 4(4):675–714, 2014.
- Jeff Calder. The game theoretic p-Laplacian and semi-supervised learning with few labels. *Nonlinearity*, 32(1):301–330, 2019.
- Jeff Calder and Nicolás García Trillos. Improved spectral convergence rates for graph Laplacians on ε -graphs and k -NN graphs. *Applied and Computational Harmonic Analysis*, 60:123–175, 2022.
- Jeff Calder, Nicolás García Trillos, and Marta Lewicka. Lipschitz regularity of graph Laplacians on random data clouds. *SIAM Journal on Mathematical Analysis*, 54(1):1169–1222, 2022.
- Yann Chaubet and Vincent Divol. Minimax spectral estimation of weighted Laplace operators. *arXiv preprint arXiv:2511.22694*, 2025.
- Minshuo Chen, Haoming Jiang, Wenjing Liao, and Tuo Zhao. Nonparametric regression on low-dimensional manifolds using deep relu networks: Function approximation and statistical recovery. *Information and Inference: A Journal of the IMA*, 11(4):1203–1253, 2022.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1597–1607, 2020a.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E. Hinton. Big self-supervised models are strong semi-supervised learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 22243–22255, 2020b.
- Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021.

- Manfredo P. do Carmo. *Riemannian geometry*. Birkhäuser, 1992.
- 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.
- Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- Lawrence C Evans. *Partial differential equations*, volume 19. American Mathematical Society, 2010.
- Nicolás García Trillos, Moritz Gerlach, Matthias Hein, and Dejan Slepčev. Error estimates for spectral convergence of the graph Laplacian on random geometric graphs toward the Laplace-Beltrami operator. *Foundations of Computational Mathematics*, 20(4):827–887, 2020.
- Nicolás García Trillos, Pengfei He, and Chenghui Li. Large sample spectral analysis of graph-based multi-manifold clustering. *Journal of Machine Learning Research*, 24(143):1–71, 2023.
- Nicolás García Trillos, Chenghui Li, and Raghavendra Venkatraman. Minimax rates for the estimation of eigenpairs of weighted Laplace-Beltrami operators on manifolds. *arXiv preprint arXiv:2506.00171*, 2025.
- David Gilbarg and Neil S. Trudinger. *Elliptic partial differential equations of second order*, volume 224. Springer, 1977.
- Victor Guillemin and Alan Pollack. *Differential topology*, volume 370. American Mathematical Soc., 2010.
- Jeff Z. HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. In *Advances in Neural Information Processing Systems*, volume 34, pages 5000–5011, 2021.
- Emmanuel Hebey. *Sobolev spaces on Riemannian manifolds*, volume 1635. Springer Science & Business Media, 1996.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.
- Svante Janson. Large deviations for sums of partly dependent random variables. *Random Structures & Algorithms*, 24(3):234–248, 2004.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- Chenghui Li, Rishi Sonthalia, and Nicolás García Trillos. Spectral neural networks: Approximation theory and optimization landscape. *Journal of Machine Learning*, 5(1):19–70, 2026.

- Alhassan Mumuni and Fuseini Mumuni. Data augmentation: A comprehensive survey of modern approaches. *Array*, 16:100258, 2022.
- A. Martina Neuman. Graph Laplacians on shared nearest neighbor graphs and graph Laplacians on k -nearest neighbor graphs having the same limit. *arXiv preprint arXiv:2302.12399*, 2023.
- Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. A theoretical analysis of contrastive unsupervised representation learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5628–5637, 2019.
- Connor Shorten, Taghi M. Khoshgoftaar, and Borko Furht. Text data augmentation for deep learning. *Journal of Big Data*, 8:101, 2021.
- Dirk Jan Struik. *Lectures on classical differential geometry*. Courier Corporation, 2012.
- Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research*, 22(281):1–31, 2021a.
- Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. In *Algorithmic Learning Theory*, pages 1179–1206, 2021b.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge University Press, 2018.
- Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- Caroline L Wormell and Sebastian Reich. Spectral convergence of diffusion maps: Improved error bounds and an alternative normalization. *SIAM Journal on Numerical Analysis*, 59(3):1687–1734, 2021.
- Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small ReLU networks are powerful memorizers: a tight analysis of memorization capacity. In *Advances in Neural Information Processing Systems*, volume 32, pages 15532–15543, 2019.