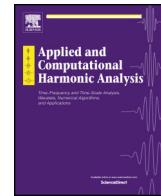




Contents lists available at ScienceDirect

Applied and Computational Harmonic Analysis

www.elsevier.com/locate/acha

The diffusion geometry of fibre bundles: Horizontal diffusion maps

Tingran Gao

Committee on Computational and Applied Mathematics, Department of Statistics, University of Chicago,
5747 S Ellis Avenue Jones 316, Chicago, IL 60637-1441, United States of America

ARTICLE INFO

Article history:

Received 3 September 2018

Received in revised form 31 March 2019

Accepted 2 August 2019

Available online xxxx

Communicated by Amit Singer

MSC:

58J65

58A30

62-07

Keywords:

Diffusion geometry

Manifold learning

Laplacian

Fibre bundles

ABSTRACT

Kernel-based nonlinear dimensionality reduction methods, such as Local Linear Embedding (LLE) and Laplacian Eigenmaps, rely heavily upon pairwise distances or similarity scores, with which one can construct and study a weighted graph associated with the data set. When each individual data object carries additional structural details, however, the correspondence relations between these structures provide extra information that can be leveraged for studying the data set using the graph. Based on this observation, we generalize *Diffusion Maps* (DM) in manifold learning and introduce the framework of *Horizontal Diffusion Maps* (HDM). We model a data set with pairwise structural correspondences as a *fibre bundle* equipped with a *connection*. We demonstrate the advantage of incorporating such additional information and study the asymptotic behavior of HDM on general fibre bundles. In a broader context, HDM reveals the sub-Riemannian structure of high-dimensional data sets, and provides a nonparametric learning framework for data sets with structural correspondences.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

Acquiring complex, massive, and often high-dimensional data sets has become a common practice in many fields of science. While inspiring and stimulating, these data sets can be challenging to analyze or understand efficiently. To gain insight despite the volume and dimension of the data, methods from a wide range of science fields have been brought into the picture, rooted in statistical inference, machine learning, signal processing, to mention just a few. Among the exploding research interests and directions in data science, the relation between the graph Laplacian [1] and the manifold Laplacian [2] has emerged as a useful guiding principle. Specifically, the field of *non-linear dimensionality reduction* has witnessed the emergence of a variety of kernel-based spectral techniques, such as Locally Linear Embedding (LLE) [3], ISOMAP [4], Hessian Eigenmaps [5], Local Tangent Space Alignment (LTSA) [6], Diffusion Maps [7], Orientable Diffusion Maps (ODM) [8], Vector Diffusion Maps (VDM) [9], and Schrödinger Eigenmaps [10]. The general practice

E-mail address: tingrangao@galton.uchicago.edu.

<https://doi.org/10.1016/j.acha.2019.08.001>

1063-5203/© 2019 Elsevier Inc. All rights reserved.

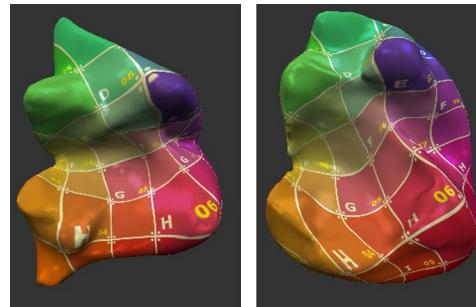


Fig. 1. An optimal structural correspondence between two lemur teeth illustrated by pushing forward a texture on the left tooth onto the right tooth. This correspondence leads to the *Continuous Procrustes Distance* [20] between shape pairs. HDM utilizes the abundant geometric information in such correspondences.

of these methods is to treat each object (images, texts, shapes, etc.) in the data set as a vertex of a graph, and two “similar” vertices are connected through an edge weighted by their similarity score. The graph is then embedded into a Euclidean space of relatively low dimensionality using the eigenvectors of the graph Laplacian (or its variant) associated with the similarity graph. Built with varying flexibility, these methods provide valuable tools for organizing complex networks and data sets by “learning” the global geometry from the local connectivity and weights.

In reality, graph-based data analysis is known to fall short of their expressiveness in capturing multiplex, heterogeneous, and time-varying pairwise relations commonly encountered in data science problems. Social network analysis has long been aware of the importance of preserving the “additional information,” such as structural, compositional, and affiliation *attributes*, for avoiding potential loss of accuracy due to the over-simplified abstraction of complex social relations into simple nodes and edges in graph models [11–14]. Recent technological advancement has also fostered an increasing trend of extending the graph-based analysis to networks of multiple types of connections, or *networks of networks* [15,16], that encode multi-modal pairwise relations as *multilayer* complex systems supported on a set of shared vertices [17–19]. These new developments essentially follow the same methodology of enriching the graph representation with structures beyond simple vertices/edges and scalar weights on them.

We propose in this paper *Horizontal Diffusion Maps* (HDM), a novel graph-based framework for analyzing complex data sets with non-scalar or functional pairwise relations, with a focus on data sets in which similarity scores between samples can be obtained from “correspondence relations” between sophisticated *individual structures* carried within each sample. We distinguish *data objects*, which constitute the vertices of the graph, from the *data points* sampled from each data object that represent the internal structure of the data object. Just like manifold learning assumes that data lie approximately on a smooth manifold, we view the data objects as approximately sampled from a smooth *base* manifold, and the data points as samples on the *fibres* of a fibre bundle over the base manifold; data points on the same data object are assumed to come from the same fibre. One such example is the biological shape data in geometric morphometrics (see Fig. 1 and Section 6), where each individual shape is a data object and each point on the shape in a data object; similar examples can be found e.g. in image analysis, where images are data objects and pixels on each image are data points. In many of these instances, the data acquired is too noisy, has huge degrees of freedom, or contains un-ordered (as opposed to sequential) features. Computing pairwise similarity between data objects typically requires optimizing some functional over the space of admissible *pairwise structural correspondences*, and the “optimal correspondence” is used to assign a distance or similarity score between the two data objects under comparison. Fig. 1 illustrates two objects from a data set of anatomical surfaces, discretized as triangular meshes; an “optimal correspondence” between the pair is a diffeomorphism between the two meshes that minimizes an energy functional whose minimum defines a distance between disk-type surfaces. Often the optimal correspondence encodes substantial information missing from the distance, which is merely a scalar condensed from the diffeomorphism. The HDM framework aims to mine this hidden

information from pairwise structural correspondences. For a data set consisting of data objects, data points, and pairwise structural correspondences, horizontal diffusion maps provide a two-level data representation that first “synchronizes” the data objects with respect to “denoised” structure correspondences by embedding the data points into a Euclidean space, and then, building on top of the first-level embedding for the data points, embed the data objects into another Euclidean space as the second level. As the second-level embedding for the data objects leverages the rich structural information at the level of data points, they are expected to be semantically more meaningful than the spectral representation obtained from standard diffusion maps which can not take advantage of the individual structural information; the synchronized spectral representation of the data points at the first level also adds to the interpretative power of HDM, enabling detailed domain-specific analysis for the data objects that is often beyond the scope of standard diffusion maps.

In the remainder of this section we relate HDM to other recent work in diffusion geometry, summarize our main theoretical contribution, and then describe the organization of the paper.

1.1. Related work

The Diffusion Map (DM) framework [7,21–24,8,9] proposes a probabilistic interpretation for graph-Laplacian-based dimensionality reduction algorithms. Under the assumption that the discrete graph is appropriately sampled from a smooth manifold, it assigns transition probabilities from a vertex to each of its neighbors (vertices connected to it) according to the edge weights, thus defining a graph random walk the continuous limit of which is a diffusion process [25,26] over the underlying manifold. The eigenvalues and eigenvectors of the graph Laplacian, which converge to those of the manifold Laplacian under appropriate assumptions [27,28], then reveal intrinsic information about the smooth manifold. More precisely, [29] proves that these eigenvectors embed the manifold into an infinite dimensional ℓ^2 space, in such a way that the ℓ^2 distance between embedded points equals to the *diffusion distance* [7] between the sample points on the manifold. Appropriate truncation of these sequences leads to an embedding of the smooth manifold into a finite dimensional Euclidean space, with small metric distortion.

Under the manifold assumption, [8,9] recently observed that estimating random walks and diffusion processes on structures associated with the original manifold (as opposed to estimates of diffusion on the manifold itself) are able to handle a wider range of tasks, or obtain improved precision or robustness for tasks considered earlier. For instance, [8] constructed a random walk on the *orientation bundle* [30, §I.7] associated with the manifold, and translated the detection of orientability into an eigenvector problem, the solution of which reveals the existence of a global section on the orientation bundle; [9] introduced a random walk on the *tangent bundle* associated with the manifold, and proposed an algorithm that embeds the manifold into an l^2 space using eigen-vector-fields instead of eigenvectors (and thus the name Vector Diffusion Maps (VDM)). Both [8] and [9] incorporate additional structures into the graph Laplacian framework: in [9] this is an extra orthogonal transformation (estimated from local tangent planes) attached to each weighted edge in the graph; in [8] the edge weights are overwritten with signs determined by this orthogonal transformation. These methods are successful, partly because they incorporate more local geometry (by estimating tangent planes) *en route* to dimensionality reduction. In [31] the VDM approach is used, analogously to [29], to embed the manifold into a finite dimensional Euclidean space. Although the VDM embedding does not reduce the dimensionality as much as standard diffusion embedding methods, it benefits from improved robustness to noise, as illustrated by the analysis of some notoriously noisy data sets [32,33].

This paper stems from the observation that it is possible to adopt the methodology of [8,9] to tackle problems in much broader contexts, where the local geometric information can be of a different type than tangent spaces. For instance, many data sets carry abundant structural details on each individual object in the data set, such as pixels in an image, vertices/faces on a triangular mesh, or a collection of persistent diagrams [34] representing a shape. Typically, kernel eigenmap methods begin by “abstracting away” these

details, encoding only pairwise similarities using a kernel function. The major advantage, like kernel methods in general, is the flexibility (no need to extract explicit features) and efficiency (most kernels are easy to compute); however, in some circumstances, the structural details may themselves be of interest. For example, in the geometry processing problem of analyzing large collections of 3D shapes, it is desirable to enable user exploration of shape variations across the collection, for which reducing each individual shape as a graph vertex completely ignores its spatial configuration. Furthermore, even when sticking to pairwise similarity scores significantly simplifies the data manipulation, the best way to score similarity (or to craft the kernel function) is not always clear. In practice, the similarity measure is often dictated by practical heuristics, which may be misguided for incompletely understood data.

Like ODM and VDM, HDM extends the diffusion map framework, but takes an essentially different path. In this paper, we are most interested in the scenario in which the sample points are themselves manifolds; the entire data set is thus modeled as a “manifold of manifolds.” To provide a mathematical model for such consideration, we first augment the manifold underlying diffusion maps, denoted as M , with extra dimensions. To each point x on M , this augmentation attaches an individual manifold, denoted as F_x ; since pairwise correspondences exist between nearby individual manifolds, we assume that around each $x \in M$ there exists an open neighborhood U such that on U the augmented structure “looks like” $U \times F$, the product of U with a “universal template” manifold F . Intuitively, M plays the role of a “parametrization” for the collection of individual manifolds $\{F_x \mid x \in M\}$. Of course, the existence of such a universal template makes sense only if the F_x ’s are compatible with each other in some appropriate sense (e.g. each F_x should at least be diffeomorphic to F); however, such compatibility is not uncommon for many data sets of interest, as we shall see in Section 2. This picture of parameterizing a family of manifolds with an underlying manifold is reminiscent of the modern differential geometric concept of a *fibre bundle*, which played an important role in the development of geometry, topology, and mathematical physics in the past century. Therefore, we shall refer to this geometric object as the underlying *fibre bundle* of the data set. Adopting the terminology from differential geometry, we call M the *base manifold*, the universal template manifold F the *fibre*, and each F_x a *fibre at x* . The fibre bundle is itself a manifold, denoted as E and referred to as the *total manifold*. We emphasize here that the fibre bundle setting we consider in this paper is even more general and flexible than the principal bundle formulation in [35], which provided a unified theoretical framework for diffusion maps and its various extensions. Whereas the principal bundle framework [35] builds upon an explicitly specified Lie group and defines the fibre bundle as a quotient space of the group action, in the framework of HDM the fibre bundles are trivialized by local parallel-transports. This flexibility allows us to analyze data sets satisfying the *fibre bundle assumption* (see Section 2) but for which the structure group can not be identified *a priori*. We shall elaborate on this in greater detail in Section 2.

A different line of research closely related to our work is the construction of adaptive *cone kernels* [36,37] in the data-driven study of dynamical systems. Unlike the geometric setting in our work (or [9,35]), the low-dimensional manifold structure lives in the phase space, and the kernels are constructed from finite differences of time-ordered data samples. In [36], the author constructed a family of nonhomogeneous and anisotropic family of kernels that assign higher affinity scores to more aligned velocity vectors; the resulting diffusion processes generate paths that asymptotically “follow along” the integral curves of the dynamical vector field. The intimate connection between the intrinsic geometry of the data and general nonhomogeneous, anisotropic kernels is characterized in great detail in [38]. The usage of these more general and flexible kernels is similar in spirit to our construction of the *coupled diffusion operator* in Section 2.2 in the specific case when the Riemannian metric on the fibre bundle splits into the direct sum of horizontal and vertical components; however, it is worth pointing out that the lack of a fibre bundle structure in [7,36,37] makes these applications of anisotropic diffusions drastically different from HDM: in our terminology, these constructions are targeted at understanding the total manifold, whereas our goal is to extract information jointly and consistently from the total manifold and the base manifold. Specifically, our definitions of *horizontal base diffusion map* (HBDM) and *horizontal base diffusion distance* (HBDD) in Section 3.2 are

meaningful only at the presence of an underlying fibre bundle structure. Most strikingly, as we point out in Remark 4.1, the HDM framework differs in an essential way from directly applying an anisotropic diffusion kernel construction to the total manifold of the fibre bundle; the two constructions coincide only in the very special case when the fibres are totally geodesically embedded into the total manifold. These subtle phenomena are characterized for the first time in the diffusion geometry literature. We thus believe that the classical differential geometric concepts of fibre bundles, Riemannian submersions, and horizontal/vertical Laplacians, though introduced into the blossoming field of geometric data analysis only for the first time, open new opportunities for gaining deeper understandings of real world data through the lens of diffusion geometry.

1.2. Main contribution

The main theoretical contribution of this paper is to provide a probabilistic interpretation of HDM as a *horizontal random walk* on the fibre bundle, extending the random walk picture of diffusion maps to a broader class of geometric objects. In one step, the transition occurs either between points on adjacent but distinct fibres, or within the same fibre. If transitions between distinct fibres depend solely on geometric proximity specified through a metric on the total manifold E , this looks no different from a direct application of diffusion maps on E . In contrast, HDM also incorporates the pairwise correspondences between individual manifolds in the fibre bundle formulation, by requiring transitions between distinct fibres to follow certain directional constraints imposed by correspondences. The resulting random walk is no longer a standard random walk on the total manifold, but rather a “horizontal lift” of a random walk on the base manifold M . Under mild assumptions, its continuous limit is a diffusion process on the total manifold E , infinitesimally generated by a *hypoelliptic differential operator* [39]. We can then map the total manifold into a Euclidean space using the eigenfunctions of this partial differential operator; discretely this corresponds to solving for the eigenvectors of *graph horizontal Laplacians*. It turns out that, by varying a couple of parameters in its construction, the family of graph horizontal Laplacians includes the discrete analogue of several important and informative partial differential operators on the fibre bundle, relating the geometry of the base manifold with that of the total manifold. Compared with [9,35], the limiting differential operators can be employed to reveal the sub-Riemannian structures of a fibre bundle (or *Riemannian submersion* [40, Chapter 9]), a task that can not be accomplished in the principal bundle framework of [9,35]. Our numerical experiments revealed intriguing geometric phenomena, such as *adiabatic limits*, when embedding the fibre bundle using eigenvectors of these new graph Laplacians; these phenomena have never been reported in any related work within the framework of [9,35].

We note that the idea of studying diffusion processes and random walks on an “augmentation” of the original data set, or extracting information from pairwise structural correspondences between sample points, has appeared elsewhere as well, in several distinct fields (e.g. shape collection analysis [41], manifold alignment [42], and neurogeometry [43]). To our knowledge, HDM is the first theoretical framework that provides the mathematical and statistical foundation for these research directions; in particular, like diffusion maps, HDM enables decoupling the probabilistic treatment of sampling from the geometry of the data set.

The rest of this paper is organized as follows: Section 2 formulates the problem and discusses the *fibre bundle assumption*; Section 3 describes the algorithmic construction; Section 4 contains the main technical results of this paper, several explicit calculations on some concrete examples of fibre bundles with totally geodesic fibres, along with a numerical example on $\text{SO}(3)$ to validate the theoretical findings; finite sampling results and applications to biological shape analysis problems will be pursued in Section 5 and Section 6, respectively; Section 7 concludes with a brief discussion and propose potentially interesting directions for future work. The differential geometry concepts essential for developing the theoretical framework, as well as technical proofs of the main results, are postponed to the appendices.

2. Horizontal diffusion maps on fibre bundles

In this section, we build the theoretical framework of horizontal diffusion maps, and relate it, where appropriate, with practical considerations for data processing.

2.1. The fibre bundle assumption

We say that the data set consists of *data objects*, and each data object contains *data points* (note that the number of data points contained in each data object may vary). Pairwise structural correspondences exist between data objects with high similarity scores; each correspondence is defined from a *source* data object (the collection of source data points) to a *target* data object (the collection of target data points), and can either be a point-to-point map or a “multi-valued map” that associates a source data point with multiple target data points. In the latter case, the correspondence may also assign similarity scores between source and target data points. To put data objects, data points, and pairwise structure correspondences in a unified geometric model, we resort to the following general definition of *fibre bundles*.

Definition 2.1 (*Fibre bundle*, [44,45]). Let $\pi : E \rightarrow M$ be a smooth map from a *total manifold* E to a *base manifold* M . We call the quadruple $\mathcal{E} = (E, M, F, \pi)$ a *fibre bundle* with *fibre manifold* F if there is an open cover $\{U_i\}$ of M with diffeomorphisms

$$\phi_i : \pi^{-1}(U_i) \longrightarrow U_i \times F$$

such that $\pi : \pi^{-1}(U_i) \rightarrow U_i$ is the composition of ϕ_i with projection onto the first factor U_i in $U_i \times F$. In other words, the following diagram is commutative:

$$\begin{array}{ccc} \pi^{-1}(U_i) & \xrightarrow{\phi_i} & U_i \times F \\ \pi \searrow & & \swarrow \text{Proj}_1 \\ & U_i & \end{array}$$

It follows immediately from this definition that $\pi^{-1}(x)$ is diffeomorphic to F for any $x \in M$. We denote F_x for $\pi^{-1}(x)$ and call it the *fibre over* $x \in M$. The diffeomorphism $\phi_i : \pi^{-1}(U_i) \rightarrow U_i \times F$ is also known as a *local trivialization* of the fibre bundle \mathcal{E} over the open set U_i . Unless otherwise stated, we assume throughout this paper that M and F are orientable Riemannian manifolds so the volume form and integration are well-defined; the dimensions of M, F will be denoted as $d = \dim(M)$, $n = \dim(F)$, respectively. Using the language of fibre bundles, our basic assumptions for the data set can be summarized as follows:

1. Data points lie approximately on a fibre bundle;
2. Data points on the same data object are sampled from the same fibre.

As stated above, the data sets of interest, to which the fibre bundle assumption applies, are those with pairwise correspondences between data objects, or fibres in the fibre bundle. This additional piece of information can now be easily incorporated into the fibre bundle framework: we interpret pairwise correspondences as *parallel-transports* along *geodesics* on the base manifold M , generated by a *connection* (see Appendix A) on the fibre bundle \mathcal{E} . For our purposes, the base manifold M plays the same role as the manifold that underlies the diffusion maps (i.e., from which data objects are drawn); additionally, we assume that each data object $x \in M$ carries a manifold structure that is diffeomorphic to a fixed fibre manifold F ; the entire data set can thus be interpreted as a collection of instantiations of the fibre F (which can be viewed as a

“template”), indexed by points on the base manifold M as $\{F_x : x \in M\}$. From a fibre bundle point of view, it is natural to study the base manifold M using the extra information in the total manifold $E = \cup_{x \in M} F_x$. In the remaining paper, unless otherwise specified, we assume all Riemannian manifolds are geodesically complete.

Roughly speaking, a data set satisfies the fibre bundle assumption if the data generation process can be viewed as first drawing fibres from the fibre bundle (equivalent to sampling on the base manifold) and then sampling on each fibre. The fibre bundle assumption admits “inconsistency” of pairwise correspondences as to the nature of the underlying geometry: though pairwise correspondences only exist (or are of high fidelity) between nearby data objects, by knitting together these correspondences along “small hops” one can still build correspondences between far-apart data objects (provided the base manifold is connected); correspondences constructed in this manner are generally inconsistent with each other in the sense that knitting together correspondences along different paths connecting the same data objects leads to different correspondences. In our framework, this inconsistency would reflect the *curvature* and *holonomy* of the *connection* on the fibre bundle; see Appendix A.

The concept of fibre bundles we chose to present above is but one of several equivalent definitions; some other popular ones can be found, for instance, in [46, Chapter 3, Chapter 10]. Our choice is based not only on the conciseness and flexibility of Definition 2.1, but also—most importantly—because there is no need to explicitly specify a *structure group*. In stark contrast is the equivalent definition of principal and associated fibre bundles, e.g. in [35, Appendix A], in which principal bundles are defined as orbit spaces of Lie group actions, and an associated bundle is obtained from a principal bundle through representations of the Lie group. The unification of all diffusion maps and variants in [35] is made possible by specifying the structure groups explicitly for each particular type of diffusion maps. Nevertheless, in most practical applications of interest to us, it is difficult to explicitly know the structure group of the fibre bundle underlying the data set. For instance, as briefly surveyed in [47, §4.1], for some data sets it may be unrealistic to model the correspondence relations between data objects as group elements; *groupoids* seem to be the more natural abstraction in those settings. Similar consideration motivated topological data analysts to propose *sheaves* as data models; see e.g. [48] and the references therein. Even in cases in which the pairwise correspondences can be modeled as group elements, the group can be too large to manipulate efficiently, such as Lie groups of diffeomorphisms or isometries commonly encountered in non-isometric collection shape analysis [49–52]. While it is not uncommon to perform *reductions* of principal bundles to reduce the structure group to smaller subgroups whenever possible, in the discrete setting this often boils down to the difficult group theoretic and combinatorial problem of understanding the rigidity or approximability of representations of discrete lattices of Lie groups [53–55]. These difficulties motivated us to take an alternative path to viewing the data sets we encountered as fibre bundles, without explicitly referring to the structure group. Fortunately, the following classical result of R. Hermann provides us with one possible route:

Theorem 1 ([56], [40, Theorem 9.3]). *Let $\pi : E \rightarrow M$ be a Riemannian submersion (cf. [40, Definition 9.8]). If E is a complete, then $\pi : E \rightarrow M$ is a fibre bundle.*

The proof of Theorem 1 is constructive. In a nutshell, Hermann explicitly constructed local trivializations around each $x \in M$, by connecting points on the fibre $\pi^{-1}(x)$ to points on any neighboring fibre $\pi^{-1}(y)$ by horizontally lifting the geodesic on M that connects x to y . Here the horizontal lifting is made possible by the Riemannian structure on E , which canonically splits the tangent bundle of E into the direct sum of a horizontal and vertical subbundles. As pointed out in [40, §9.E], the horizontal subbundle is an *Ehresmann connection* (see Appendix A) on the fibre bundle. The structure group of the fibre bundle can then be determined from the *holonomy* of the Ehresmann connection; see [40, §9.47] for more details. Obviously, the data required in Theorem 1 to fully specify the fibre bundle structure can be provided in a slightly different order: if we are given a Riemannian manifold M and another manifold E but without a prescribed

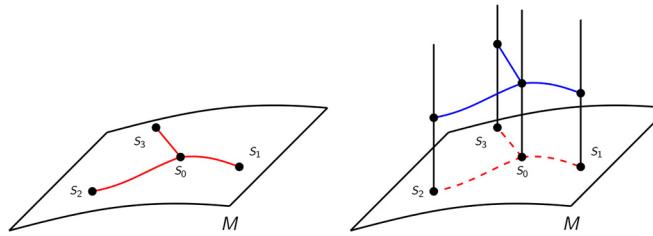


Fig. 2. Left: A random walk on the base manifold M jumps in one step from a point s_0 on M to one of its neighboring points s_1, s_2, s_3 . Right: The horizontal lift of the same random walk from M to the fibre bundle \mathcal{E} , which jumps in one step from a point on fibre F_{s_0} to a neighboring point on the fibre F_{s_1}, F_{s_2} , or F_{s_3} .

Riemannian structure, and $\pi : E \rightarrow M$ is a smooth submersion with an Ehresmann connection on E , then we can define a product Riemannian structure on E which imposes the orthogonality between the horizontal and vertical subbundles of the tangent bundle TE . It is straightforward to verify that $\pi : E \rightarrow M$ is a Riemannian submersion with such a Riemannian structure on E . In other words, a fibre bundle can be defined equivalently by a smooth submersion between the total and base manifold (with appropriate completeness assumptions), a Riemannian structure on the base manifold, and an Ehresmann connection. We close the discussion in this section by emphasizing that, though it might appear that our fibre bundle framework “discards” the notion of structure groups compared with the fibre bundle formulation pioneered in [9,35], structure groups indeed are specified, just in an indirect manner.

2.2. Horizontal random walks and diffusion processes on fibre bundles

Equipped with the geometric notion of fibre bundles, we are now ready to define a random walk tailored to a data set with pairwise correspondences. Starting from a point $e \in E$, in one step a random walker is allowed to jump to a neighboring $e' \in E$ only $\pi(e') \neq \pi(e)$ and e, e' can be joined by a horizontally lifted image of a piecewise geodesic connecting $\pi(e)$ to $\pi'(e)$ on M . More specifically, just as a standard random walk on M jumps from $x \in M$ to a neighbor $y \in M$ following a transition probability $\mathbb{P}(y | x)$, a horizontal random walk jumps from $e \in F_x \subset E$ to $P_{yx}(e) \in F_y \subset E$ with transition probability $\mathbb{P}(y | x)$; note in particular that this transition probability depends only on the projections $x = \pi(e)$ and $y = \pi(P_{yx}(e))$. In this sense, a horizontal random walk on the fibre bundle \mathcal{E} can be viewed as “driven” by an underlying random walk on the base manifold M (see Fig. 2 for an illustration). Passing to the continuous limit (in the weak sense as the random walk step size approaches zero, see [57] and Section 5.1), both random walks on the fibre bundle and the base manifold converge to diffusion processes. For the convenience of exposition, hereafter we refer to the limit diffusion process on the fibre bundle as the *horizontal lift* of the limit diffusion process on the base manifold. In the Riemannian setting, this construction is reminiscent of the notion of *stochastic parallel transport* [58,59] in stochastic differential geometry.

The following is a precise description of the horizontal diffusion processes on the fibre bundle in the language of symmetric Markov semigroups. For clarity, let us assume M and F are both orientable. Let kernel $K : \mathbb{R} \rightarrow \mathbb{R}^{\geq 0}$ be a smooth function compactly supported on the unit interval $[0, 1]$. For bandwidth parameter $\epsilon > 0$ and any pairs of $x, y \in M$, define

$$K_\epsilon(x, y) = K\left(\frac{d_M^2(x, y)}{\epsilon}\right),$$

where $d_M(\cdot, \cdot)$ stands for the geodesic distance on M . Note that $K_\epsilon(\cdot, \cdot)$ is non-zero only if x, y are sufficiently close to each other under the Riemannian metric on M , due to the compactness of the kernel function K . For any $f \in C^\infty(E)$, define the diffusion operator $H_\epsilon : C^\infty(E) \rightarrow C^\infty(E)$ as

$$H_\epsilon f(x, v) = \int_M K_\epsilon(x, y) f(y, P_{yx}v) d\text{vol}_M(y), \quad \forall x \in M, v \in F_x,$$

where $d\text{vol}_M$ stands for the Riemannian volume element on M . Intuitively, at each point $(x, v) \in E$, H_ϵ averages the value of f on a neighborhood around (x, v) expanded by parallel-transporting (x, v) along geodesics on M . Compared with the diffusion semigroup characterization of a diffusion process on the base manifold, H_ϵ incorporates the extra information provided by the connection.

Variants of H_ϵ that involve the sampling density can be similarly constructed, which is useful since in practice it is difficult to uniformly sample from M . Consider a density function $p \in C^\infty(M)$ with respect to which the samples are generated. For simplicity, assume p is bounded away from zero from below, i.e.,

$$\int_M p(y) d\text{vol}_M(y) = 1 \quad \text{and} \quad p(y) \geq p_0 > 0 \quad \forall y \in M. \quad (2.1)$$

Similar to the construction of diffusion maps [7], we can decouple the geometry of the manifold from the influence of sampling density by normalizing the integral kernel. To this end, we set

$$p_\epsilon(x) = \int_M K_\epsilon(x, y) p(y) d\text{vol}_M(y)$$

and denote for any *normalization parameter* $\alpha \in [0, 1]$

$$K_\epsilon^{(\alpha)}(x, y) = \frac{K_\epsilon(x, y)}{p_\epsilon^\alpha(x) p_\epsilon^\alpha(y)},$$

then define the *horizontal diffusion operator*

$$H_\epsilon^{(\alpha)} f(x, v) = \frac{\int_M K_\epsilon^{(\alpha)}(x, y) f(y, P_{yx}v) p(y) d\text{vol}_M(y)}{\int_M K_\epsilon^{(\alpha)}(x, y) p(y) d\text{vol}_M(y)}, \quad \forall x \in M, v \in F_x \quad (2.2)$$

for any $f \in C^\infty(E)$. As we shall see later, the infinitesimal generator of $H_\epsilon^{(\alpha)}$ is a second order partial differential operator in which all derivatives, as vector fields on E , are horizontal.

A different practical consideration is that pairwise correspondences can be relaxed from maps to couplings of probability measures when fibres are discretized. Examples for such relaxed pairwise correspondences include the *soft-assign Procrustes matching* [60] in medical imaging, the *soft maps* [61] in geometry processing, the *transport plans* [62,63] in optimal transportation, to name just a few. In the HDM framework, these relaxed correspondences also define diffusion processes on the fibre bundle, now consisting of two ingredients: a horizontal lift of a diffusion process on the base manifold, composed with another diffusion process within the fibre. In this setting, it is an interesting question to “learn” the connection from the composition of two diffusion processes; practically, this amounts to “recovering” maps from couplings in a collection of data objects. In some applications (see e.g. Section 6), one can also “learn” the structure of the template fibre from the connection. Making an analogy with the terminology *manifold learning*, we call this type of learning problems *fibre learning*. Similar to (2.2), we can write the diffusion process considered in fibre learning in the language of Markov semigroups. Let $K : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$ be a smooth bi-variate function compactly supported on the unit square $[0, 1] \times [0, 1]$, and let $\epsilon > 0, \delta > 0$ be *bandwidth parameters*. Define

$$K_{\epsilon,\delta}(x, v; y, w) = K \left(\frac{d_M^2(x, y)}{\epsilon}, \frac{d_{F_y}^2(P_{yx}v, w)}{\delta} \right) \quad (2.3)$$

for $(x, v) \in E$, $(y, w) \in E$, where $d_M(\cdot, \cdot)$, $d_{F_y}(\cdot, \cdot)$ are the geodesic distances on M , F_y respectively. Assume $p \in C^\infty(E)$ is a density function bounded away from zero from below, i.e.

$$\iint_{M F_y} p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y) = 1 \quad (2.4)$$

and

$$p(y, w) \geq p_0 > 0, \quad \forall (y, w) \in E. \quad (2.5)$$

For $\alpha \in [0, 1]$, if we set

$$p_{\epsilon,\delta}(x, v) = \iint_{M F_y} K_{\epsilon,\delta}(x, v; y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y),$$

and

$$K_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) = \frac{K_{\epsilon,\delta}(x, v; y, w)}{p_{\epsilon,\delta}^\alpha(x, v) p_{\epsilon,\delta}^\alpha(y, w)},$$

then the *coupled diffusion operator* for all $(x, v) \in E$ can be written as

$$H_{\epsilon,\delta}^{(\alpha)} f(x, v) = \frac{\iint_{M F_y} K_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) f(y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y)}{\iint_{M F_y} K_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y)}. \quad (2.6)$$

The infinitesimal generator of $H_{\epsilon,\delta}^{(\alpha)}$ has to be considered differently from that of $H_\epsilon^{(\alpha)}$ due to the appearance of two (instead of one) bandwidth parameters ϵ, δ . It turns out that the relative rate with which ϵ and δ approach 0 affects the type of the infinitesimal generator associated with the diffusion process, see Section 4.

3. The HDM algorithm

In this section, we describe the manifold learning framework of HDM that extracts feature information in a data set with pairwise similarity and structural correspondences, based on the geometric intuition explained in Section 2. The construction of graph horizontal Laplacians and spectral embeddings apply to any fibred graph and symmetric similarity measure satisfying the structural assumptions in this section; the theoretical results to be presented in Section 4 and Section 5 apply to the concrete scenario where the graph arises from sampling the fibre bundle as an embedded submanifold of an ambient Euclidean space and the similarity measure encodes the connection information (see Section 5.1 for more details).

3.1. Graph horizontal Laplacians

The data set considered in the HDM framework is a triplet (\mathcal{X}, ρ, G) , where

(1) The *total* data set \mathcal{X} can be partitioned into a collection of *data objects* X_1, \dots, X_n

$$\mathcal{X} = \bigcup_{j=1}^n X_j, \quad X_j \cap X_k = \emptyset \text{ for all } 1 \leq j \neq k \leq n,$$

where each data object X_j is referred to as the j -th *fibre* of \mathcal{X} , which contains κ_j *data points*

$$X_j = \{x_{j,1}, x_{j,2}, \dots, x_{j,\kappa_j}\}.$$

We call the collection of fibres the *base* data set

$$\mathcal{B} = \{X_1, X_2, \dots, X_n\},$$

and let $\pi : \mathcal{X} \rightarrow \mathcal{B}$ be the *canonical projection* from \mathcal{X} to \mathcal{B}

$$\begin{aligned} \pi : \mathcal{X} &\longrightarrow \mathcal{B} \\ x_{j,k} &\longmapsto X_j, \quad 1 \leq j \leq n, 1 \leq k \leq \kappa_j. \end{aligned}$$

Denote the total number of points in \mathcal{X} as

$$\kappa = \kappa_1 + \kappa_2 + \dots + \kappa_n.$$

(2) The *mutual similarity measure* $\rho : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{\geq 0}$ is a symmetric non-negative function that vanishes on each fibre, i.e.

$$\rho(\xi, \eta) \geq 0, \quad \rho(\xi, \eta) = \rho(\eta, \xi) \quad \forall \xi, \eta \in \mathcal{X}$$

and

$$\rho(\xi, \eta) = 0 \quad \text{if } \xi, \eta \in X_j \text{ for some } 1 \leq j \leq n.$$

For simplicity of notation, we denote the restriction of ρ on $X_i \times X_j$ as

$$\rho_{ij}(s, t) := \rho(x_{i,s}, x_{j,t}) \quad \forall x_{i,s} \in X_i, x_{j,t} \in X_j.$$

In words, ρ_{ij} is an $\kappa_i \times \kappa_j$ matrix on \mathbb{R} , to which we will refer as the *mutual similarity matrix* between X_i and X_j . Note that $\rho_{ij} = 0$ if $i = j$.

(3) The *affinity graph* $G = (V, E)$ has κ vertices, with each $v_{i,s}$ corresponding to a point $x_{i,s} \in \mathcal{X}$. Without loss of generality, assume G is connected. (In our applications, each $x_{i,s}$ is typically connected to several $x_{j,t}$'s on neighboring fibres.) If there is an edge between $v_{i,s}$ and $v_{j,t}$ in G , then $x_{i,s}$ is a *neighbor* of $x_{j,t}$ (and $x_{j,t}$ is a neighbor of $x_{i,s}$); X_i is called a *neighbor* of X_j (and similarly X_j a neighbor of X_i) if there is an edge in G linking one point in X_i with one point in X_j . Implicitly, these define a graph $G_B = (V_B, E_B)$ in which vertices of V_B are in one-to-one correspondences with fibres of \mathcal{X} , and E_B encodes the neighborhood relations between pairs of fibres. G_B will be called as the *base affinity graph*.

With the triplet (\mathcal{X}, ρ, G) specified, we detail below the construction of the *graph horizontal Laplacian*. Let $W \in \mathbb{R}^{\kappa \times \kappa}$ be the *weighted adjacency matrix* of the graph G , i.e., W is a block matrix in which the (i, j) -th block is ρ_{ij} . The (s, t) entry in W_{ij} stands for the edge weight $\rho_{ij}(s, t)$ between $v_{i,s}$ and $v_{j,t}$. Since $\rho_{ij} = \rho_{ji}^\top$, W is a symmetric matrix. Let D be the $\kappa \times \kappa$ diagonal matrix in which the j -th diagonal entry

equals to the j -th row sum of W . We define the *graph horizontal Laplacian* for the triplet (\mathcal{X}, ρ, G) as the weighted graph Laplacian of G with edge weights W , i.e.

$$L^H := D - W. \quad (3.1)$$

Since G is connected, the diagonal elements of D are all non-zero. Thus D is invertible and we can define the *random-walk* and *normalized* version of L^H :

$$L_{\text{rw}}^H := D^{-1}L^H = I - D^{-1}W, \quad (3.2)$$

$$L_*^H := D^{-1/2}L^H D^{-1/2} = I - D^{-1/2}WD^{-1/2}. \quad (3.3)$$

Following [7], we can also repeat these constructions on a renormalized graph of G by setting for some $\alpha \in [0, 1]$

$$W_\alpha := D^{-\alpha}WD^{-\alpha} \quad (3.4)$$

and constructing the graph horizontal Laplacians from W_α instead of W . More precisely, let D_α be the $\kappa \times \kappa$ diagonal matrix in which the j -th diagonal entry equals to the j -th row sum of W_α , and set

$$L_\alpha^H := D_\alpha - W_\alpha, \quad (3.5)$$

$$L_{\alpha, \text{rw}}^H := D_\alpha^{-1}L_\alpha^H = I - D_\alpha^{-1}W_\alpha, \quad (3.6)$$

$$L_{\alpha,*}^H := D_\alpha^{-1/2}L_\alpha^H D_\alpha^{-1/2} = I - D_\alpha^{-1/2}W_\alpha D_\alpha^{-1/2}. \quad (3.7)$$

Remark 3.1. The block structure in the matrix W is reminiscent of the *graph connection Laplacian* [9, Section 3], but the constraints on the blocks are different: blocks of the graph connection Laplacian are built from matrix representations of a Lie group, but blocks of the graph horizontal Laplacian represent similarity between data objects and are matrices with non-negative entries. The normalization we apply to W is the same as for standard diffusion maps [7]. Formally, the constructions of L_α^H , $L_{\alpha, \text{rw}}^H$, and $L_{\alpha,*}^H$, as well as the embeddings derived from their eigen-decompositions, appears identical to their counterparts in standard diffusion maps, but we will show below that the unique fibred structure of the graph G allows us to characterize more subtle geometry in (\mathcal{X}, ρ, G) than standard diffusion maps could (see Remark 4.1 and Remark 5.1).

3.2. Spectral distances and embeddings

Spectral distances are defined via the eigen-decompositions of graph Laplacians. Since $L_{\alpha, \text{rw}}^H$ differs from $L_{\alpha,*}^H$ only by a similarity transformation

$$L_{\alpha,*}^H = D_\alpha^{1/2}L_{\alpha, \text{rw}}^H D_\alpha^{-1/2},$$

the two Laplacians have essentially the same eigen-decomposition. We shall focus on $L_{\alpha,*}^H$ for the rest of this section due to its computational advantage as a real symmetric matrix.

Any right eigenvector $v \in \mathbb{R}^\kappa$ of $L_{\alpha,*}^H$ defines a function on the vertices of G . By the construction of $L_{\alpha,*}^H$, the length- κ vector v , when written as the concatenation of n segments of length $\kappa_1, \dots, \kappa_n$ respectively, defines a function on each of the n fibres X_1, \dots, X_n . We assume eigenvectors are always column vectors, and write

$$v = \left(v_{[1]}^\top, \dots, v_{[n]}^\top \right)^\top$$

where each column vector $v_{[j]} \in \mathbb{R}^{\kappa_j}$ defines a function on the fibre X_j . Now let

$$\lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{\kappa-1}$$

be the κ eigenvalues of $L_{\alpha,*}^H$ in ascending order, and denote the eigenvector corresponding to eigenvalue λ_j as v_j . By the connectivity assumption for G , we know from spectral graph theory [1] that $\lambda_0 = 0$, $\lambda_0 < \lambda_1$, and v_0 is a constant multiple of the column vector with all entries equal to 1; we have thus

$$0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{\kappa-1}.$$

By the spectral decomposition of $L_{\alpha,*}^H$,

$$L_{\alpha,*}^H = \sum_{l=0}^{\kappa-1} \lambda_l v_l v_l^\top, \quad (3.8)$$

and for any fixed *diffusion time* $t \in \mathbb{R}^{>0}$,

$$(L_{\alpha,*}^H)^t = \sum_{l=0}^{\kappa-1} \lambda_l^t v_l v_l^\top, \quad (3.9)$$

with the (i,j) -th block

$$\left((L_{\alpha,*}^H)^t \right)_{ij} = \sum_{l=0}^{\kappa-1} \lambda_l^t v_{l[i]} v_{l[j]}^\top. \quad (3.10)$$

In general, this block is not a square matrix. Its Frobenius norm can be computed as

$$\begin{aligned} \left\| \left((L_{\alpha,*}^H)^t \right)_{ij} \right\|_F^2 &= \text{Tr} \left[\left((L_{\alpha,*}^H)^t \right)_{ij} \left((L_{\alpha,*}^H)^t \right)_{ij}^\top \right] = \text{Tr} \left[\sum_{l,m=0}^{\kappa-1} \lambda_l^t \lambda_m^t v_{l[i]} v_{l[j]}^\top v_{m[j]} v_{m[i]}^\top \right] \\ &= \text{Tr} \left[\sum_{l,m=0}^{\kappa-1} \lambda_l^t \lambda_m^t v_{m[i]}^\top v_{l[i]} v_{l[j]}^\top v_{m[j]} \right] = \sum_{l,m=0}^{\kappa-1} \lambda_l^t \lambda_m^t v_{m[i]}^\top v_{l[i]} v_{l[j]}^\top v_{m[j]}. \end{aligned} \quad (3.11)$$

Define the *horizontal base diffusion map* (HBDM) as

$$\begin{aligned} V^t : \mathcal{B} &\longrightarrow \mathbb{R}^{\kappa^2} \\ X_j &\longmapsto \left(\lambda_l^{t/2} \lambda_m^{t/2} v_{l[j]}^\top v_{m[j]} \right)_{0 \leq l,m \leq \kappa-1} \end{aligned} \quad (3.12)$$

with which

$$\left\| \left((L_{\alpha,*}^H)^t \right)_{ij} \right\|_F^2 = \langle V^t(X_i), V^t(X_j) \rangle, \quad (3.13)$$

where $\langle \cdot, \cdot \rangle$ is the standard Euclidean inner product on \mathbb{R}^{κ^2} . Furthermore, we define the *horizontal base diffusion distance* (HBDD) on \mathcal{B} as

$$\begin{aligned} d_{\text{HBDM},t}(X_i, X_j) &= \|V^t(X_i) - V^t(X_j)\| \\ &= \{ \langle V^t(X_i), V^t(X_i) \rangle + \langle V^t(X_j), V^t(X_j) \rangle - 2 \langle V^t(X_i), V^t(X_j) \rangle \}^{\frac{1}{2}}. \end{aligned} \quad (3.14)$$

From a learning point of view, the map $V^t : \mathcal{B} \rightarrow \mathbb{R}^{\kappa^2}$ is equivalent to the unsupervised features learned from the data set with structural correspondences. Note also that HDM embeds the base data set \mathcal{B} into a Euclidean space of dimension κ^2 , which is of much higher dimensionality than the size of the original data set. In practice, however, one often truncates the spectrum of graph Laplacians, thus embedding the data set into a Euclidean of reduced dimensionality. In our numerical experiments and applications (see Section 6), we found it is usually sufficient to retain the first $O(\sqrt{\kappa})$ to $O(\kappa)$ eigenvalues. Even though this truncation still involves higher spatial complexity than diffusion maps, our results show that HDM significantly outperforms DM for our purposes; we thus believe that the high-dimensional embedding is a modest price to pay for extracting the hidden information in the structural correspondences.

In addition to handling the base data set \mathcal{B} , HDM is also capable of embedding the total data set \mathcal{X} into Euclidean spaces. Define for each *diffusion time* $t \in \mathbb{R}^+$ the *horizontal diffusion map* (HDM)

$$\begin{aligned} H^t : \mathcal{X} &\longrightarrow \mathbb{R}^{\kappa-1} \\ x_{j,s} &\longmapsto (\lambda_1^t v_{1[j]}(s), \lambda_2^t v_{2[j]}(s), \dots, \lambda_{\kappa-1}^t v_{(\kappa-1)[j]}(s)), \end{aligned} \tag{3.15}$$

where $v_{l[j]}(s)$ is the s -th entry of the j -th segment of the l -th eigenvector, with $j = 1, \dots, n, s = 1, \dots, \kappa_j$. We could also have written

$$v_{l[j]}(s) = v_l(s_j + s), \quad \text{where } s_1 = 0 \text{ and } s_j = \sum_{p=1}^{j-1} \kappa_p \text{ for } j \geq 2.$$

Following a similar argument as in [7], we can define the *horizontal diffusion distance* (HDD) on \mathcal{X} as

$$d_{\text{HDM},t}(x_{i,s}, x_{j,t}) = \|H^t(x_{i,s}) - H^t(x_{j,t})\|. \tag{3.16}$$

As it stands, H^t embeds the total data set \mathcal{X} into a Euclidean space preserving the horizontal diffusion distance on \mathcal{X} . Moreover, this embedding automatically suggests a global registration for all fibres that respects the mutual similarity measure ρ ; similar ideas were already implicit in [41]. For simplicity of notation, let us write

$$H_j^t := H^t \upharpoonright X_j$$

for the restriction of H^t to fibre X_j , and call this the j -th *component* of H^t . Up to scaling, the components of H^t bring the fibres of \mathcal{X} to a common “template”, such that points $x_{i,s}$ and $x_{j,t}$ with a high similarity measure $\rho_{ij}(s,t)$ tend to be close to each other in the embedded Euclidean space. One can then reconstruct pairwise structural correspondences between fibres X_i, X_j in the embedded Euclidean space, now between the embedded point clouds in \mathbb{R}^{κ^2} . With appropriate truncation of the spectrum of the graph horizontal Laplacian, these reconstructed structural correspondences are the “denoised version” of the original correspondences. Moreover, recalling that each X_j is sampled from some manifold F_j , one can often estimate a *template fibre* $F \subset \mathbb{R}^m$ from the embedded images

$$H_1^t(X_1), \dots, H_n^t(X_n),$$

and extend (by interpolation) H_j^t from a discrete correspondence to a continuous bijective map from F_j to F , then build correspondence maps between an arbitrary pair X_i, X_j by composing (the interpolated continuous maps) H_i^t with $(H_j^t)^{-1}$. Pairwise correspondences reconstructed in this manner are globally consistent, since they all go through the common template manifold F . We discuss in greater detail an application of HDM and HDD to a data set of shapes in geometric morphometrics in Section 6.

4. Infinitesimal generators for horizontal and coupled diffusion operators

We are now ready to present the main technical results of this paper. First, we characterize the infinitesimal generator of the horizontal diffusion operator $H_\epsilon^{(\alpha)}$ in (2.2).

Theorem 2. *Suppose $\mathcal{E} = (E, M, F, \pi)$ is a fibre bundle, M is a smooth Riemannian manifold without boundary, and E is equipped with the Riemannian metric (B.6). For any $f \in C^\infty(E)$ and $(x, v) \in E$,*

$$\lim_{\epsilon \rightarrow 0} \frac{H_\epsilon^{(\alpha)} f(x, v) - f(x, v)}{\epsilon} = \frac{m_2}{2m_0} \frac{[\Delta_H(f\bar{p}^{1-\alpha}) - f\Delta_H\bar{p}^{1-\alpha}](x, v)}{\bar{p}^{1-\alpha}(x)}, \quad (4.1)$$

where m_0, m_2 are positive constants depending only on the base manifold M and the kernel K , Δ_H is the rough horizontal Laplacian on E in (B.8), and $\bar{p} = p \circ \pi \in C^\infty(E)$.

The proof of Theorem 2 can be found in Appendix C.

Corollary 3. *Under the same assumptions as in Theorem 2, when $\alpha = 1$,*

$$\lim_{\epsilon \rightarrow 0} \frac{H_\epsilon^{(1)} f(x, v) - f(x, v)}{\epsilon} = \frac{m_2}{2m_0} \Delta_H f(x, v). \quad (4.2)$$

Characterizing the infinitesimal generator of the coupled diffusion operator $H_{\epsilon, \delta}^{(\alpha)}$ is slightly more subtle: the generator of the diffusion process depends on the relative speed at which the two bandwidth parameters ϵ, δ approach 0. For clarity, we first state the result for the case when the ratio δ/ϵ remains bounded as $\epsilon, \delta \rightarrow 0$.

Theorem 4 (Bounded ratio δ/ϵ). *Suppose $\mathcal{E} = (E, M, F, \pi)$ is a fibre bundle, M is a smooth Riemannian manifold without boundary, and E is equipped with the Riemannian metric (B.6). For any $f \in C^\infty(E)$ and $(x, v) \in E$, if $\delta = O(\epsilon)$ as $\epsilon \rightarrow 0$, then*

$$\begin{aligned} H_{\epsilon, \delta}^{(\alpha)} f(x, v) &= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \frac{[\Delta_H(f\bar{p}^{1-\alpha}) - f\Delta_H\bar{p}^{1-\alpha}](x, v)}{\bar{p}^{1-\alpha}(x, v)} \\ &\quad + \delta \frac{m_{22}}{2m_0} \frac{[\Delta_E^V(f\bar{p}^{1-\alpha}) - f\Delta_E^V\bar{p}^{1-\alpha}](x, v)}{\bar{p}^{1-\alpha}(x, v)} + O(\epsilon^2 + \epsilon\delta + \delta^2), \end{aligned} \quad (4.3)$$

where m_0, m_{21}, m_{22} are positive constants depending only on the total manifold E and the kernel K , Δ_H is the rough horizontal Laplacian on E defined in (B.8), and Δ_E^V is the vertical Laplacian of the fibre bundle \mathcal{E} defined in (B.11).

For a proof of Theorem 4, see Appendix C. Note that in the distribution sense Theorem 2 can be interpreted as a special case of Theorem 4 when $\delta = o(\epsilon)$ as $\epsilon \rightarrow 0$. From a different point of view, Theorem 4 can also be interpreted as [7, Theorem 2] applied on a fibre bundle (E, M, F, π) with a family of varying Riemannian metrics

$$g_{\delta/\epsilon}^E = g^M \oplus \frac{\delta}{\epsilon} g^F,$$

which is known as the *canonical variation* in the literature of Riemannian submersion [40, §9.G] [64, §2.7.5]. If $\delta/\epsilon \rightarrow 0$, then the rescaled metric

$$\frac{\epsilon}{\delta} g_{\delta/\epsilon}^E = \frac{\epsilon}{\delta} g^M \oplus g^F$$

is said to approach its *adiabatic limit*, or taking adiabatic limits amounts to blowing up or contracting the fibres, which is very useful in studying foliations. In the horizontal diffusion maps framework, the adiabatic limits can be indirectly taken by adjusting the relative magnitudes of the horizontal and vertical bandwidth parameters; see Fig. 3 for an illustration. An in-depth discussion of adiabatic limits is beyond the scope of this paper, and we refer interested readers to [65,66] and references therein.

Corollary 5. *Under the same assumptions as in Theorem 4, if the limit of the ratio δ/ϵ exists and is finite, i.e.,*

$$\beta := \lim_{\epsilon \rightarrow 0} \delta/\epsilon < \infty,$$

then

$$\lim_{\epsilon \rightarrow 0} \frac{H_{\epsilon,\delta}^{(\alpha)} f(x, v) - f(x, v)}{\epsilon} = \frac{1}{2} \frac{[L_\beta(f p^{1-\alpha}) - f L_\beta p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} \quad (4.4)$$

where L_β is a second order partial differential operator on E given by

$$L_\beta = \frac{m_{21}}{m_0} \Delta_H + \beta \frac{m_{22}}{m_0} \Delta_E^V. \quad (4.5)$$

In particular, if $m_{21} = \beta m_{22}$ and $\pi : E \rightarrow M$ is a harmonic map, then $L = c \Delta_E$ where Δ_E is the Laplace-Beltrami operator on E and c a multiplicative constant. In addition, if $\alpha = 1$, then

$$\lim_{\epsilon \rightarrow 0} \frac{H_{\epsilon,\delta}^{(1)} f(x, v) - f(x, v)}{\epsilon} = \frac{c}{2} \Delta_E f(x, v).$$

Proof. If $\pi : E \rightarrow M$ is a harmonic map, the fibres of π are minimal submanifolds of E (*vice versa*, see e.g. [64, Lemma 2.2.4]) and $\Delta_H = \Delta_E^H$ (see Remark 7.2). \square

Remark 4.1. Corollary 5 clearly indicates that the coupled diffusion operator $H_{\epsilon,\delta}^{(\alpha)}$ differs from the anisotropic diffusion operators considered in [21,7] and the dynamical system literature [36,37] in an essential way: in general, when the fibres are not totally geodesic submanifolds of the fibre bundle, the infinitesimal generator (4.5) will never equal to the Laplace-Beltrami operator of the total manifold, regardless of the relative ratio between δ and ϵ — even when the two constants in front of Δ_H and Δ_E^V coincide. This is essentially due to the difference between the rough horizontal Laplacian Δ_H and the *bona fide* “horizontal Laplacian” Δ_E^H commonly encountered in sub-Riemannian geometry and Riemannian submersions; see Remark 7.2 for more details. The HDM framework is thus by no means a straightforward application of the anisotropic diffusion maps to the total manifold of the fibre bundle.

In order to state the result for the case when the ratio δ/ϵ is not asymptotically bounded as $\epsilon \rightarrow 0$, let us define the *fibre average* of any function $f \in C^\infty(E)$ as

$$\langle f \rangle(x) = \int_{F_x} f(x, v) \, d\text{vol}_{F_x}(v) \quad (4.6)$$

whenever the integral converges. If $\langle f \rangle(x)$ exists for all $x \in M$ (e.g. when the fibre is compact or f is integrable), obviously $\langle f \rangle \in C^\infty(M)$.

Consider now the probability density function p in the definition of $H_{\epsilon,\delta}^{(\alpha)}$. The fibre average $\langle p \rangle$ is a probability density function on M , since

$$\int_M \langle p \rangle(x) \, d\text{vol}_M(x) = \int_M \int_{F_x} p(x, v) \, d\text{vol}_{F_x}(v) \, d\text{vol}_M(x) = 1.$$

Note that $\langle p \rangle$ is bounded away from 0 from below according to our assumption (2.5). We can thus divide p by $\langle p \rangle$ and define the *conditional probability density function* on E as

$$p(v | x) := \frac{p(x, v)}{\langle p \rangle(x)}. \quad (4.7)$$

The name comes from the observation that $p(v | x)$ defines a probability density function when restricted to a single fibre:

$$\int_{F_x} p(v | x) \, d\text{vol}_{F_x}(v) = \frac{\int_{F_x} p(x, v) \, d\text{vol}_{F_x}(v)}{\langle p \rangle(x)} = 1.$$

The last piece of notation we need for Theorem 6 is

$$\langle f \rangle_p(x) := \int_{F_x} f(x, v) p(v | x) \, d\text{vol}_{F_x}(v), \quad (4.8)$$

for any function $f \in C^\infty(E)$. We shall refer to $\langle f \rangle_p$ as the *fibre average of f with respect to the probability density function p* .

Theorem 6 (Unbounded ratio δ/ϵ). *Suppose $\mathcal{E} = (E, M, F, \pi)$ is a fibre bundle, M is a smooth Riemannian manifold without boundary, and E is equipped with the Riemannian metric (B.6). Define $\gamma := \delta/\epsilon$ (equivalently $\delta = \gamma\epsilon$). For any $f \in C^\infty(E)$ and $(x, v) \in E$, as $\epsilon \rightarrow 0$,*

$$\begin{aligned} & \lim_{\gamma \rightarrow \infty} H_{\epsilon, \gamma\epsilon}^{(\alpha)} f(x, v) \\ &= \langle f \rangle_p(x) + \epsilon \frac{m'_2}{2m'_0} \frac{[\Delta_M(\langle f \rangle_p \langle p \rangle^{1-\alpha}) - \langle f \rangle_p \Delta_M \langle p \rangle^{1-\alpha}](x)}{\langle p \rangle^{1-\alpha}(x)} + O(\epsilon^2), \end{aligned} \quad (4.9)$$

where m'_0, m'_2 are positive constants depending only on the base manifold M and the kernel K , Δ_M is the Laplace-Beltrami operator on M , $\langle p \rangle$ is the fibre average of the probability density function p , and $\langle f \rangle_p$ is the fibre average of f with respect to the density p . In particular, if $\alpha = 1$, then

$$\lim_{\gamma \rightarrow \infty} H_{\epsilon, \gamma\epsilon}^{(1)} = \langle f \rangle_p(x) + \epsilon \frac{m'_2}{2m'_0} \Delta_M \langle f \rangle_p + O(\epsilon^2).$$

The proof of Theorem 6 can be found in Appendix C. Intuitively, Theorem 6 states that if the vertical bandwidth parameter $\delta \rightarrow \infty$ then the coupled diffusion operator contains little information about the fibres. Comparing Theorem 6 with Theorem 4, one can see that in general

$$\lim_{\epsilon \rightarrow 0} \lim_{\gamma \rightarrow \infty} \frac{H_{\epsilon, \gamma\epsilon}^{(\alpha)} f(x, v) - f(x, v)}{\epsilon} \neq \lim_{\gamma \rightarrow \infty} \lim_{\epsilon \rightarrow 0} \frac{H_{\epsilon, \gamma\epsilon}^{(\alpha)} f(x, v) - f(x, v)}{\epsilon},$$

thus an asymptotic expansion of $H_{\epsilon,\delta}^{(\alpha)} f(x, v)$ for small ϵ, δ is not well-defined without careful consideration of the behavior of δ/ϵ if it is not asymptotically bounded.

Remark 4.2. The subtlety in the characterization of the infinitesimal generator $H_{\epsilon,\delta}^{(\alpha)}$ speaks of the peculiarity of the nonhomogeneous, anisotropic diffusion processes considered in [7,67], at the presence of an underlying fibre bundle structure. These phenomena not only indicate that the horizontal and coupled diffusion operators are capable of unveiling richer geometric structures in complex real world data sets, but also imply that additional care has to be taken when tuning the bandwidth parameters in practice — the flexibility in choosing the appropriate relative scale between δ and ϵ adapts the HDM framework to a myriad of scenarios in which the relative importance of the structural information in the data objects vary drastically. The dependence of the infinitesimal generators on the ratio δ/ϵ is also reminiscent of recent trends of studying “big data” in high-dimensional statistics [68,69], where new paradigms of estimation and inference arise as the ratio between the number of features and the number of samples becomes unbounded asymptotically.

5. Finite sampling results on unit tangent bundles

The algorithm and theoretical results discussed so far are very general — we assumed that the diffusion kernel (2.3) is constructed from abstract, geodesic distances on the base and fibre manifolds. This section investigates the finite sampling aspects of horizontal diffusion maps, which connects the discrete, graph construction in Section 3 with the continuous, infinitesimal characterization in Section 4. We focus on analyzing the finite sample rate of convergence for *unit tangent bundles*, the fibre bundle with compact fibres that is as prevalent as manifolds. This is a subbundle of the tangent bundle TM (which is non-compact) defined as

$$UTM := \coprod_{x \in M} S_x, \quad S_x := \{v \in T_x M \mid g_x(v, v) = 1\} \subset T_x M.$$

In particular, UTM is a hypersurface of TM equipped with a metric induced from TM . The volume form on UTM with respect to the induced metric

$$d\Theta(x, v) = d\text{vol}_{S_x}(v) d\text{vol}_M(x)$$

is often known as the *Liouville measure* or the *kinematic density* [70, Chapter VII]. It is the only invariant measure on UTM under geodesic flows. The coupled diffusion operator on UTM can be written with the Liouville measure:

$$H_{\epsilon,\delta}^{(\alpha)} f(x, v) = \frac{\int_{UTM} K_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) f(y, w) p(y, w) d\Theta(y, w)}{\int_{UTM} K_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) p(y, w) d\Theta(y, w)}, \quad \forall f \in C^\infty(UTM).$$

The horizontal and vertical Laplacians on UTM can be defined from Δ_{TM}^H and Δ_{TM}^V by extending $f \in C^\infty(UTM)$ to $C^\infty(TM)$ and restricting the result back to UTM . Therefore, for any $f \in C^\infty(TM)$, if $\delta = O(\epsilon)$,

$$\begin{aligned} H_{\epsilon,\delta}^{(\alpha)} f(x, v) &= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \frac{[\Delta_{UTM}^H(f p^{1-\alpha}) - f \Delta_{UTM}^H p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} \\ &\quad + \delta \frac{m_{22}}{2m_0} \frac{[\Delta_{UTM}^V(f p^{1-\alpha}) - f \Delta_{UTM}^V p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} + O(\epsilon^2 + \epsilon\delta + \delta^2). \end{aligned}$$

This is consistent with the conclusion obtained in [71, Chapter 3].

The theory of HDM on tangent and unit tangent bundles are parallel to each other, but a general theory for sampling from fibre bundles of arbitrary fibre type will find it easier to consider sampling from the unit tangent bundle due to the compactness of its fibres. Sampling from tangent bundles is special, since its fibres are vector spaces and thus determined by estimating a basis; this is considered in [9, §5]. We thus study the behavior of HDM on unit tangent bundles under finite sampling. In this section, we first consider sampling without noise, i.e. where we sample exactly on unit tangent bundles; next, we study the case where the tangent spaces are empirically estimated from samples on the base manifold. The latter scenario is a proof-of-concept for applying HDM to general fibre bundles in practical situations where data representing each fibre are often acquired with noise. The proofs of Theorem 9 and Theorem 10 can be found in Appendix D. In Section 5.2, we shall demonstrate a numerical experiment on $\text{SO}(3)$ (the unit tangent bundle of the 2-sphere in \mathbb{R}^3) that addresses the two sampling strategies. Throughout this section, recall from Remark 7.2 that $\Delta_H = \Delta_{UTM}^H$ since the fibres of UTM are totally geodesic.

5.1. Rate of convergence from finite samples

5.1.1. Sampling without noise

We begin with some assumptions and definitions. Assumption 7 includes our technical assumptions, and Assumption 8 specifies the noiseless sampling strategy.

Assumption 7.

- (1) $\iota : M \hookrightarrow \mathbb{R}^D$ is an isometric embedding of a d -dimensional closed Riemannian manifold into \mathbb{R}^D , with $D \gg d$.
- (2) Let the bi-variate smooth kernel function $K : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$ be compactly supported within the unit square $[0, 1] \times [0, 1]$. The partial derivatives $\partial_1 K, \partial_2 K$ are therefore automatically compactly supported on the unit square as well. (In fact, a similar result still holds if K and its first order derivatives decay faster at infinity than any inverse polynomials; to avoid technicalities and focus on demonstrating the idea, we use compactly supported K .)

Assumption 8.

The $(N_B \times N_F)$ data points

$$\begin{matrix} x_{1,1}, & x_{1,2}, & \cdots, & x_{1,N_F} \\ x_{2,1}, & x_{2,2}, & \cdots, & x_{2,N_F} \\ \vdots & \vdots & \cdots & \vdots \\ x_{N_B,1}, & x_{N_B,2}, & \cdots, & x_{N_B,N_F} \end{matrix}$$

are sampled from UTM with respect to a probability density function $p(x, v)$ satisfying (2.5), following a two-step strategy: (i) sample N_B points ξ_1, \dots, ξ_{N_B} i.i.d. on M with respect to $\langle p \rangle$, the fibre average of p on M ; (ii) sample N_F points $x_{j,1}, \dots, x_{j,N_F}$ on S_{ξ_j} with respect to $p(\cdot | \xi_j)$, the conditional probability density.

Definition 5.1.

- (1) For $\epsilon > 0, \delta > 0$ and $1 \leq i, j \leq N_B, 1 \leq r, s \leq N_F$, define

$$\hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) = \begin{cases} K\left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} x_{i,r} - x_{j,s}\|^2}{\delta}\right), & i \neq j, \\ 0, & i = j, \end{cases}$$

where $P_{\xi_j, \xi_i} : S_{\xi_i} \rightarrow S_{\xi_j}$ is the parallel transport from S_{ξ_i} to S_{ξ_j} . Note the difference between $\hat{K}_{\epsilon, \delta}$ and $K_{\epsilon, \delta}$ defined in (2.3): $\hat{K}_{\epsilon, \delta}$ uses Euclidean distance while $K_{\epsilon, \delta}$ uses geodesic distance.

(2) For $0 \leq \alpha \leq 1$, define

$$\hat{p}_{\epsilon, \delta}(x_{i,r}) = \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s})$$

and the *empirical α -normalized kernel* $\hat{K}_{\epsilon, \delta}^\alpha$

$$\hat{K}_{\epsilon, \delta}^\alpha(x_{i,r}, x_{j,s}) = \frac{\hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s})}{\hat{p}_{\epsilon, \delta}^\alpha(x_{i,r}) \hat{p}_{\epsilon, \delta}^\alpha(x_{j,s})}, \quad 1 \leq i, j \leq N_B, 1 \leq r, s \leq N_F.$$

(3) For $0 \leq \alpha \leq 1$ and $f \in C^\infty(UTM)$, denote the α -normalized *empirical horizontal diffusion operator* by

$$\hat{H}_{\epsilon, \delta}^\alpha f(x_{i,r}) = \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}^\alpha(x_{i,r}, x_{j,s}) f(x_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}^\alpha(x_{i,r}, x_{j,s})}.$$

Theorem 9 (*Finite sampling without noise*). *Under Assumption 7 and Assumption 8, if*

- (i) $\delta = O(\epsilon)$ as $\epsilon \rightarrow 0$;
- (ii)

$$\lim_{\substack{N_B \rightarrow \infty \\ N_F \rightarrow \infty}} \frac{N_F}{N_B} = \beta \in (0, \infty),$$

then for any $x_{i,r}$ with $1 \leq i \leq N_B$ and $1 \leq r \leq N_F$, as $\epsilon \rightarrow 0$ (and thus $\delta \rightarrow 0$), with high probability

$$\begin{aligned} \hat{H}_{\epsilon, \delta}^\alpha f(x_{i,r}) &= f(x_{i,r}) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_H [fp^{1-\alpha}](x_{i,r})}{p^{1-\alpha}(x_{i,r})} - f(x_{i,r}) \frac{\Delta_H p^{1-\alpha}(x_{i,r})}{p^{1-\alpha}(x_{i,r})} \right] \\ &\quad + \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_{UTM}^V [fp^{1-\alpha}](x_{i,r})}{p^{1-\alpha}(x_{i,r})} - f(x_{i,r}) \frac{\Delta_{UTM}^V p^{1-\alpha}(x_{i,r})}{p^{1-\alpha}(x_{i,r})} \right] \\ &\quad + O\left(\epsilon^2 + \delta^2 + \theta_*^{-1} N_B^{-\frac{1}{2}} \epsilon^{-\frac{d}{4}}\right), \end{aligned} \quad (5.1)$$

where

$$\theta_* = 1 - \frac{1}{1 + \epsilon^{\frac{d}{4}} \delta^{\frac{d-1}{4}} \sqrt{\frac{N_F}{N_B}}}.$$

The proof of Theorem 9 is deferred to Appendix D.

Remark 5.1. Theorem 9 reflects the difference in the finite-sample rate of convergence between considering horizontal diffusion and standard diffusion on the total manifold of the fibre bundle. For instance, in the special case $\epsilon = \delta$, by [72], the variance error associated with the standard diffusion maps on the total

manifold is $O\left(N_F^{-1/2}N_B^{-1/2}\epsilon^{-(2d-1)/4}\right)$, while the variance error in (5.1) is $O\left(\theta_*^{-1}N_B^{-1/2}\epsilon^{-d/4}\right)$. This is another evidence demonstrating the difference between horizontal diffusion maps and standard diffusion maps on the total manifold of the fibre bundle; see also Remark 4.1.

5.1.2. Sampling from empirical tangent spaces

In practice, it has been shown in [9] that, under the manifold assumption, a local PCA procedure can be used for estimating tangent spaces from a point cloud; we are using PCA here as a procedure that determines the dimension of a local good linear approximation to the manifold, and also, conveniently, provides a good basis, which can be viewed as a basis for each tangent plane. To sample on these tangent spaces, it suffices to repeatedly sample coordinate coefficients from a fixed standard unit sphere; each sample can be interpreted as giving the coordinates of a point (approximately) on the tangent space. Parallel-transports will take the corresponding point that truly lies on the tangent space at ξ to the tangent space at ζ , another point on the manifold. This new tangent space is, however, again known only approximately; points in this approximate space are characterized by coordinates with respect to the local PCA basis at ζ . We can thus express the whole (approximate) parallel-transport procedure by maps between coordinates with respect to PCA basis at ξ to sets of coordinates at ζ ; these changes of coordinates incorporate information on the choices of basis at each end as well as on the parallel-transport itself.

Let us now describe this in more detail, setting up notations along the way. Throughout this section, Assumption 7 still holds. Let $\{\xi_1, \dots, \xi_{N_B}\}$ be a collection of i.i.d. samples from M ; then the local PCA procedure can be summarized as follows: for any ξ_j , $1 \leq j \leq N_B$, let $\xi_{j_1}, \dots, \xi_{j_k}$ be its k nearest neighboring points. Then

$$X_j = [\xi_{j_1} - \xi_j, \dots, \xi_{j_k} - \xi_j]$$

is a $D \times k$ matrix. Let K_{PCA} be a positive monotonic decreasing function supported on the unit interval, e.g. the *Epanechnikov kernel* [73]

$$K_{\text{PCA}}(u) = (1 - u^2) \chi_{[0,1]},$$

where χ is the indicator function. Fix a scale parameter $\epsilon_{\text{PCA}} > 0$, let D_j be the $k \times k$ diagonal matrix

$$D_j = \text{diag} \left(\sqrt{K_{\text{PCA}} \left(\frac{\|\xi_j - \xi_{j_1}\|}{\sqrt{\epsilon_{\text{PCA}}}} \right)}, \dots, \sqrt{K_{\text{PCA}} \left(\frac{\|\xi_j - \xi_{j_k}\|}{\sqrt{\epsilon_{\text{PCA}}}} \right)} \right)$$

and carry out the singular value decomposition (SVD) of matrix $X_j D_j$ as

$$X_j D_j = U_j \Sigma_j V_j^\top.$$

An estimated basis B_j for the local tangent plane at ξ_j is formed by the first d left singular vectors (corresponding to the d largest singular values in Σ_j), arranged into a matrix as follows:

$$B_j = [u_j^{(1)}, \dots, u_j^{(d)}] \in \mathbb{R}^{D \times d}.$$

Note that the intrinsic dimension d is generally not known *a priori*. The authors of [9] proposed a procedure that first estimates local dimensions from the decay of singular values in Σ_j and then sets d to be the median of all local dimensions; [74] proposed a different approach based on multi-scale singular value decomposition.

Once a pair of estimated bases B_i, B_j is obtained for neighboring points ξ_i, ξ_j , one estimates a parallel-transport from $T_{\xi_i} M$ to $T_{\xi_j} M$ as

$$O_{ji} := \arg \min_{O \in O(d)} \|O - B_j^\top B_i\|_{\text{HS}},$$

where $\|\cdot\|_{\text{HS}}$ is the Hilbert-Schmidt norm. Though this minimization problem is non-convex, it has an efficient closed-form solution via the SVD of $B_i^\top B_j$, namely

$$O_{ji} = UV^\top, \quad \text{where } B_j^\top B_i = U\Sigma V^\top \text{ is the SVD of } B_j^\top B_i.$$

It is worth noting that O_{ji} depends on the bases; it operates on the coordinates of tangent vectors under B_i and B_j , as explained above. O_{ji} approximates the true parallel-transport P_{ξ_j, ξ_i} (composed with the bases-expansions) with an error of $O(\epsilon_{\text{PCA}})$, in the sense of [9, Lemma B.1].

We summarize our sampling strategy for this section (with some new notations) in the following definition.

Definition 5.2.

(1) Let $\{\xi_1, \dots, \xi_{N_B}\}$ be a collection of samples from the base manifold M , i.i.d. with respect to some probability density function $\bar{p} \in C^\infty(M)$. For each ξ_j , $1 \leq j \leq N_B$, sample N_F points uniformly from the $(d-1)$ -dimensional standard unit sphere S^{d-1} in \mathbb{R}^d , and denote the set of samples as $\mathcal{C}_j = \{c_{j,1}, \dots, c_{j,N_F}\}$, where each $c_{j,s}$ is a $d \times 1$ column vector. Using the basis B_j estimated from the local PCA procedure, each $c_{j,s}$ corresponds to an “approximate tangent vector at ξ_j ”, denoted as

$$\tau_{j,s} := B_j c_{j,s}.$$

We use the notation \mathcal{S}_j for the unit sphere in the estimated tangent space (i.e., the column space of B_j). Note that the $\tau_{j,1}, \dots, \tau_{j,N_F}$ are uniformly distributed on \mathcal{S}_j .

(2) By [9, lemma B.1], for any B_j there exists a $D \times d$ matrix Q_j , such that the columns of Q_j constitutes an orthonormal basis for $\iota_* T_{\xi_j} M$ and

$$\|B_j - Q_j\|_{\text{HS}} = O(\epsilon_{\text{PCA}}).$$

We define the *tangent projection* from $\iota_* S_{\xi_j}$ to the estimated tangent plane as

$$\tau_{j,s} \mapsto \bar{\tau}_{j,s} = \frac{Q_j Q_j^\top \tau_{j,s}}{\|Q_j Q_j^\top \tau_{j,s}\|}.$$

This map is well-defined for sufficiently small ϵ_{PCA} , and then it is an isometry. Its inverse is given by

$$\bar{\tau}_{j,s} \mapsto \tau_{j,s} = \frac{B_j B_j^\top \bar{\tau}_{j,s}}{\|B_j B_j^\top \bar{\tau}_{j,s}\|}.$$

Note that we have

$$\|\tau_{j,s} - \bar{\tau}_{j,s}\| \leq C \epsilon_{\text{PCA}}$$

for some constant $C > 0$ independent of indices j, s . Since we sample each \mathcal{S}_j uniformly and the projection map $\tau_{j,s} \mapsto \bar{\tau}_{j,s}$ is an isometry, the points $\{\bar{\tau}_{j,1}, \dots, \bar{\tau}_{j,N_F}\}$ are also uniformly distributed on S_{ξ_j} . The points

$$\begin{aligned} \bar{\tau}_{1,1}, & \quad \bar{\tau}_{1,2}, \quad \dots, \quad \bar{\tau}_{1,N_F} \\ \bar{\tau}_{2,1}, & \quad \bar{\tau}_{2,2}, \quad \dots, \quad \bar{\tau}_{2,N_F} \\ & \quad \vdots \quad \vdots \quad \dots \quad \vdots \\ \bar{\tau}_{N_B,1}, & \quad \bar{\tau}_{N_B,2}, \quad \dots, \quad \bar{\tau}_{N_B,N_F} \end{aligned}$$

are therefore distributed on UTM according to a joint probability density function p on UTM defined as

$$p(x, v) = \bar{p}(x), \quad \forall (x, v) \in UTM.$$

As in Assumption 8, we assume p satisfies (2.5), i.e.,

$$0 < p_m \leq p(x, v) = \bar{p}(x) \leq p_M < \infty, \quad \forall (x, v) \in UTM$$

for positive constants p_m, p_M .

(3) For $\epsilon > 0, \delta > 0$ and $1 \leq i, j \leq N_B, 1 \leq r, s \leq N_F$, define

$$\mathcal{H}_{\epsilon, \delta}(\bar{\tau}_{i,r}, \bar{\tau}_{j,s}) = \begin{cases} K\left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|O_{ji}c_{i,r} - c_{j,s}\|^2}{\delta}\right), & i \neq j, \\ 0, & i = j, \end{cases}$$

where O_{ji} is the estimated parallel-transport from $T_{\xi_i}M$ to $T_{\xi_j}M$.

(4) For $0 \leq \alpha \leq 1$, define

$$\hat{q}_{\epsilon, \delta}(\bar{\tau}_{i,r}) = \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \mathcal{H}_{\epsilon, \delta}(\bar{\tau}_{i,r}, \bar{\tau}_{j,s})$$

and

$$\mathcal{H}_{\epsilon, \delta}^{\alpha}(\bar{\tau}_{i,r}, \bar{\tau}_{j,s}) = \frac{\mathcal{H}_{\epsilon, \delta}(\bar{\tau}_{i,r}, \bar{\tau}_{j,s})}{\hat{q}_{\epsilon, \delta}^{\alpha}(\bar{\tau}_{i,r}) \hat{q}_{\epsilon, \delta}^{\alpha}(\bar{\tau}_{j,s})}, \quad 1 \leq i, j \leq N_B, 1 \leq r, s \leq N_F.$$

(5) For $0 \leq \alpha \leq 1$ and $f \in C^{\infty}(UTM)$, denote

$$\mathcal{H}_{\epsilon, \delta}^{\alpha} f(\bar{\tau}_{i,r}) = \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \mathcal{H}_{\epsilon, \delta}^{\alpha}(\bar{\tau}_{i,r}, \bar{\tau}_{j,s}) f(\bar{\tau}_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \mathcal{H}_{\epsilon, \delta}^{\alpha}(\bar{\tau}_{i,r}, \bar{\tau}_{j,s})}.$$

Theorem 10 (*Finite sampling from empirical tangent planes*). *In addition to Assumption 7, suppose*

- (i) $\epsilon_{PCA} = O\left(N_B^{-\frac{2}{d+2}}\right)$ as $N_B \rightarrow \infty$;
- (ii) *As $\epsilon \rightarrow 0$, $\delta = O(\epsilon)$ and $\delta \gg \left(\epsilon_{PCA}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right)$;*
- (iii)

$$\lim_{\substack{N_B \rightarrow \infty \\ N_F \rightarrow \infty}} \frac{N_F}{N_B} = \beta \in (0, \infty).$$

Then for any $\tau_{i,r}$ with $1 \leq i \leq N_B$ and $1 \leq r \leq N_F$, as $\epsilon \rightarrow 0$ (and thus $\delta \rightarrow 0$), with high probability

$$\begin{aligned} \mathcal{H}_{\epsilon,\delta}^{\alpha} f(\bar{\tau}_{i,r}) &= f(\bar{\tau}_{i,r}) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_{UTM}^H [fp^{1-\alpha}](\bar{\tau}_{i,r})}{p^{1-\alpha}(\bar{\tau}_{i,r})} - f(\bar{\tau}_{i,r}) \frac{\Delta_{UTM}^H p^{1-\alpha}(\bar{\tau}_{i,r})}{p^{1-\alpha}(\bar{\tau}_{i,r})} \right] \\ &\quad + \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_{UTM}^V [fp^{1-\alpha}](\bar{\tau}_{i,r})}{p^{1-\alpha}(\bar{\tau}_{i,r})} - f(\bar{\tau}_{i,r}) \frac{\Delta_{UTM}^V p^{1-\alpha}(\bar{\tau}_{i,r})}{p^{1-\alpha}(\bar{\tau}_{i,r})} \right] \\ &\quad + O\left(\epsilon^2 + \epsilon\delta + \delta^2 + \theta_*^{-1} N_B^{-\frac{1}{2}} \epsilon^{-\frac{d}{4}} + \delta^{-1} \left(\epsilon_{PCA}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right)\right), \end{aligned} \quad (5.2)$$

where

$$\theta_* = 1 - \frac{1}{1 + \epsilon^{\frac{d}{4}} \delta^{\frac{d-1}{4}} \sqrt{\frac{N_F}{N_B}}}.$$

We give a proof of Theorem 10 in Appendix D.

5.2. Numerical experiments

The unit tangent bundle is of special interest since UTM is a compact Riemannian manifold whenever M is compact, enabling finite sampling and numerically validating Theorem 4 and Theorem 6. We present in below a numerical experiment on $SO(3)$, the unit tangent bundle of the standard two-dimensional sphere in \mathbb{R}^3 , along with an analysis of sampling errors on general unit tangent bundles. In the first step, we uniformly sample $N_B = 2,000$ points $\{\xi_1, \dots, \xi_{N_B}\}$ on the unit sphere S^2 , and find for each sample point the $K_B = 100$ nearest neighbors in the point cloud. Next, we sample $N_F = 50$ vectors of unit length tangent to the unit sphere at each sample point (which in this case is a circle), thus collecting a total of $N_B \times N_F = 100,000$ points on $UTS^2 = SO(3)$, denoted as

$$\{x_{j,s} \mid 1 \leq j \leq N_B, 1 \leq s \leq N_F\}.$$

The horizontal diffusion matrix H is then constructed as an $N_B \times N_B$ block matrix with block size $N_F \times N_F$, and H_{ij} (the (i, j) -th block of H) is non-zero only if the sample points ξ_i, ξ_j are each among the K_B -nearest neighbors of the other; when H_{ij} is non-zero, its (r, s) -entry ($1 \leq r, s \leq N_F$) is non-zero only if $P_{\xi_j, \xi_i} x_{i,r}$ and $x_{j,s}$ are each among the $K_F = 50$ nearest neighbors of the other, and in that case for all $i \neq j$

$$H_{ij}(r, s) = \exp \left[- \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon} + \frac{\|P_{\xi_j, \xi_i} x_{i,r} - x_{j,s}\|^2}{\delta} \right) \right], \quad (5.3)$$

where the choices of ϵ, δ will be explained below. The diagonal blocks are set to zero as in Definition 5.1. Note that for the unit sphere S^2 the parallel-transport from $T_{\xi_i} S^2$ to $T_{\xi_j} S^2$ can be explicitly constructed as a rotation along the axis $\xi_i \times \xi_j$. Finally, we form the α -normalized horizontal diffusion matrix H_{α} by

$$(H_{\alpha})_{ij}(r, s) = \frac{H_{ij}(r, s)}{\left(\sum_{l=1}^{N_B} \sum_{m=1}^{N_F} H_{il}(r, m) \right)^{\alpha} \left(\sum_{k=1}^{N_B} \sum_{n=1}^{N_F} H_{jk}(r, n) \right)^{\alpha}}, \quad (5.4)$$

and solve the eigenvalue problem

$$\left(D^{-\frac{1}{2}} H_{\alpha} D^{-\frac{1}{2}} \right) U = U \Lambda \quad (5.5)$$

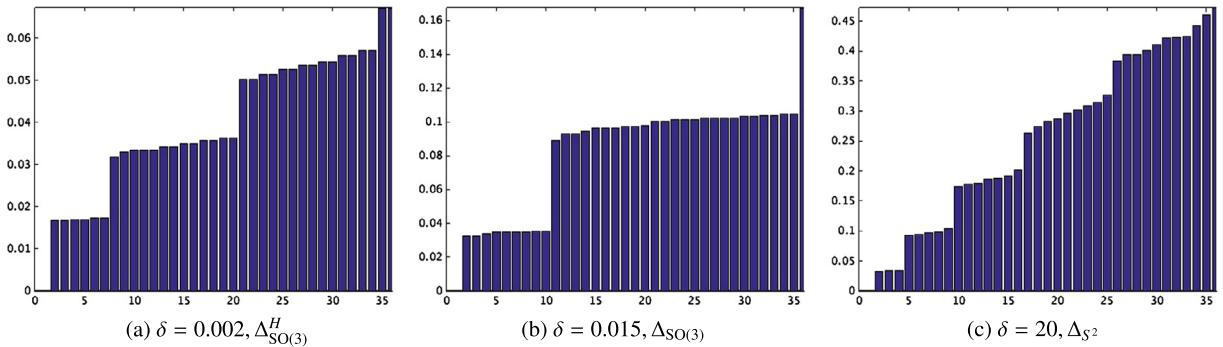


Fig. 3. Bar plots of the smallest 36 eigenvalues of $I - D^{-1}H_{\alpha}$ with $\alpha = 1$, $\epsilon = 0.2$, and varying δ values (sampling without noise). *Left:* When $\delta \ll \epsilon$, $H_{\epsilon,\delta}^1$ approximates the heat kernel of $\Delta_{SO(3)}^H$, of which the multiplicities of largest eigenvalues are $1, 6, 13, \dots$; *Middle:* When $\delta \approx \epsilon$, $H_{\epsilon,\delta}^1$ approximates the heat kernel of $\Delta_{SO(3)}$, of which the multiplicities of largest eigenvalues are $1, 9, 25, \dots$; *Right:* When $\delta \gg \epsilon$, $H_{\epsilon,\delta}^1$ approximates the heat kernel of Δ_{S^2} , of which the multiplicities of largest eigenvalues are $1, 3, 5, \dots$.

where D is the $(N_B N_F) \times (N_B N_F)$ diagonal matrix with entry (k, k) equal to the k -th column sum of H_{α} :

$$D(k, k) = \sum_{v=1}^{N_B N_F} H_{\alpha}(k, v),$$

and Λ is a diagonal matrix of the same dimensions. Throughout this experiment, we fix $\alpha = 1$, $\epsilon = 0.2$ and choose various values of δ ranging from 0.0005 to 50, and observe the spacing of the eigenvalues stored in Λ .

The purpose of this experiment is to investigate the influence of the ratio $\gamma = \delta/\epsilon$ on the spectral behavior of graph horizontal Laplacians. As shown in Fig. 3, the spacing in the spectrum of these graph horizontal Laplacians follow patterns similar to the multiplicities of the eigenvalues of corresponding Laplacians on $SO(3)$ (governed by the relative size of δ and ϵ). In Fig. 3(a), $\delta \ll \epsilon$, hence the graph horizontal Laplacian approximates the horizontal Laplacian on $SO(3)$ (according to Theorem 4 and Corollary 5), in which the smallest eigenvalues have multiplicities $1, 6, 13, \dots$; in Fig. 3(b), $\delta = O(\epsilon)$, hence the graph horizontal Laplacian approximates the total Laplacian on $SO(3)$ (again, according to Theorem 4 and Corollary 5), with eigenvalue multiplicities $1, 9, 25, \dots$; in Fig. 3(c), $\delta \gg \epsilon$, hence the graph horizontal Laplacian approximates the Laplacian on the base manifold S^2 (according to Theorem 6), with eigenvalue multiplicities $1, 3, 5, \dots$. Note that in Fig. 3(c) we fixed ϵ and pushed δ to ∞ , which essentially corresponds to the limit process in (4.9) rather than (4.1). Moreover, if in each figure we divide the sequence of eigenvalues by the smallest non-zero eigenvalue, the resulting sequence coincides with the list of eigenvalues of the corresponding manifold Laplacian up to numerical error. For a description of the spectrum of these partial differential operators, see [75, Chapter 2].

Similar numerical results have been observed for sampling from empirically estimated tangent spaces; we refer interested readers to [71, §3.5.2].

6. Application to automated geometric morphometrics

The HDM framework can be applied to any data set with pairwise structural correspondences. In many applications, such structural correspondences are readily available through a registration procedure, and have been used to compute similarity scores or distances between objects of interest. In this section, we sketch the application of HDM to automated geometric morphometrics. In a nutshell, this is an unsupervised learning problem with heterogeneous or unorganized data, for which feature engineering is particularly difficult; moreover, it is hard to apply kernel methods due to the lack of an informative kernel function.

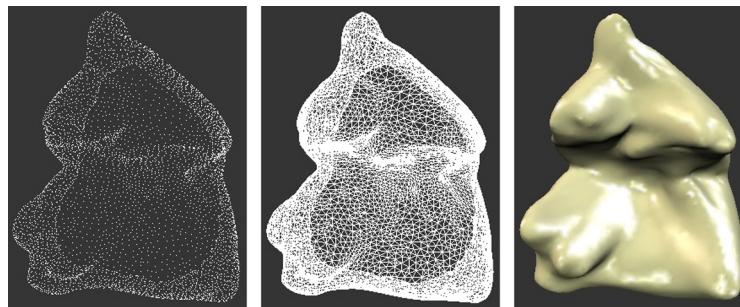


Fig. 4. A second mandibular molar of a Philippine flying lemur (*Cynocephalus volans*), represented as a point cloud (left), wireframe (middle), and a piecewise linear surface (right).

We expect problems arising from machine learning, pattern recognition, and computer vision facing similar difficulties to benefit from the proposed approach.

Geometric morphometrics is the quantitative analysis of shape variation and their correlation with other traits for biological organisms. For instance, it is often of interest to geometric morphometricians to understand quantitatively the amount of the shape variation explained by geometric features within a collection of shapes. They typically select equal numbers of consistently *homologous* landmark points on each surface [76], corresponding to a mental model of a latent “template,” of which every individual shape is an instantiation. In statistical shape analysis, this landmark-based approach is developed in the framework of *Procrustes analysis* [77]. Obviously, such an analysis is limited by the knowledge of landmark placement. From a mathematical point of view, extracting a limited number of landmarks from a continuous surface inevitably loses geometric information, unless when the shapes under consideration are solely determined by the landmarks (e.g. polygonal shapes, as considered in [78] [79]), which is rarely the case for geometric morphometricians in biology; from a practical point of view, the requirement that an equal number of landmarks must be chosen on each shape is sometimes unrealistic due to the complex evolutionary and developmental process. Manually placing landmarks on each shape among a large collection is also a tedious task, and the skill to perform it “correctly” typically requires years of professional training; even then the “correctness” or the number of landmarks one should fix for a collection of shapes can be subject to debate among experts. These difficulties are gradually and continuously being addressed by a recent trend that advocates *automated* workflows to bypass the repetitive, laborious, and time-consuming process of manual landmark placement on large collections of 3D digitized anatomical surfaces (see [20,80–86,71,87–91] and references therein).

The digitized morphological data set contains hundreds of triangular meshes (see Fig. 4) of diverse size, topology, and quality, each representing an anatomical surface reconstructed from MicroCT images. In [20], the authors introduced *Continuous Procrustes Distance* (CPD) between surfaces with disk-type topology, and used conformal parameterization (uniformization) to design a fast algorithm that computes automatically (without landmarks) distances between pairs of morphological surfaces that would be at least as effective, for species discrimination, as Procrustes distances computed from user-defined landmarks [80]. Upon computing a distance between pairs of surfaces, the algorithm in [20] minimize an energy functional depending on the pair, over an admissible set of *correspondence maps*; the distance is indeed the value of the functional at the optimal correspondence map (Fig. 1). This approach has recently been followed by other authors as well [86].

Detailed inspection of the optimal correspondence maps generated from the algorithm [20] showed that some of them presented serious anomalies, such as reversed alignments of the anteroposterior/buccolingual axes [89]. It may seem surprising that the algorithms, despite sometimes producing these erroneous maps, nevertheless were sufficiently successful in capturing sample geometry to achieve the success rate reported in [80]. As we extend the algorithm in [20] in different directions, the correspondence maps became an

important explicit goal of the algorithm, as opposed to an interesting by-product. While CPD automates the traditional Procrustes analysis, the optimal correspondence maps parallel the landmark-identification procedure performed mentally by geometric morphometricians. Moreover, these maps provide detailed information about correlations, often not fully retained when an energy functional summarizes a similarity measure, between functional or developmental regions on different shapes.

HDM is a natural algorithmic framework for unsupervised learning from structural correspondences maps. In this section, we apply HDM to a data set consisting of 50 discretized triangular meshes of the second mandibular molar of prosimian primates and nonprimate close relatives. The 50 meshes are evenly divided into 5 genus groups: *Alouatta*, *Ateles*, *Brachyteles*, *Callicebus*, and *Saimiri*; each mesh contains about 5,000 vertices and 10,000 faces. We compute first all pairwise CPD and correspondence maps, then all pairwise *Horizontal Base Diffusion Distance* (HBDD) from the distances and maps. The 50×50 distance matrices are finally embedded into \mathbb{R}^3 for comparison via *multi-dimensional scaling* (MDS).

The HBDD is constructed from CPD as follows. For each pair of triangular meshes S_i, S_j in the data set, denote their CPD as d_{ij} , and the optimal correspondence from S_i to S_j as f_{ij} . Note that $d_{ij} = d_{ji}$ and $f_{ij} = f_{ji}^{-1}$. In the first step, we discretize each surface area measure $\mu_j = d\text{vol}_{S_j}$ into a linear combination of Dirac delta measures supported on vertices of S_j , where each vertex of S_j is assigned 1/3 of the surface area of its *one-ring neighborhood*. We then *soften* each bijective smooth map f_{ij} into a *transport plan* matrix w_{ij} , the s -th row of which records the transition probability from vertex $x_{i,s}$ of S_i to each vertex on S_j ; moreover, the specific softening we choose here allows each $x_{i,s}$ to jump (in one step) only to the three vertices of the unique¹ triangular face on S_j that contains $f_{ij}(x_{i,s})$. If $x_{j,r}$ is a vertex on S_j that can be reached from $x_{i,s}$ in one step of the random walk, we set the transition probability between $x_{i,s}$ and $x_{j,r}$ proportional to

$$\exp\left(-\frac{\|f_{ij}(x_{i,s}) - x_{j,r}\|^2}{\epsilon_F}\right),$$

where ϵ_F is a prescribed positive constant playing the role of the vertical bandwidth parameter δ in (2.6). For this specific data set, we choose $\epsilon_F = 0.001$ which is the order of magnitude of the average distance between adjacent vertices on each mesh in the data set. Next, we construct the horizontal diffusion matrix H as a 50×50 block matrix, with block (i, j)

$$H(i, j) = \begin{cases} \exp\left(-\frac{d_{ij}^2}{\epsilon_B}\right) \cdot w_{ij} & \text{if } S_j \text{ is within the } N_B\text{-neighborhood of } S_i, \\ 0 & \text{otherwise.} \end{cases}$$

We chose for this data set $N_B = 4$ and $\epsilon_B = 0.03$. These parameters are picked empirically, where 0.03 is usually the maximum CPD between surfaces that belong to the same species group. We then construct the normalized graph horizontal Laplacian $L_{\alpha,*}^H$ from H , as in (3.7), and solve for its largest 100 eigenvalues and corresponding eigenvectors. From this eigen-decomposition we compute the *horizontal base diffusion map* (HBDM) as in (3.12), obtaining an embedding of the data set into $\mathbb{R}^{\binom{100}{2}} = \mathbb{R}^{4950}$. Though this embedding is still high dimensional, it is only 1/3 of the original dimensionality (approximately $5000 \times 3 = 15000$). The HBDD between each pair S_i, S_j is then defined as the Euclidean distance between their images embedded in \mathbb{R}^{4950} , as in (3.14). For comparison, we also embed the standard *Diffusion Distance* matrix in to \mathbb{R}^3 using MDS. As shown in Fig. 5, HBDD demonstrates the most clear pattern of species clusters among the three

¹ It is conceivable that $f_{ij}(x_{i,s})$ could fall on the edge shared by two triangles in S_j , or even on a vertex of S_j shared by more than 2 triangles. While this rarely happens in practice, in our implementation for this application we resolve such conflicts by assigning $f_{ij}(x_{i,s})$ randomly to any of the qualified triangles. This is because we express $f_{ij}(x_{i,s})$ as a barycentric combination of the vertices of the triangle to which it is assigned, and thus the softening is in fact independent of the specific choice made.

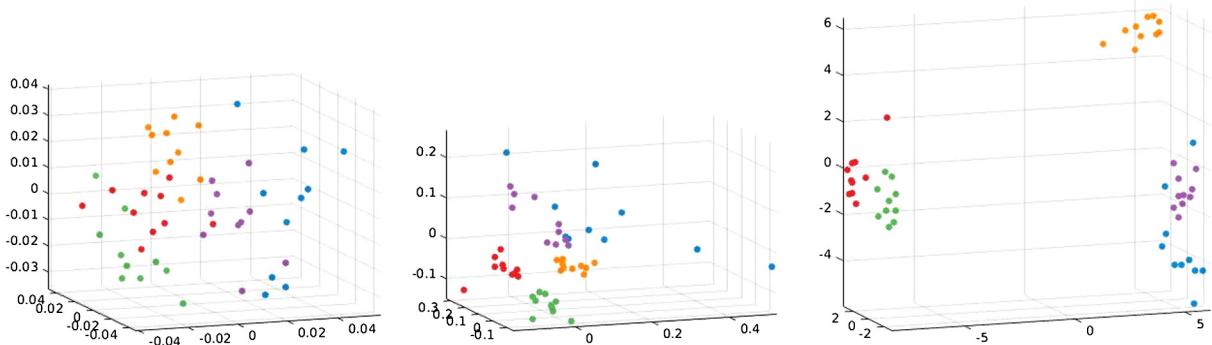


Fig. 5. Embeddings of CPD (left), DM (middle), and HBDD (right) matrices into \mathbb{R}^3 using *Multi-dimensional Scaling* (MDS). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

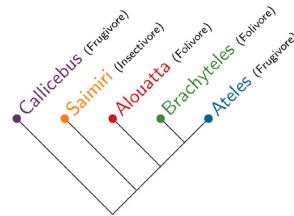


Fig. 6. Phylogeny of the 5 species groups *Alouatta*, *Ateles*, *Brachyteles*, *Callicebus*, and *Saimiri*. HBDD (see Fig. 5) reflects the dietary categories but not the phylogeny.

distances. It is even more interesting to notice that HBDD reflects the dietary groups within the data set (see Fig. 6): folivores *Alouatta* (red) and *Brachyteles* (green) are adjacent to each other in the rightmost panel of Fig. 5, so are frugivores *Ateles* (blue) and *Callicebus* (purple); the insectivore *Saimiri* (yellow) is far from the other herbivorous groups.

For applications in geometric morphometrics, a major advantage of HDM over persistence-diagram-based methods is the morphological interpretability. This interpretability amounts to a globally consistent manner to identify corresponding regions on each shape in the data set and is potentially useful for subsequent studies of the evolutionary and developmental history. In standard morphologists' practice, such correspondences are assessed visually and manually; recent progress in techniques for generating and analyzing digital representations led to major advances [92–94] but still require the input of anatomical landmarks from the user. In contrast, by spectral clustering on the point cloud embedded into \mathbb{R}^{100} by HDM, we can easily obtain a globally consistent segmentation for all surfaces, see Fig. 7.

7. Discussion and future work

This paper introduced *horizontal diffusion maps* (HDM), a novel semi-supervised learning framework for the analysis and organization of a class of complex data sets, in which individual structures at each data point carry abundant information that can not be easily abstracted away by a pairwise similarity measure. We also introduced the fibre bundle assumption, a generalization of the manifold assumption, and showed that under this assumption HDM provides embeddings for both the base and the total manifold; furthermore, the flexibility of the HDM framework enables us to view VDM and the standard diffusion maps (DM) as special cases. The rest of the paper focused on analyzing the asymptotic behavior of HDM, with convergence rate estimated for finite sampling on unit tangent bundles. These results provide the mathematical foundation for HDM on fibre bundles, and motivate further studies concerning both wider applicability and deeper mathematical understanding of the algorithmic framework. We conclude this paper by listing a few potential directions for further exploration.

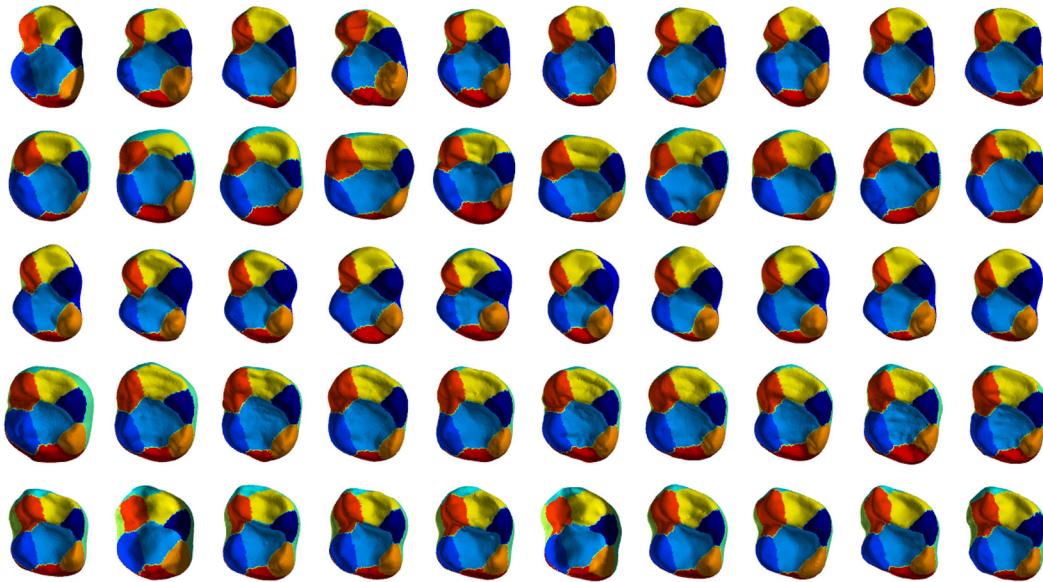


Fig. 7. Automated landmarking: consistent segmentation of 50 lemur teeth by spectral clustering in the Euclidean space to which HDM embeds. From the top row to the bottom row: Alouatta, Ateles, Brachyteles, Callicebus, Saimiri.

- 1) *Spectral Convergence of HDM.* The convergence results in this paper are pointwise; as in [28,35], we believe that it is possible to show the convergence of the eigenvalues and eigenvectors of the graph horizontal Laplacians to the eigenvalues and eigenvectors of the manifold horizontal Laplacians, thus establishing the mathematical foundation for the spectral analysis of the HDM framework. Moreover, the horizontal diffusion maps differ from diffusion maps and vector diffusion maps in that the fibres tend to be registered to a common “template”, which, to our knowledge, is a new phenomenon addressed here for the first time.
- 2) *Spectral Clustering and Cheeger-Type Inequalities.* An important application of graph Laplacian is spectral clustering (graph partitioning). In a simple case, for a connected graph, the eigenvector corresponding to the smallest positive eigenvalue of the graph Laplacian partitions the graph vertices into two similarly sized subsets, in such a way that the number of edges across the subsets is as small as possible. In spectral graph theory [1], the classical Cheeger’s Inequality provides upper and lower bounds for the performance of the partition; recently, [95] established similar results for the graph connection Laplacian, the central object of VDM. We believe that similar inequalities can be established for graph horizontal Laplacians as well, with potentially more interesting behavior of the eigenvectors. For instance, we observed in practice that the eigenvector corresponding to the smallest positive eigenvalue of the graph horizontal Laplacian stably partitions all the fibres in a globally consistent manner.
- 3) *Multiscale Analysis and Hierarchical Coarse-Graining.* Multiscale representation of massive, complex data sets based on similarity graphs is an interesting and fruitful application of diffusion operators [96,24]. Based on HDM, one can build a similar theory for data sets possessing fibre bundle structures, providing a natural framework for coarse-graining that is meaningful (or even possible) only when performed simultaneously on the base and fibre manifolds. Moreover, since the horizontal diffusion matrix is often of high dimensionality, an efficient approach to store and compute its powers will significantly improve the applicability of the HDM algorithm. We thus expect to develop a theory of *horizontal diffusion wavelets* and investigate their performance on real data sets with underlying fibre bundle structures.

Software Matlab code accompanying this paper can be found at <https://github.com/trgao10/HDM>.

Acknowledgments

The author would like to thank Ingrid Daubechies, Robert L. Bryant, and Yaron Lipman for inspirational discussions from the very beginning of this project. This work is partially supported by NSF DMS-1854831, Simons Math+X Investigators Award 400837, and NSF BCS-1552848.

Appendix A. Fibre bundles and connections

There are different ways to define a connection on a fibre bundle. For the sake of generality, we adopt here the treatment in [45] or [97] from a Riemannian submersion point of view; see also [40,64] for more detailed discussions.

For any fibre bundle $\mathcal{E} = (E, M, F, \pi)$, the bundle projection map $\pi : E \rightarrow M$ descends canonically to its differential $d\pi$ between tangent bundles TE and TM , defining linear surjective homomorphisms between tangent planes $T_e E$ and $T_{\pi(e)} M$ for any $e \in E$. We denote VE for the *vertical bundle*, a sub-bundle of TE defined as the kernel of the differential map $d\pi : TE \rightarrow TM$. A *horizontal bundle* HE is a sub-bundle of TE that is supplementary to VE in the sense that $TE = HE \oplus VE$, or equivalently

$$T_e E = H_e E \oplus V_e E \quad \text{for all } e \in E.$$

Here $H_e E$, $V_e E$ stand for the fibres of HE , VE above $e \in E$, respectively; we shall refer to $H_e E$, $V_e E$ as the *horizontal tangent space* and *vertical tangent space* at $e \in E$ for future convenience, and denote

$$\mathcal{H} : TE \rightarrow HE, \quad \mathcal{V} : TE \rightarrow VE \tag{A.1}$$

for the corresponding *horizontal projection* and *vertical projection*. Note that although VE is canonically defined, the choice of HE is arbitrary at this point. Since $d\pi|_{H_e E} : H_e E \rightarrow T_{\pi(e)} M$ is a linear isomorphism, for any tangent vector $u \in T_{\pi(e)} M$ there exists a unique tangent vector $\bar{u} \in H_e E$ such that $d\pi_e(\bar{u}) = u$; we call \bar{u} the *horizontal lift* of u . Furthermore, we know from simple ODE theory (and the smoothness of HE) that for any vector field $X \in \Gamma(M, TM)$ there exists a unique horizontal lift $\bar{X} \in \Gamma(E, HE)$ such that $d\pi_e(\bar{X}_e) = X_{\pi(e)}$ for all $e \in E$.

In the rest of this paper, a path $\gamma : [0, T] \rightarrow E$ is *horizontal* if all tangent vectors along γ are in HE . Given a path $c : [0, T] \rightarrow M$, a *horizontal lift* of c is any horizontal path \bar{c} in E that projects to c under the bundle projection π , i.e. $\pi \circ \bar{c} = c$. Again, by horizontally lifting the tangent vector field along the path from TM to HE and solving the ODE system (where the overline again stands for horizontally lifted tangent vectors)

$$\frac{d\bar{c}}{dt} = \overline{\left(\frac{dc}{dt} \right)}, \quad t \in [0, T]$$

we can uniquely lift any piecewise smooth path c in M starting at $\pi(e) \in M$ to a horizontal path \bar{c} in E starting at $e \in E$, at least locally around $c(0)$. We call HE a *Ehresmann connection* [97], or *connection* hereafter, if any path in M starting at $m \in M$ can be *globally* horizontally lifted to E with any given initial point $e \in E$ satisfying $e \in \pi^{-1}(m)$. Such a lifting property is guaranteed, for instance, on any Riemannian submersion $\pi : E \rightarrow M$ with geodesically complete total space E , in which case the submersion is known to be a locally trivial fibration [56].

We shall focus on Ehresmann connections so that the horizontal lift of any path in M is uniquely determined once the starting point on E is specified. Therefore, given a smooth curve $\gamma : [0, T] \rightarrow M$ that connects $\gamma(0)$ to $\gamma(T)$ on M , there exists a smooth map from $F_{\gamma(0)}$ to $F_{\gamma(T)}$ (at least when $\gamma(0)$ and $\gamma(T)$ are sufficiently close), defined as

$$F_{\gamma(0)} \ni e \mapsto \bar{\gamma}_e(T) \in F_{\gamma(T)},$$

where $\bar{\gamma}_e$ denotes the horizontal lift of γ with starting point p . We call this construction of maps between fibres, obviously depending on the choice of path γ , the *parallel transport along γ (with respect to the connection)*, and denote $P_{yx}^\gamma : F_x \rightarrow F_y$ for the parallel transport from fibre F_x to fibre F_y . When γ is a unique geodesic on M that connects x to y , we drop the superscript γ and simply write $P_{yx} : F_x \rightarrow F_y$. For future reference, we give the precise definition of the operator P_{yx} here.

Definition 7.1 (*Parallel transport on fibre bundles*). Let $\mathcal{E} = (E, M, F, \pi)$ be a fibre bundle, $x \in M$, $v \in F_x$, and U a geodesic normal neighborhood of x on the base manifold M . For any $y \in U$, denote the geodesic distance between x and y as $d_M(x, y)$. Let $\gamma : [0, d_M(x, y)] \rightarrow M$ be the unique unit-speed geodesic on M connecting x to y , i.e., $\gamma(0) = x$, $\gamma(d_M(x, y)) = y$; let $\bar{\gamma}$ be the unique horizontal lift of γ starting at $v \in F_x$, i.e.,

$$\begin{cases} \bar{\gamma}'(t) = \overline{\gamma}'(t), & t \in [0, d_M(x, y)], \\ \bar{\gamma}(0) = v. \end{cases}$$

The parallel-transport of v from x to y , denoted as $P_{yx}v$, is defined as

$$P_{yx}v = \bar{\gamma}(d_M(x, y)) \in F_y.$$

The probabilistic interpretation of HDM (and even VDM) implicitly depends on lifting from the base manifold a path that is continuous but not necessarily smooth. Though this can not be trivially achieved by the ODE-based approach, stochastic differential geometers developed tools appropriate for tackling this technicality (see e.g. [98, §5.1.2]).

Appendix B. Horizontal and vertical Laplacians

Assume (M, g^M) is a d -dimensional Riemannian manifold, and denote ∇^M for the canonical Levi-Civita connection on M . The *Laplace-Beltrami operator* on M , or *Laplacian* for short, is the analogy of the usual Laplace operator on the Euclidean space defined by

$$\Delta_M f(x) = \text{Trace } \nabla^M \nabla f(x)$$

for all $f \in C^\infty(M)$, $x \in M$. For an orthonormal local frame $\{X_1, \dots, X_d\}$ near $x \in M$, Δ_M can also be written as

$$\Delta_M f(x) = \sum_{j=1}^d g^M \left(\nabla_{X_j}^M \nabla f, X_j \right) (x) = \sum_{j=1}^d X_j^2 f(x) - \left(\sum_{j=1}^d \nabla_{X_j}^M X_j(x) \right) f(x). \quad (\text{B.1})$$

If we further pick the frame to be a local geodesic frame centered at $x \in M$, then $\nabla_{X_j}^M X_k(x) = 0$ for all $1 \leq j, k \leq d$ and thus Δ_M takes the following sum-of-squares form

$$\Delta_M f(x) = \sum_{j=1}^d X_j^2 f(x). \quad (\text{B.2})$$

The infinitesimal generator of the horizontal diffusion (2.2) turns out to be a differential operator on E that is a “horizontal lift” of Δ_M in a sense to be made clear later in this section. To characterize this infinitesimal

generator, let us first introduce a Riemannian metric on E that is adapted to the connection HE . For any $x \in M$, recall from Section 2.1 that F_x (the fibre at point $x \in M$) is a Riemannian submanifold of E , thus vertical tangent vectors at $e \in \pi^{-1}(x)$ can be canonically identified with tangent vectors to F_x ; if each F_x is equipped with a Riemannian metric g^{F_x} , we define for any $U, V \in V_e E$

$$g^E(U, V) = g^{F_x}(U, V). \quad (\text{B.3})$$

For any $X, Y \in H_e E$, by the linear isomorphism between $H_e E$ and $T_x M$ we define

$$g^E(X, Y) = g^M(d\pi_e(X), d\pi_e(Y)) \quad (\text{B.4})$$

where g^M stands for the Riemannian metric on M . Finally, impose orthogonality between $H_e E$ and $V_e E$ by setting for any $X \in H_e E$, $U \in V_e E$

$$g^E(X, U) = 0. \quad (\text{B.5})$$

The smoothness of g^E with respect to $e \in E$ follows from the smoothness of g^M and g^{F_x} . In other words, g^E is constructed so as to make the decomposition $TE = HE \oplus VE$ orthogonal. Some authors [66,99] abbreviate this construction as

$$g^E = g^M \oplus g^F. \quad (\text{B.6})$$

For future convenience, let us use superscripts to denote the horizontal and vertical components of tangent vectors to E , i.e. for any $Z \in T_e E$

$$Z = Z^H + Z^V$$

where $Z^H \in H_e E$, $Z^V \in V_e E$ are uniquely determined due to the direct sum decomposition $T_e E = H_e E \oplus V_e E$. Thus for any $W, Z \in T_e E$

$$g^E(W, Z) = g^M(d\pi_e(W^H), d\pi_e(Z^H)) + g^{F_{\pi(e)}}(W^V, Z^V).$$

We also write the horizontal and vertical components of the gradient of any smooth function $f \in C^\infty(E)$ as

$$\nabla^H f := (\nabla f)^H, \quad \nabla^V f := (\nabla f)^V. \quad (\text{B.7})$$

Let ∇^E denote the Levi-Civita connection with respect to g^E . Define the *rough horizontal Laplacian* Δ_H on E for $f \in C^\infty(E)$ as the following second order partial differential operator:

$$\Delta_H f(e) = \text{Trace}(\nabla^E \nabla^H f)^H(e) \quad \text{for all } e \in E. \quad (\text{B.8})$$

Let $\{\bar{X}_1, \dots, \bar{X}_d\}$ be the horizontal lift of an orthonormal frame $\{X_1, \dots, X_d\}$ near $\pi(e) = x \in M$. Since $g^E(\bar{X}_j, \bar{X}_k) = g^M(X_j, X_k)$ for $1 \leq j, k \leq d$, the tangent vectors $\bar{X}_1(e'), \dots, \bar{X}_d(e')$ form an orthonormal basis for $H_{e'} E$ for all e' sufficiently close to e . We can write (B.8) in terms of these horizontally lifted vector fields as

$$\begin{aligned} \Delta_H f(e) &= \sum_{j=1}^d g^E \left(\nabla_{\bar{X}_j}^E \nabla^H f, \bar{X}_j \right) (e) \\ &= \sum_{j=1}^d \bar{X}_j^2 f(e) - \left(\sum_{j=1}^d \left(\nabla_{\bar{X}_j}^E \bar{X}_j \right)^H \right) f(e). \end{aligned} \quad (\text{B.9})$$

Loosely speaking, Δ_H is the “horizontal lift” of Δ_M from M to E , since (B.9) can be obtained from (B.1) by replacing each X_j with its horizontal lift \bar{X}_j and noting that $\left(\nabla_{\bar{X}_j}^E \bar{X}_j\right)^H$ is the horizontal lift of $\nabla_{X_j}^M X_j$ (see e.g. [56, Proposition 3.1]). More precisely, for any $g \in C^\infty(M)$, denote $\bar{g} = g \circ \pi \in C^\infty(E)$, then for any $e \in E$ and $x = \pi(e) \in M$ we have

$$\Delta_M g(x) = \Delta_H \bar{g}(e). \quad (\text{B.10})$$

Remark 7.1. When $E = \mathcal{O}(M)$ is the frame bundle of M , the rough horizontal Laplacian Δ_H coincides with the *Bochner horizontal Laplacian* $\Delta_{\mathcal{O}(M)}$ in stochastic differential geometry [59, Chapter 3]. The classical Eells-Elworthy-Malliavin approach intrinsically defines a Brownian motion on manifolds as a *horizontal Brownian motion* on $\mathcal{O}(M)$ generated by $\Delta_{\mathcal{O}(M)}$.

Remark 7.2. In general, the rough horizontal Laplacian Δ_H differs from the concept of “horizontal Laplacian” commonly seen in sub-Riemannian geometry by a mean curvature term [99,100]; the two types of horizontal Laplacian coincide only when the fibres of E are minimal submanifolds of E . In fact, for any $f \in C^\infty(E)$, the Laplace-Beltrami operator on E with respect to g^E splits into two parts

$$\Delta_E f = \text{Trace } \nabla^E \nabla f = \text{Trace } \nabla^E \nabla^H f + \text{Trace } \nabla^E \nabla^V f$$

Define the *horizontal Laplacian* Δ_E^H and the *vertical Laplacian* Δ_E^V as

$$\Delta_E^H f := \text{Trace } \nabla^E \nabla^H f, \quad \Delta_E^V f := \text{Trace } \nabla^E \nabla^V f, \quad (\text{B.11})$$

then

$$\Delta_E = \Delta_E^H + \Delta_E^V. \quad (\text{B.12})$$

Recalling the definition of Δ_H from (B.8), we have

$$\begin{aligned} \Delta_E^H f &= \text{Trace } (\nabla^E \nabla^H f)^H + \text{Trace } (\nabla^E \nabla^H f)^V \\ &= \Delta_H f + \text{Trace } (\nabla^E \nabla^H f)^V \end{aligned}$$

and $\Delta_E^H = \Delta_H$ if and only if

$$\text{Trace } (\nabla^E \nabla^H f)^V = 0 \quad \text{for all } f \in C^\infty(E)$$

which turns out to be equivalent to the requirement that F_x are minimal submanifolds of E for all $x \in M$. This holds, for instance, when all fibres of the Riemannian submersion $\pi : E \rightarrow M$ are totally geodesic, a scenario of great theoretic interest since it implies that all fibres are isometric [56]; we do not make such an assumption in the HDM framework since this particularly simple case is obviously too restricted for practical purposes.

Remark 7.3. For any $x \in M$ and $e \in \pi^{-1}(x)$, if $\{X_1, \dots, X_d\}$ is a geodesic frame on M near x , then the horizontal lifts $\{\bar{X}_1, \dots, \bar{X}_d\}$ near e also constitute a “horizontal geodesic frame” in the sense that

$$\left(\nabla_{\bar{X}_j}^E \bar{X}_k\right)^H(e) = 0, \quad \text{for all } 1 \leq j, k \leq d,$$

which simplifies (B.9) into a sum-of-squares form analogous to (B.2)

$$\Delta_H f(e) = \sum_{j=1}^d \bar{X}_j^2 f(e) \quad \text{for all } f \in C^\infty(E). \quad (\text{B.13})$$

Remark 7.4. We make the observation that the vertical Laplacian Δ_E^V , which turns out to characterize the “vertical component” of the coupled diffusion operator $H_{\epsilon,\delta}^{(\alpha)}$ on the fibre bundle, coincides with the Laplace-Beltrami operator on each fibre F_x . This fact will be needed in the proof of Theorem 4 in Appendix C. More precisely, for any $f \in C^\infty(E)$ and $e \in E$,

$$\Delta_E^V f(e) = [\Delta_{F_{\pi(e)}}(f \upharpoonright F_{\pi(e)})](e) = \text{Trace}(\nabla^E \nabla^V f)^V. \quad (\text{B.14})$$

At a first glance this might seem a bit surprising since one may expect a mean curvature term in Δ_E^V from (B.11) (as is the case for Δ_E^H):

$$\Delta_E^V f = \text{Trace} \nabla^E \nabla^V f = \text{Trace}(\nabla^E \nabla^V f)^H + \text{Trace}(\nabla^E \nabla^V f)^V. \quad (\text{B.15})$$

However, the first trace term in (B.15) vanishes for the following reason. Let $\{X_1, \dots, X_d\}$ be a local horizontal orthonormal frame around $e \in E$, and $\{U_1, \dots, U_n\}$ a local vertical orthonormal frame (recall that $\dim(F) = n$); $\{X_1, \dots, X_d, U_1, \dots, U_n\}$ is then a local orthonormal frame on E . We have

$$\begin{aligned} \text{Trace}(\nabla^E \nabla^V f)^H &= \sum_{j=1}^d \langle (\nabla_{X_j}^E \nabla^V f)^H, X_j \rangle + \sum_{k=1}^n \langle (\nabla_{U_k}^E \nabla^V f)^H, U_k \rangle = \sum_{j=1}^d \langle (\nabla_{X_j}^E \nabla^V f)^H, X_j \rangle \\ &= \sum_{j=1}^d \langle \nabla_{X_j}^E \nabla^V f, X_j \rangle = \langle \nabla^V f, \sum_{j=1}^d -\nabla_{X_j}^E X_j \rangle = \langle \nabla^V f, \sum_{j=1}^d -(\nabla_{X_j}^E X_j)^V \rangle = 0, \end{aligned}$$

where the last equality follows from [101, Lemma 2]:

$$(\nabla_{X_j}^E X_j)^V = \frac{1}{2} ([X_j, X_j])^V = 0 \quad \text{for all } 1 \leq j \leq d.$$

In the remaining section we consider horizontal and coupled diffusion operators on a few classical examples. All fibre bundles in this section are Riemannian submersions with totally geodesic fibres, for which, as explained in Remark 7.2, the rough horizontal Laplacian Δ_H equals to the horizontal Laplacian Δ_E^H . See [99] [40, §9.F] for more details about Riemannian submersions with totally geodesic fibres.

Example 11 (Heisenberg group). The Heisenberg group

$$\mathbb{H}^{2n+1} = \{(x, y, z) \in \mathbb{R}^{2n+1} \mid x \in \mathbb{R}^n, y \in \mathbb{R}^n, z \in \mathbb{R}\}$$

is essentially \mathbb{R}^{2n+1} endowed with the following group structure:

$$(x_1, y_1, z_1) \cdot (x_2, y_2, z_2) = \left(x_1 + x_2, y_1 + y_2, z_1 + z_2 + \frac{1}{2} (x_1 \cdot y_2 - x_2 \cdot y_1) \right).$$

The projection

$$\begin{aligned} \pi : \mathbb{H}^{2n+1} &\longrightarrow \mathbb{R}^2 \\ (x, y, z) &\longmapsto (x, y) \end{aligned}$$

is a Riemannian submersion with totally geodesic fibres [99]. Since \mathbb{H}^{2n+1} is complete, it follows from [56, Theorem 1] that $(\mathbb{H}^{2n+1}, \mathbb{R}^2, \mathbb{R}, \pi)$ is a fibre bundle. In fact, \mathbb{H}^{2n+1} is a Lie group, and its Lie algebra of left invariant vector fields at (x, y, z) is spanned by

$$\frac{\partial}{\partial z}, \quad \frac{\partial}{\partial x_j} - \frac{1}{2} y_j \frac{\partial}{\partial z}, \quad \frac{\partial}{\partial y_j} - \frac{1}{2} x_j \frac{\partial}{\partial z}, \quad j = 1, \dots, n.$$

These invariant vector fields define a connection on \mathbb{H}^{2n+1} in the sense of Ehresmann [97]. The horizontal and vertical Laplacians on \mathbb{H}^{2n+1} with respect to this connection are

$$\begin{aligned} \Delta_{\mathbb{H}^{2n+1}}^H &= \sum_{j=1}^n \left[\frac{\partial^2}{\partial x_j^2} + \frac{\partial^2}{\partial y_j^2} + \frac{1}{4} (x_j^2 + y_j^2) \frac{\partial^2}{\partial z^2} - y_j \frac{\partial^2}{\partial x_j \partial z} + x_j \frac{\partial^2}{\partial y_j \partial z} \right], \\ \Delta_{\mathbb{H}^{2n+1}}^V &= \frac{\partial^2}{\partial z^2}. \end{aligned}$$

By Theorem 4, for any $f \in C^\infty(\mathbb{H}^{2n+1})$, if $\delta = O(\epsilon)$,

$$\begin{aligned} H_{\epsilon, \delta}^{(\alpha)} f(x, v) &= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \frac{[\Delta_{\mathbb{H}^{2n+1}}^H(f p^{1-\alpha}) - f \Delta_{\mathbb{H}^{2n+1}}^H p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} \\ &\quad + \delta \frac{m_{22}}{2m_0} \frac{[\Delta_{\mathbb{H}^{2n+1}}^V(f p^{1-\alpha}) - f \Delta_{\mathbb{H}^{2n+1}}^V p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} + O(\epsilon^2 + \epsilon \delta + \delta^2). \end{aligned}$$

When $n = 1$, this is consistent with the conclusion obtained in [71, Chapter 4].

Example 12 (Tangent bundles). Tangent bundles play an important role in Riemannian geometry. For a closed d -dimensional Riemannian manifold (M, g) , its tangent bundle TM is defined as

$$TM = \coprod_{x \in M} T_x M$$

equipped with a natural smooth structure (see e.g. [102]). In a local coordinate chart $(U; x_1, \dots, x_d)$ of M , $\{E_j = \partial/\partial x_j \mid 1 \leq j \leq d\}$ is a local frame on M , and we write $v \in T_x M$ as $v = v_j E_j(x)$. A local trivialization on U can be chosen as

$$(x, v) \mapsto (x_1, \dots, x_d, v_1, \dots, v_d), \quad \forall x \in U, v \in T_x M,$$

and the corresponding basis for $T_{(x, v)} TM$ can be written as

$$\left\{ \frac{\partial}{\partial x_1} \Big|_{(x, v)}, \dots, \frac{\partial}{\partial x_d} \Big|_{(x, v)}, \frac{\partial}{\partial v_1} \Big|_{(x, v)}, \dots, \frac{\partial}{\partial v_d} \Big|_{(x, v)} \right\}.$$

Let $\Gamma_{\alpha j}^\beta$ be the connection coefficients of the Levi-Civita connection on M . The horizontal subbundle of TTM determined by this connection is

$$HTM := \coprod_{(x, v) \in TM} HT_{(x, v)} M = \coprod_{(x, v) \in TM} \text{span} \left\{ \frac{\partial}{\partial x_j} \Big|_{(x, v)} - \Gamma_{\alpha j}^\beta(x) v_\alpha \frac{\partial}{\partial v_\beta} \Big|_{(x, v)}, \quad j = 1, \dots, d \right\}.$$

The metric (B.6) on TM given by this construction is the *Sasaki metric* [103, 104]. The horizontal and vertical Laplacians acts on any $f \in C^\infty(TM)$ as

$$\begin{aligned}\Delta_{TM}^H f(x, v) &= \frac{1}{\sqrt{|g(x)|}} \left(\frac{\partial}{\partial x_j} - \Gamma_{\alpha j}^\beta(x) v_\alpha \frac{\partial}{\partial v_\beta} \right) \left[\sqrt{|g(x)|} g^{jk}(x) \left(\frac{\partial f}{\partial x_k} - \Gamma_{\alpha k}^\beta(x) v_\alpha \frac{\partial f}{\partial v_\beta} \right) \right], \\ \Delta_{TM}^V f(x, v) &= \frac{1}{\sqrt{|g(x)|}} \frac{\partial}{\partial v_j} \left(\sqrt{|g(x)|} g^{jk}(x) \frac{\partial f}{\partial v_k} \right) = g^{jk}(x) \frac{\partial^2 f}{\partial v_j \partial v_k}.\end{aligned}$$

According to Theorem 4, for any $f \in C^\infty(TM)$, if $\delta = O(\epsilon)$,

$$\begin{aligned}H_{\epsilon, \delta}^{(\alpha)} f(x, v) &= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \frac{[\Delta_{TM}^H (fp^{1-\alpha}) - f\Delta_{TM}^H p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} \\ &\quad + \delta \frac{m_{22}}{2m_0} \frac{[\Delta_{TM}^V (fp^{1-\alpha}) - f\Delta_{TM}^V p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} + O(\epsilon^2 + \epsilon\delta + \delta^2).\end{aligned}$$

This is consistent with the conclusion obtained in [71, Chapter 3].

Appendix C. Proofs of Theorem 2, Theorem 4, and Theorem 6

Throughout this appendix we assume the Einstein summation convention unless otherwise specified. Our starting point is the following lemma, in reminiscent of [7, Lemma 8] and [9, Lemma B.10].

Lemma 13. *Let $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ be a smooth function compactly supported in $[0, 1]$. Assume M is a d -dimensional compact Riemannian manifold without boundary, with injectivity radius $\text{Inj}(M) > 0$. For any $\epsilon > 0$, define kernel function*

$$\Phi_\epsilon(x, y) = \Phi\left(\frac{d_M^2(x, y)}{\epsilon}\right) \quad (\text{C.1})$$

on $M \times M$, where $d_M^2(\cdot, \cdot)$ is the geodesic distance on M . For sufficiently small ϵ satisfying $0 \leq \epsilon \leq \sqrt{\text{Inj}(M)}$, the integral operator associated with kernel Φ_ϵ

$$(\Phi_\epsilon g)(x) := \int_M \Phi_\epsilon(x, y) g(y) d\text{vol}_M(y) \quad (\text{C.2})$$

has the following asymptotic expansion as $\epsilon \rightarrow 0$:

$$(\Phi_\epsilon g)(x) = \epsilon^{\frac{d}{2}} \left[m_0 g(x) + \epsilon \frac{m_2}{2} \left(\Delta_M g(x) - \frac{1}{3} \text{Scal}^M(x) g(x) \right) + O(\epsilon^2) \right], \quad (\text{C.3})$$

where m_0, m_2 are constants that depend on the moments of Φ and the dimension d of the Riemannian manifold M , Δ_M is the Laplace-Beltrami operator on M , and $\text{Scal}^M(x)$ is the scalar curvature of M at x .

Proof. Consider geodesic normal coordinates near $x \in M$. Let $\{e_1, \dots, e_d\}$ be an orthonormal basis for $T_x M$, (s_1, \dots, s_d) the geodesic normal coordinates, and write $r = d_M(x, y)$. Then $r^2 = s_1^2 + \dots + s_d^2$. Note that

$$\begin{aligned}\int_M \Phi_\epsilon(x, y) g(y) d\text{vol}_M(y) &= \int_{B_{\sqrt{\epsilon}}(x)} \Phi\left(\frac{d_M^2(x, y)}{\epsilon}\right) g(y) d\text{vol}_M(y) \\ &= \int_{B_{\sqrt{\epsilon}}(0)} \Phi\left(\frac{r^2}{\epsilon}\right) \tilde{g}(s) d\text{vol}_M(s)\end{aligned} \quad (\text{C.4})$$

where

$$\tilde{g}(s) = \tilde{g}(s_1, \dots, s_d) = g \circ \exp_x(s_1 e_1 + \dots + s_d e_d),$$

$$d\text{vol}_M(s) = d\text{vol}_M(s_1, \dots, s_d) = d\text{vol}_M(\exp_x(s_1 e_1 + \dots + s_d e_d)).$$

By a further change of variables

$$\tilde{s}_1 = \frac{s_1}{\sqrt{\epsilon}}, \dots, \tilde{s}_d = \frac{s_d}{\sqrt{\epsilon}}; \quad \tilde{r} = \frac{r}{\sqrt{\epsilon}}, \quad (\text{C.5})$$

we have

$$\int_{B_{\sqrt{\epsilon}}(0)} \Phi\left(\frac{r^2}{\epsilon}\right) \tilde{g}(s) d\text{vol}_M(s) = \int_{B_1(0)} \Phi(\tilde{r}^2) \tilde{g}(\sqrt{\epsilon} \tilde{s}) d\text{vol}_M(\sqrt{\epsilon} \tilde{s}). \quad (\text{C.6})$$

On the other hand, in geodesic normal coordinates the Riemannian volume form has asymptotic expansion (see e.g. [105])

$$d\text{vol}_M(s_1, \dots, s_d) = \left[1 - \frac{1}{6} R_{k\ell}(x) s_k s_\ell + O(r^3)\right] ds_1 \cdots ds_d \quad (\text{C.7})$$

where $R_{k\ell}$ is the Ricci curvature tensor

$$R_{k\ell}(x) = g^{ij} R_{kij\ell}(x).$$

Thus

$$d\text{vol}_M(\sqrt{\epsilon} \tilde{s}_1, \dots, \sqrt{\epsilon} \tilde{s}_d) = \left[1 - \frac{\epsilon}{6} R_{k\ell}(x) \tilde{s}_k \tilde{s}_\ell + O(\epsilon^{\frac{3}{2}} \tilde{r}^3)\right] \cdot \epsilon^{\frac{d}{2}} d\tilde{s}_1 \cdots d\tilde{s}_d. \quad (\text{C.8})$$

In the meanwhile, the Taylor expansion of $\tilde{g}(s)$ near x reads

$$\tilde{g}(s) = \tilde{g}(0) + \frac{\partial \tilde{g}}{\partial s_j}(0) s_j + \frac{1}{2} \frac{\partial^2 \tilde{g}}{\partial s_k \partial s_\ell}(0) s_k s_\ell + O(r^3)$$

and thus

$$\tilde{g}(\sqrt{\epsilon} \tilde{s}) = g(x) + \sqrt{\epsilon} \cdot \frac{\partial \tilde{g}}{\partial s_j}(0) \tilde{s}_j + \epsilon \cdot \frac{1}{2} \frac{\partial^2 \tilde{g}}{\partial s_k \partial s_\ell}(0) \tilde{s}_k \tilde{s}_\ell + O(\epsilon^{\frac{3}{2}} \tilde{r}^3). \quad (\text{C.9})$$

By the symmetry of the kernel and the domain of integration $B_1(0)$,

$$\int_{B_1(0)} \Phi(\tilde{r}^2) \tilde{s}_j d\tilde{s}_1 \cdots d\tilde{s}_d = 0 \quad \text{for all } 1 \leq j \leq d, \quad (\text{C.10})$$

$$\int_{B_1(0)} \Phi(\tilde{r}^2) \tilde{s}_k \tilde{s}_\ell d\tilde{s}_1 \cdots d\tilde{s}_d = 0 \quad \text{for all } 1 \leq k \neq \ell \leq d.$$

Combining (C.4)–(C.10), we have

$$\begin{aligned} \int_M \Phi_\epsilon(x, y) g(y) d\text{vol}_M(y) &= \int_{B_1(0)} \Phi(\tilde{r}^2) \tilde{g}(\sqrt{\epsilon} \tilde{s}_1, \dots, \sqrt{\epsilon} \tilde{s}_d) d\text{vol}_M(\sqrt{\epsilon} \tilde{s}) \\ &= \epsilon^{\frac{d}{2}} \left[g(x) \int_{B_1(0)} \Phi(\tilde{r}^2) d\tilde{s} + \frac{\epsilon}{2} \sum_{k=1}^d \left(\frac{\partial^2 \tilde{g}}{\partial s_k^2}(0) - \frac{1}{3} g(x) R_{kk}(x) \right) \int_{B_1(0)} \Phi(\tilde{r}^2) \tilde{s}_k^2 d\tilde{s} + O(\epsilon^2) \right] \end{aligned}$$

Note that $O(\epsilon^{3/2})$ term vanishes again by symmetry (the same argument given in [72, §2] applies). Define constants

$$\begin{aligned} m_0 &:= \int_{B_1(0)} \Phi(\tilde{r}^2) d\tilde{s}_1 \cdots d\tilde{s}_d = \omega^{d-1} \int_0^1 \Phi(\tilde{r}^2) \tilde{r}^{d-1} d\tilde{r}, \\ m_2 &:= \int_{B_1(0)} \Phi(\tilde{r}^2) (\tilde{s}_k)^2 d\tilde{s}_1 \cdots d\tilde{s}_d \quad \text{for any } k \in \{1, \dots, d\}. \end{aligned} \tag{C.11}$$

Then

$$\begin{aligned} \int_M \Phi_\epsilon(x, y) g(y) d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \left[m_0 g(x) + \epsilon \frac{m_2}{2} \left(\Delta_M g(x) - \frac{1}{3} \text{Scal}^M(x) g(x) \right) + O(\epsilon^2) \right], \end{aligned}$$

where we used the fact that in geodesic normal coordinates

$$\sum_{k=1}^d \frac{\partial^2 \tilde{g}}{\partial s_k^2}(0) = \Delta_M g(x), \quad \sum_{k=1}^d R_{kk}(x) = \text{Scal}^M(x). \quad \square$$

Before applying Lemma 13 to compute the infinitesimal generators of $H_\epsilon^{(\alpha)}$ and $H_{\epsilon, \delta}^{(\alpha)}$, we need more local information about $f(x, P_{yx}v)$ near (x, v) . To this end, let $\{X_1, \dots, X_d\}$ be a local geodesic frame on U at x , and denote $\{\bar{X}_1, \dots, \bar{X}_d\}$ for the horizontal lift of this frame; in addition, let $\{V_1, \dots, V_n\}$ be vertical vector fields on E such that

$$\{\bar{X}_1(e), \dots, \bar{X}_d(e), V_1(e), \dots, V_n(e)\}$$

constitutes an orthonormal basis for all e in a sufficiently small neighborhood of (x, v) contained in $\pi^{-1}(U)$. Write $\{\theta^1, \dots, \theta^d, \phi^1, \dots, \phi^n\}$ for the 1-forms dual to the vector fields $\{\bar{X}_1, \dots, \bar{X}_d, V_1, \dots, V_n\}$, i.e.,

$$\begin{aligned} \theta^j(\bar{X}_k) &= \delta_k^j, \quad \theta^j(V_\ell) = 0, \\ \phi^m(\bar{X}_k) &= 0, \quad \phi^m(V_\ell) = \delta_\ell^m, \end{aligned}$$

for all $1 \leq j, k \leq d, 1 \leq \ell, m \leq n$.

If γ is a unit speed geodesic on M starting at x , recall from Definition 7.1 that $t \mapsto P_{\gamma(t), x}v$ is the unique horizontal lift of γ with starting point $v \in F_x$, i.e.,

$$\bar{\gamma}(t) = P_{\gamma(t), x}v.$$

Since $\bar{\gamma}$ is horizontal, $\phi^m(\bar{\gamma}'(t)) = 0$ for all $1 \leq m \leq n$ and thus (adopting Einstein summation convention)

$$\bar{\gamma}'(t) = \theta^j(\bar{\gamma}'(t)) \bar{X}_j(t). \quad (\text{C.12})$$

Here, as well as in the rest of this appendix, we set

$$X_j(t) = X_j(\gamma(t)), \quad \bar{X}_j(t) = \bar{X}_j(\gamma(t)).$$

By [56, Proposition 3.1], $\bar{\gamma}(t)$ is a geodesic on E , thus

$$0 = \nabla_{\bar{\gamma}'(t)}^E \bar{\gamma}'(t) = \frac{d}{dt} [\theta^j(\bar{\gamma}'(t))] \bar{X}_j(t) + \theta^j(\bar{\gamma}'(t)) \theta^k(\bar{\gamma}'(t)) \nabla_{\bar{X}_k(t)}^E \bar{X}_j(t),$$

which implies

$$\frac{d}{dt} [\theta^j(\bar{\gamma}'(t))] \bar{X}_j(t) = -\theta^j(\bar{\gamma}'(t)) \theta^k(\bar{\gamma}'(t)) \nabla_{\bar{X}_k(t)}^E \bar{X}_j(t). \quad (\text{C.13})$$

In particular, the right hand side of (C.13) is horizontal. It follows that

$$\begin{aligned} \frac{d}{dt} [\theta^j(\bar{\gamma}'(t))] \bar{X}_j(t) &= -\theta^j(\bar{\gamma}'(t)) \theta^k(\bar{\gamma}'(t)) \mathcal{H} \nabla_{\bar{X}_k(t)}^E \bar{X}_j(t) \\ &= -\theta^j(\bar{\gamma}'(t)) \theta^k(\bar{\gamma}'(t)) \langle \nabla_{\bar{X}_k(t)}^E \bar{X}_j(t), \bar{X}_i(t) \rangle_{\bar{\gamma}(t)} \bar{X}_i(t) \\ &= -\theta^j(\bar{\gamma}'(t)) \theta^k(\bar{\gamma}'(t)) \langle \nabla_{X_k(t)}^M X_j(t), X_i(t) \rangle_{\gamma(t)} \bar{X}_i(t), \end{aligned}$$

where \mathcal{H} is the horizontal projection as defined in (A.1). By linear independence,

$$\frac{d}{dt} [\theta^j(\bar{\gamma}'(t))] = -\theta^i(\bar{\gamma}'(t)) \theta^k(\bar{\gamma}'(t)) \Gamma_{ik}^j(\gamma(t)), \quad (\text{C.14})$$

where Γ_{ik}^j are the connection coefficients for the frame $\{X_1, \dots, X_d\}$ on M

$$\Gamma_{ik}^j = \langle \nabla_{X_k}^M X_j, X_i \rangle, \quad \forall 1 \leq i, j, k \leq d.$$

Setting $t = 0$ in (C.14) to get

$$\frac{d}{dt} \Big|_{t=0} [\theta^j(\bar{\gamma}'(t))] = -\theta^i(\bar{\gamma}'(0)) \theta^k(\bar{\gamma}'(0)) \Gamma_{ik}^j(\gamma(0)) = 0 \quad (\text{C.15})$$

where $\Gamma_{ik}^j(x) = 0$ since we picked $\{X_j \mid 1 \leq j \leq d\}$ as a geodesic frame at x .

Now for any $f \in C^\infty(E)$ write

$$f(t) := f(\bar{\gamma}(t)) = f(\gamma(t), P_{\gamma(t), x} v).$$

Using (C.12) and (C.15), the first and second derivatives of $f(t)$ at $t = 0$ can be written as

$$\begin{aligned} f'(0) &= \theta^j(\bar{\gamma}'(0)) \bar{X}_j f(0), \\ f''(0) &= \frac{d}{dt} \Big|_{t=0} [\theta^j(\bar{\gamma}'(t))] \bar{X}_j f(t) + \theta^i(\bar{\gamma}'(0)) \theta^k(\bar{\gamma}'(0)) \bar{X}_k \bar{X}_j f(0) \\ &= \theta^i(\bar{\gamma}'(0)) \theta^k(\bar{\gamma}'(0)) \bar{X}_k \bar{X}_i f(0). \end{aligned}$$

Furthermore, if we denote $\pi^* : \Lambda^* M \rightarrow \Lambda^* E$ for the pullback map, and write $\{\psi^j \mid 1 \leq j \leq d\}$ for the dual 1-forms to the geodesic frame $\{X_j \mid 1 \leq j \leq d\}$ on M , then $\theta^j = \pi^* \psi^j$ for all $1 \leq j \leq d$ and

$$\theta^j(\bar{\gamma}'(0)) = \pi^* \psi^j(\bar{\gamma}'(0)) = \psi^j(\gamma'(0)).$$

Thus $(\theta^1(\bar{\gamma}'(0)), \dots, \theta^d(\bar{\gamma}'(0)))$ is $\gamma'(0)$ represented in the geodesic normal coordinate system associated with the geodesic frame $\{X_j \mid 1 \leq j \leq d\}$. If we write $\sigma_j = \theta^j(\bar{\gamma}'(0))$ and $s_j(t) = t\sigma_j$ for all $j = 1, \dots, d$, then $\sum_{j=1}^d \sigma_j^2 = 1$ and (s_1, \dots, s_d) are the geodesic coordinates of $\gamma(t)$ on M with respect to the geodesic frame $\{X_j \mid 1 \leq j \leq d\}$. With this notation,

$$\begin{aligned} f'(0) &= \sigma_j \bar{X}_j f(0), \\ f''(0) &= \sigma_i \sigma_k \bar{X}_i \bar{X}_k f(0). \end{aligned}$$

Using (C.12), (C.15), and $f''(t)$, it is straightforward to compute the third order derivative of f at $t = 0$:

$$f'''(0) = \sigma_i \sigma_j \sigma_k \bar{X}_i \bar{X}_j \bar{X}_k f(0),$$

hence the Taylor expansion of $f(t)$ near $t = 0$ is

$$\begin{aligned} f(t) &= f(0) + t f'(0) + \frac{t^2}{2} f''(0) + \frac{t^3}{6} f'''(0) + O(t^4) \\ &= f(x, v) + t \sigma_j \bar{X}_j f(x, v) + \frac{t^2}{2} \sigma_i \sigma_k \bar{X}_i \bar{X}_k f(x, v) + \frac{t^3}{6} \sigma_i \sigma_j \sigma_k \bar{X}_i \bar{X}_j \bar{X}_k f(x, v) + O(t^4). \end{aligned} \quad (\text{C.16})$$

This expansion immediately leads to the following lemma:

Lemma 14. *Following Lemma 13, let $P_{yx} : F_x \rightarrow F_y$ be as defined in Definition 7.1. For any $f \in C^\infty(E)$ and $v \in F_x$, as $\epsilon \rightarrow 0$,*

$$\begin{aligned} &\int_M \Phi_\epsilon(x, y) f(y, P_{yx} v) \, d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \left\{ m_0 f(x, v) + \epsilon \frac{m_2}{2} \left[\Delta_H f(x, v) - \frac{1}{3} \text{Scal}^M(x) f(x, v) \right] + O(\epsilon^2) \right\}, \end{aligned} \quad (\text{C.17})$$

where m_0, m_2 are constants, $\text{Scal}^M(x)$ is the scalar curvature of M at x , and Δ_H is the rough horizontal Laplacian on E defined in (B.8).

Proof. Let $U \subset M$ be a geodesic normal neighborhood around $x \in M$, and ϵ sufficiently small that any point in U can be connected to x with a geodesic of length less than $\epsilon^{1/2}$. Let $\{X_j \mid 1 \leq j \leq d\}$ be a geodesic frame on E , s_1, \dots, s_d geodesic normal coordinates on U with respect to this geodesic frame, and

$$r > 0, \quad r^2 = \sum_{j=1}^d s_j^2, \quad \sigma_j = \frac{s_j}{r} \quad \text{for all } 1 \leq j \leq d.$$

Following the proof of Lemma 13, let \tilde{s}_j, \tilde{r} be as defined in (C.5) and use (C.16) in place of (C.12),

$$\begin{aligned} \int_M \Phi_\epsilon(x, y) f(y, P_{yx} v) \, d\text{vol}_M(y) &= \int_{B_1(0)} \Phi(\tilde{r}^2) \tilde{f}(\sqrt{\epsilon} \tilde{s}, v) \, d\text{vol}_M(\sqrt{\epsilon} \tilde{s}) \\ &= f(x, v) \int_{B_1(0)} \Phi(\tilde{r}^2) \, d\text{vol}_M(\sqrt{\epsilon} \tilde{s}) + \epsilon^{\frac{1}{2}} \bar{X}_j f(x, v) \int_{B_1(0)} \tilde{s}_j \Phi(\tilde{r}^2) \, d\text{vol}_M(\sqrt{\epsilon} \tilde{s}) \end{aligned}$$

$$\begin{aligned}
& + \frac{\epsilon}{2} \bar{X}_i \bar{X}_k f(x, v) \int_{B_1(0)} \tilde{s}_i \tilde{s}_k \Phi(\tilde{r}^2) d\text{vol}_M(\sqrt{\epsilon} \tilde{s}) \\
& + \frac{\epsilon^{\frac{3}{2}}}{6} \bar{X}_i \bar{X}_j \bar{X}_k f(x, v) \int_{B_1(0)} \tilde{s}_i \tilde{s}_j \tilde{s}_k \Phi(\tilde{r}^2) d\text{vol}_M(\sqrt{\epsilon} \tilde{s}) + O(\epsilon^2).
\end{aligned}$$

Again by symmetry of these integrals and (C.7), this reduces to

$$\begin{aligned}
& \epsilon^{\frac{d}{2}} \left[m_0 f(x, v) - \frac{\epsilon}{6} m_2 \text{Scal}^M(x) f(x, v) + \frac{\epsilon}{2} m_2 \sum_{k=1}^d \bar{X}_k^2 f(x, v) + O(\epsilon^2) \right] \\
& = \epsilon^{\frac{d}{2}} \left[m_0 f(x, v) + \epsilon \frac{m_2}{2} \left(\Delta_H f(x, v) - \frac{1}{3} \text{Scal}^M(x) f(x, v) \right) + O(\epsilon^2) \right],
\end{aligned}$$

where m_0, m_2 are constants defined in (C.11), and

$$\sum_{k=1}^d \bar{X}_k^2 f(x, v) = \Delta_H f(x, v)$$

as explained in (B.13). \square

We are now ready to give the proof of Theorem 2.

Proof of Theorem 2. By the definition of $H_\epsilon^{(\alpha)}$ in (2.2), for any $f \in C^\infty(E)$,

$$\begin{aligned}
H_\epsilon^{(\alpha)} f(x, v) &= \frac{\int_M K_\epsilon^{(\alpha)}(x, y) f(y, P_{yx} v) p(y) d\text{vol}_M(y)}{\int_M K_\epsilon^{(\alpha)}(x, y) p(y) d\text{vol}_M(y)} \\
&= \frac{\int_M K_\epsilon(x, y) f(y, P_{yx} v) p(y) p_\epsilon^{-\alpha}(y) d\text{vol}_M(y)}{\int_M K_\epsilon(x, y) p(y) p_\epsilon^{-\alpha}(y) d\text{vol}_M(y)}.
\end{aligned}$$

By Lemma 13,

$$\begin{aligned}
p_\epsilon(y) &= \int_M K_\epsilon(x, y) p(\eta) d\text{vol}_M(\eta) \\
&= \epsilon^{\frac{d}{2}} \left\{ m_0 p(y) + \epsilon \frac{m_2}{2} \left(\Delta_M p(y) - \frac{1}{3} \text{Scal}^M(y) p(y) \right) + O(\epsilon^2) \right\}.
\end{aligned}$$

Using this expansion of p_ϵ and applying Lemma 13 to the denominator of $H_\epsilon^{(\alpha)}$,

$$\begin{aligned}
& \int_M K_\epsilon(x, y) p(y) p_\epsilon^{-\alpha}(y) d\text{vol}_M(y) \\
&= \epsilon^{\frac{(1-\alpha)d}{2}} m_0^{-\alpha} \left[m_0 p^{1-\alpha}(x) + \epsilon \frac{m_2}{2} \left(\Delta_M p^{1-\alpha}(x) - \frac{1}{3} \text{Scal}^M(x) p^{1-\alpha}(x) \right) \right]
\end{aligned}$$

$$- \alpha \epsilon \frac{m_2}{2} p^{-\alpha}(x) \left(\Delta_M p(x) - \frac{1}{3} \text{Scal}^M(x) p(x) \right) + O(\epsilon^2) \Big].$$

Similarly, apply Lemma 14 to the numerator of $H_\epsilon^{(\alpha)}$ to get

$$\begin{aligned} & \int_M K_\epsilon(x, y) f(y, P_{yx} v) p(y) p_\epsilon^{-\alpha}(y) d\text{vol}_M(y) \\ &= \epsilon^{\frac{(1-\alpha)d}{2}} m_0^{-\alpha} \left\{ m_0(f\bar{p}^{1-\alpha})(x, v) \right. \\ & \quad \left. + \epsilon \frac{m_2}{2} \left[\Delta_H(f\bar{p}^{1-\alpha})(x, v) - \frac{1}{3} \text{Scal}^M(x) (f\bar{p}^{1-\alpha})(x, v) \right] \right. \\ & \quad \left. - \alpha \epsilon \frac{m_2}{2} f(x, v) p^{-\alpha}(x) \left(\Delta_M p(x) - \frac{1}{3} \text{Scal}^M(x) p(x) \right) + O(\epsilon^2) \right\}. \end{aligned}$$

Noting that $\bar{p} = p \circ \pi$ and by (B.10)

$$\Delta_H \bar{p}^{1-\alpha} = \Delta_M p^{1-\alpha},$$

a direct computation (plus assumption (2.1) for the density p) concludes

$$H_\epsilon^{(\alpha)} f(x, v) = f(x, v) + \epsilon \frac{m_2}{2m_0} \frac{[\Delta_H(f\bar{p}^{1-\alpha}) - f\Delta_H \bar{p}^{1-\alpha}](x, v)}{\bar{p}^{1-\alpha}(x, v)} + O(\epsilon^2),$$

whence (4.1) follows. \square

We now turn to the proof of Theorem 4. The basic idea is to apply Lemma 13 and Lemma 14 repeatedly in both vertical and horizontal directions.

Lemma 15. *Suppose $\mathcal{E} = (E, M, F, \pi)$ is a fibre bundle, M is a smooth closed Riemannian manifold with $\text{Inj}(M) > 0$, and E equipped with the Riemannian metric (B.6). Assume $\dim M = d$ and $\dim F = n$. Let $K_{\epsilon, \delta}$ be defined as in (2.3) with $\epsilon \in (0, \text{Inj}(M)^2)$, $\delta = O(\epsilon)$. For any function $f \in C^\infty(E)$,*

$$\begin{aligned} & \iint_{M \times F_y} K_{\epsilon, \delta}(x, v; y, w) f(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \delta^{\frac{n}{2}} \left\{ m_0 f(x, v) + \epsilon \frac{m_{21}}{2} \left(\Delta_H f(x, v) - \frac{1}{3} \text{Scal}^M(x) f(x, v) \right) \right. \\ & \quad \left. + \delta \frac{m_{22}}{2} \left(\Delta_E^V f(x, v) - \frac{1}{3} \text{Scal}^{F_x}(v) f(x, v) \right) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}, \end{aligned} \tag{C.18}$$

where m_0, m_{21}, m_{22} are positive constants depending only on the kernel K and the fibre bundle, Scal^M , Scal^{F_x} are scalar curvatures of M , F_x respectively, and Δ_H , Δ_E^V are defined in (B.8) and (B.11).

Proof. By definition of $K_{\epsilon,\delta}$,

$$\begin{aligned} & \int_M \int_{F_y} K_{\epsilon,\delta}(x, v; y, w) f(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y) \\ &= \int_M \int_{F_y} K\left(\frac{d_M^2(x, y)}{\epsilon}, \frac{d_{F_y}^2(P_{yx}v, w)}{\delta}\right) f(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y). \end{aligned}$$

For any fixed $y \in M$, apply Lemma 13 to the inner integral over F_y with

$$\Phi(p, q) = K\left(\frac{d_M^2(x, y)}{\epsilon}, \frac{d_{F_y}^2(p, q)}{\delta}\right)$$

then the constants m_0, m_2 will depend on $d_M^2(x, y)/\epsilon$. More specifically, if we set

$$\begin{aligned} M_0(r^2) &= \int_{B_1^n(0)} K(r^2, \rho^2) d\theta_1 \cdots d\theta_n, \quad M_2(r^2) = \int_{B_1^n(0)} \theta_1^2 K(r^2, \rho^2) d\theta_1 \cdots d\theta_n, \\ M_3(r^2) &= \int_{B_1^n(0)} \theta_1^4 K(r^2, \rho^2) d\theta_1 \cdots d\theta_n, \quad \text{where } \rho^2 = \sum_{j=1}^n \theta_j^2 \end{aligned}$$

and recall from (B.14) that Δ_E^V coincides with Δ_{F_y} if one restricts a smooth function in $C^\infty(E)$ to F_y , then Lemma 13 leads to

$$\begin{aligned} & \int_{F_y} K_{\epsilon,\delta}(x, v; y, w) f(y, w) d\text{vol}_{F_y}(w) \\ &= \delta^{\frac{n}{2}} \left\{ M_0\left(\frac{d_M^2(x, y)}{\epsilon}\right) f(y, P_{yx}v) + \frac{\delta}{2} M_2\left(\frac{d_M^2(x, y)}{\epsilon}\right) \times \right. \\ & \quad \left. \left[\Delta_E^V f(y, P_{yx}v) - \frac{1}{3} \text{Scal}^{F_y}(P_{yx}v) f(y, P_{yx}v) \right] + O\left(\delta^2 M_3\left(\frac{d_M^2(x, y)}{\epsilon}\right)\right) \right\}. \end{aligned}$$

Now integrate over M and apply Lemma 14 multiple times:

$$\begin{aligned} & \int_M M_0\left(\frac{d_M^2(x, y)}{\epsilon}\right) f(y, P_{yx}v) d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \left\{ m_0 f(x, v) + \epsilon \frac{m_{21}}{2} \left(\Delta_H f(x, v) - \frac{1}{3} \text{Scal}^M(x) f(x, v) \right) + O(\epsilon^2) \right\}, \\ & \int_M M_2\left(\frac{d_M^2(x, y)}{\epsilon}\right) \left[\Delta_E^V f(y, P_{yx}v) - \frac{1}{3} \text{Scal}^{F_y}(P_{yx}v) f(y, P_{yx}v) \right] d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \left\{ m_{22} \left[\Delta_E^V f(x, v) - \frac{1}{3} \text{Scal}^{F_x}(v) f(x, v) \right] + O(\epsilon) \right\}, \end{aligned}$$

where the constants m_0, m_{21}, m_{22} are determined by (writing $r^2 = \sum_{j=1}^d s_j^2$)

$$m_0 = \int_{B_1^d(0)} M_0(r^2) ds_1 \cdots ds_d,$$

$$m_{21} = \int_{B_1^d(0)} M_0(r^2) s_1^2 ds_1 \cdots ds_d, \quad m_{22} = \int_{B_1^d(0)} M_2(r^2) ds_1 \cdots ds_d.$$

Therefore

$$\begin{aligned} & \int_M \int_{F_y} K_{\epsilon, \delta}(x, v; y, w) f(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \delta^{\frac{n}{2}} \left\{ m_0 f(x, v) + \epsilon \frac{m_{21}}{2} \left(\Delta_H f(x, v) - \frac{1}{3} \text{Scal}^M(x) f(x, v) \right) \right. \\ & \quad \left. + \delta \frac{m_{22}}{2} \left(\Delta_E^V f(x, v) - \frac{1}{3} \text{Scal}^{F_x}(v) f(x, v) \right) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}. \quad \square \end{aligned}$$

Proof of Theorem 4. Note that

$$\begin{aligned} H_{\epsilon, \delta}^{(\alpha)} f(x, v) &= \frac{\int_M \int_{F_y} K_{\epsilon, \delta}^{(\alpha)}(x, v; y, w) f(y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y)}{\int_M \int_{F_y} K_{\epsilon, \delta}^{(\alpha)}(x, v; y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y)} \\ &= \frac{\int_M \int_{F_y} K_{\epsilon, \delta}(x, v; y, w) f(y, w) p_{\epsilon, \delta}^{-\alpha}(y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y)}{\int_M \int_{F_y} K_{\epsilon, \delta}(x, v; y, w) p_{\epsilon, \delta}^{-\alpha}(y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y)}. \end{aligned} \quad (\text{C.19})$$

Applying Lemma 15 to $p_{\epsilon, \delta}$ to get

$$\begin{aligned} p_{\epsilon, \delta}(y, w) &= \epsilon^{\frac{d}{2}} \delta^{\frac{n}{2}} \left\{ m_0 p(y, w) + \epsilon \frac{m_{21}}{2} \left(\Delta_H p(y, w) - \frac{1}{3} \text{Scal}^M(y) p(y, w) \right) \right. \\ & \quad \left. + \delta \frac{m_{22}}{2} \left(\Delta_E^V p(y, w) - \frac{1}{3} \text{Scal}^{F_y}(w) p(y, w) \right) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}. \end{aligned}$$

Using this and applying Lemma 13, Lemma 14 to the denominator and numerator of (C.19) respectively:

$$\begin{aligned} & \int_M \int_{F_y} K_{\epsilon, \delta}(x, v; y, w) p_{\epsilon, \delta}^{-\alpha}(y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y) \\ &= \epsilon^{\frac{(1-\alpha)d}{2}} \delta^{\frac{(1-\alpha)n}{2}} m_0^{1-\alpha} p^{1-\alpha}(x, v) \left\{ 1 + \epsilon \frac{m_{21}}{2m_0} \left(\frac{\Delta_H p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} - \frac{1}{3} \text{Scal}^M(x) \right) \right. \\ & \quad + \delta \frac{m_{22}}{2m_0} \left(\frac{\Delta_E^V p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} - \frac{1}{3} \text{Scal}^{F_x}(x) \right) - \alpha \epsilon \frac{m_{21}}{2m_0}(x, v) \left(\frac{\Delta_H p(x, v)}{p(x, v)} - \frac{1}{3} \text{Scal}^M(x) \right) \\ & \quad \left. - \alpha \delta \frac{m_{22}}{2m_0} \left(\frac{\Delta_E^V p(x, v)}{p(x, v)} - \frac{1}{3} \text{Scal}^{F_x}(w) \right) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}, \end{aligned}$$

$$\int_M \int_{F_y} K_{\epsilon, \delta}(x, v; y, w) f(y, w) p_{\epsilon, \delta}^{-\alpha}(y, w) p(y, w) d\text{vol}_{F_y}(w) d\text{vol}_M(y)$$

$$\begin{aligned}
 &= \epsilon^{\frac{(1-\alpha)d}{2}} \delta^{\frac{(1-\alpha)n}{2}} m_0^{1-\alpha} (fp^{1-\alpha})(x, v) \left\{ 1 + \epsilon \frac{m_{21}}{2m_0} \left(\frac{\Delta_H(fp^{1-\alpha})(x, v)}{(fp^{1-\alpha})(x, v)} - \frac{1}{3} \text{Scal}^M(x) \right) \right. \\
 &+ \delta \frac{m_{22}}{2m_0} \left(\frac{\Delta_E^V(fp^{1-\alpha})(x, v)}{(fp^{1-\alpha})(x, v)} - \frac{1}{3} \text{Scal}^{F_x}(x) \right) - \alpha \epsilon \frac{m_{21}}{2m_0} (x, v) \left(\frac{\Delta_H p(x, v)}{p(x, v)} - \frac{1}{3} \text{Scal}^M(x) \right) \\
 &\left. - \alpha \delta \frac{m_{22}}{2m_0} \left(\frac{\Delta_E^V p(x, v)}{p(x, v)} - \frac{1}{3} \text{Scal}^{F_x}(w) \right) + O(\epsilon^2 + \epsilon \delta + \delta^2) \right\}.
 \end{aligned}$$

Combining these two expansions, a direct computation concludes

$$\begin{aligned}
 H_{\epsilon, \delta}^{(\alpha)} f(x, v) &= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \frac{[\Delta_H(fp^{1-\alpha}) - f \Delta_H p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} \\
 &+ \delta \frac{m_{22}}{2m_0} \frac{[\Delta_E^V(fp^{1-\alpha}) - f \Delta_E^V p^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} + O(\epsilon^2 + \epsilon \delta + \delta^2). \quad \square
 \end{aligned}$$

Proof of Theorem 6. Since $P_{y,x}v$ does not depend on γ ,

$$\begin{aligned}
 \lim_{\gamma \rightarrow \infty} p_{\epsilon, \gamma \epsilon}(x, v) &= \lim_{\gamma \rightarrow \infty} \int_M \int_{F_y} K \left(\frac{d_M^2(x, y)}{\epsilon}, \frac{d_{S_y}^2(P_{y,x}v, w)}{\gamma \epsilon} \right) p(y, w) d\text{vol}(w) d\text{vol}_M(y) \\
 &= \int_M K \left(\frac{d_M^2(x, y)}{\epsilon}, 0 \right) \left[\int_{F_y} p(y, w) d\text{vol}(w) \right] d\text{vol}_M(y).
 \end{aligned}$$

Define

$$\overline{K}_\epsilon(x, y) = \overline{K} \left(\frac{d_M^2(x, y)}{\epsilon} \right) = K \left(\frac{d_M^2(x, y)}{\epsilon}, 0 \right), \quad \overline{K}_\epsilon^{(\alpha)}(x, y) = \frac{\overline{K}_\epsilon(x, y)}{\langle p \rangle_\epsilon^\alpha(x) \langle p \rangle_\epsilon^\alpha(y)}.$$

By direct computation,

$$\begin{aligned}
 \lim_{\gamma \rightarrow \infty} H_{\epsilon, \gamma \epsilon}^{(\alpha)} f(x, v) &= \frac{\int_M \overline{K}_\epsilon^{(\alpha)}(x, y) \left[\int_{F_y} f(y, w) \frac{p(y, w)}{\langle p \rangle(y)} d\text{vol}(w) \right] \langle p \rangle(y) d\text{vol}_M(y)}{\int_M \overline{K}_\epsilon^{(\alpha)}(x, y) \langle p \rangle(y) d\text{vol}_M(y)} \\
 &= \frac{\int_M \overline{K}_\epsilon^{(\alpha)}(x, y) \langle f \rangle_p(y) \langle p \rangle(y) d\text{vol}_M(y)}{\int_M \overline{K}_\epsilon^{(\alpha)}(x, y) \langle p \rangle(y) d\text{vol}_M(y)}. \tag{C.20}
 \end{aligned}$$

By [7, Theorem 2], as $\epsilon \rightarrow 0$

$$\lim_{\gamma \rightarrow \infty} H_{\epsilon, \gamma \epsilon}^{(\alpha)} f(x, v) = \langle f \rangle_p(x) + \epsilon \frac{m'_2}{2m'_0} \frac{[\Delta_M(\langle f \rangle_p \langle p \rangle^{1-\alpha}) - \langle f \rangle_p \Delta_M \langle p \rangle^{1-\alpha}](x)}{\langle p \rangle^{1-\alpha}(x)} + O(\epsilon^2), \tag{C.21}$$

where

$$m'_0 = \int_{B_1^d(0)} \overline{K}(r^2) ds_1 \cdots ds_d, \quad m'_2 = \int_{B_1^d(0)} \overline{K}(r^2) s_1^2 ds_1 \cdots ds_d. \quad \square$$

Appendix D. Proofs of Theorem 9 and Theorem 10

In this appendix, we prove the two finite sampling theorems on unit tangent bundles in Section 5.1, following the paths paved by [27,106,72,9]. Recall from Section 5.2 that for any $f \in C^\infty(UTM)$

$$H_{\epsilon,\delta}^{(\alpha)} f(x, v) = \frac{\int_{UTM} K_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) f(y, w) p(y, w) d\Theta(y, w)}{\int_{UTM} K_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) p(y, w) d\Theta(y, w)},$$

where $d\Theta(x, v) = d\text{vol}_{S_x} d\text{vol}_M(v)$ is the *Liouville measure*. Since S_x is a unit ball in $T_x M$, we shall also write $d\sigma_x = d\text{vol}_{S_x}$ for convenience.

D.1. Sampling without noise

The following lemma builds the bridge between the geodesic distance on the manifold and the Euclidean distance in the ambient space.

Lemma 16. *Let $\iota : M \hookrightarrow \mathbb{R}^D$ be an isometric embedding of the smooth d -dimensional closed Riemannian manifold M into \mathbb{R}^D . For any $x, y \in M$ such that $d_M(x, y) < \text{Inj}(M)$, we have*

$$d_M^2(x, y) = \|\iota(x) - \iota(y)\|^2 + \frac{1}{12} d_M^4(x, y) \|\Pi(\theta, \theta)\|^2 + O(d_M^5(x, y)), \quad (\text{D.1})$$

where $\theta \in T_x M$, $\|\theta\|_x = 1$ comes from the geodesic polar coordinates of y in a geodesic normal neighborhood of x :

$$y = \exp_x r\theta, \quad r = d_M(x, y).$$

Proof. See [107, Proposition 6]. \square

For proving Theorem 9, it is convenient to introduce the “Euclidean distance version” of the diffusion operators introduced in Section 2.2. Note that in Definition 5.1 the hat “ $\hat{\cdot}$ ” is used for empirical quantities; for the remainder of this appendix, the tilde “ $\tilde{\cdot}$ ” will be used for quantities in the definition of $H_\epsilon^{(\alpha)}$ and $H_{\epsilon,\delta}^{(\alpha)}$ with Euclidean distance in place of geodesic distance. For instance,²

$$\begin{aligned} \tilde{K}_{\epsilon,\delta}(x, v; y, w) &= K\left(\frac{\|x - y\|^2}{\epsilon}, \frac{\|P_{y,x}v - w\|_y^2}{\delta}\right), \\ \tilde{p}_{\epsilon,\delta}(x, v) &= \int_{UTM} \tilde{K}_{\epsilon,\delta}(x, v; y, w) p(y, w) d\Theta(y, w), \\ \tilde{K}_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) &= \frac{\tilde{K}_{\epsilon,\delta}(x, v; y, w)}{\tilde{p}_{\epsilon,\delta}^\alpha(x, v) \tilde{p}_{\epsilon,\delta}^\alpha(y, w)}, \end{aligned}$$

and eventually

² Note that here $\hat{K}_{\epsilon,\delta}^{(\alpha)} = \tilde{K}_{\epsilon,\delta}^{(\alpha)}$, but this equality no longer holds in next subsection where $\hat{K}_{\epsilon,\delta}^{(\alpha)}$ is constructed from estimated parallel-transport.

$$\tilde{H}_{\epsilon,\delta}^{(\alpha)} f(x, v) = \frac{\int_{UTM} \tilde{K}_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) f(y, w) p(y, w) d\Theta(y, w)}{\int_{UTM} \tilde{K}_{\epsilon,\delta}^{(\alpha)}(x, v; y, w) p(y, w) d\Theta(y, w)}.$$

The next step is to establish an asymptotic expansion of type (4.3) for $\tilde{H}_{\epsilon,\delta}^{(\alpha)}$. We deduce the following Lemma 17, the “Euclidean distance version” of Lemma 13, from Lemma 16 and Lemma 13 itself.

Lemma 17. *Let $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ be a smooth function compactly supported in $[0, 1]$. Assume M is a d -dimensional closed Riemannian manifold isometrically embedded in \mathbb{R}^D , with injectivity radius $\text{Inj}(M) > 0$. For any $\epsilon > 0$, define kernel function*

$$\hat{\Phi}_\epsilon(x, y) = \Phi\left(\frac{\|x - y\|^2}{\epsilon}\right) \quad (\text{D.2})$$

on $M \times M$, where $\|\cdot\|$ is the Euclidean distance on \mathbb{R}^D . If the parameter ϵ is sufficiently small such that $0 \leq \epsilon \leq \sqrt{\text{Inj}(M)}$, then the integral operator associated with kernel $\hat{\Phi}_\epsilon$

$$(\hat{\Phi}_\epsilon g)(x) := \int_M \hat{\Phi}_\epsilon(x, y) g(y) d\text{vol}_M(y) \quad (\text{D.3})$$

has the following asymptotic expansion as $\epsilon \rightarrow 0$

$$(\hat{\Phi}_\epsilon g)(x) = \epsilon^{\frac{d}{2}} \left[m_0 g(x) + \epsilon \frac{m_2}{2} (\Delta_M g(x) + E(x) g(x)) + O(\epsilon^2) \right], \quad (\text{D.4})$$

with

$$E(x) = -\frac{1}{3} \text{Scal}^M(x) + \frac{d(d+2)}{12} A(x)$$

where m_0, m_2 are constants that depend on the moments of Φ and the dimension d of the Riemannian manifold M , Δ_M is the Laplace-Beltrami operator on M , $\text{Scal}^M(x)$ is the scalar curvature of M at x , and $A(x)$ is a scalar function on M that only depends on the intrinsic dimension d and the second fundamental form of the isometric embedding $\iota : M \hookrightarrow \mathbb{R}^D$.

Proof. Since we already established Lemma 13, it suffices to expand the difference

$$\int_M \left[\Phi\left(\frac{\|x - y\|^2}{\epsilon}\right) - \Phi\left(\frac{d_M^2(x, y)}{\epsilon}\right) \right] g(y) d\text{vol}_M(y). \quad (\text{D.5})$$

Put y in geodesic polar coordinates in a geodesic normal neighborhood of $x \in M$:

$$y = \exp_x r\theta, \quad \text{with } r = d_M(x, y), \theta \in T_x M, \|\theta\|_x = 1,$$

and denote the geodesic normal coordinates around x as (s_1, \dots, s_d) . By Lemma 16,

$$\|x - y\|^2 - d_M^2(x, y) = -\frac{1}{12} d_M^4(x, y) \|\Pi(\theta, \theta)\|^2 + O(d_M^5(x, y))$$

thus

$$\begin{aligned} & \Phi\left(\frac{\|x-y\|^2}{\epsilon}\right) - \Phi\left(\frac{d_M^2(x, y)}{\epsilon}\right) \\ &= \Phi'\left(\frac{d_M^2(x, y)}{\epsilon}\right) \cdot \left(-\frac{1}{12\epsilon} d_M^4(x, y) \|\Pi(\theta, \theta)\|^2\right) + O\left(\frac{d_M^8(x, y)}{\epsilon^2}\right). \end{aligned} \quad (\text{D.6})$$

Recall that Φ is supported on the unit interval, which implies that in (D.5) only those $y \in M$ satisfying $\|x-y\| \leq \sqrt{\epsilon}$ or $d_M(x, y) \leq \sqrt{\epsilon}$ are involved. According to Lemma 16, for sufficiently small $\epsilon > 0$, $\|x-y\| \leq \sqrt{\epsilon}$ implies $d_M(x, y) < 2\sqrt{\epsilon}$, thus the higher order error in (D.6) is indeed

$$O\left(\frac{d_M^8(x, y)}{\epsilon^2}\right) = O\left(\frac{(\sqrt{\epsilon})^8}{\epsilon^2}\right) = O(\epsilon^2).$$

Therefore,

$$\begin{aligned} & \int_M \left[\Phi\left(\frac{\|x-y\|^2}{\epsilon}\right) - \Phi\left(\frac{d_M^2(x, y)}{\epsilon}\right) \right] g(y) d\text{vol}_M(y) \\ &= -\frac{1}{12\epsilon} \int_M \Phi'\left(\frac{d_M^2(x, y)}{\epsilon}\right) d_M^4(x, y) \|\Pi(\theta, \theta)\|^2 g(y) d\text{vol}_M(y) + \epsilon^{\frac{d}{2}} \cdot O(\epsilon^2) \\ &= -\frac{1}{12\epsilon} \int_M \Phi'\left(\frac{r^2}{\epsilon}\right) r^4 \|\Pi(\theta, \theta)\|^2 g(y) d\text{vol}_M(y) + \epsilon^{\frac{d}{2}} \cdot O(\epsilon^2). \end{aligned} \quad (\text{D.7})$$

In geodesic normal coordinates (s_1, \dots, s_d) ,

$$\begin{aligned} & \int_M \Phi'\left(\frac{r^2}{\epsilon}\right) r^4 \|\Pi(\theta, \theta)\|^2 g(y) d\text{vol}_M(y) \\ &= \int_{B_{\sqrt{\epsilon}}(0)} \Phi'\left(\frac{r^2}{\epsilon}\right) r^4 \|\Pi(\theta, \theta)\|^2 \tilde{g}(s) \left[1 - \frac{1}{6} R_{k\ell}(x) s_k s_\ell + O(r^3)\right] ds_1 \cdots ds_d. \end{aligned} \quad (\text{D.8})$$

Using the Taylor expansion of $\tilde{g}(s)$ around $s = 0$ and the symmetry of the integral, (D.8) reduces to

$$\begin{aligned} & \int_{B_{\sqrt{\epsilon}}(0)} \Phi'\left(\frac{r^2}{\epsilon}\right) r^4 \|\Pi(\theta, \theta)\|^2 g(x) ds_1 \cdots ds_d \\ &+ \int_{B_{\sqrt{\epsilon}}(0)} \Phi'\left(\frac{r^2}{\epsilon}\right) r^4 \|\Pi(\theta, \theta)\|^2 O(r^2) ds_1 \cdots ds_d \\ &= \epsilon^{\frac{d}{2}} \cdot \epsilon^2 g(x) \int_{S_1(0)} \|\Pi(\theta, \theta)\|^2 d\theta \int_0^1 \Phi'(\tilde{r}^2) \tilde{r}^{3+d} d\tilde{r} + \epsilon^{\frac{d}{2}} \cdot O(\epsilon^3). \end{aligned}$$

Let m_2 be the constant as in Lemma 14, ω_{d-1} the volume of the standard unit sphere in \mathbb{R}^d . Note that

$$\omega_{d-1} \int_0^1 \Phi(\tilde{r}^2) \tilde{r}^{d+1} d\tilde{r} = \int_{B_1(0)} \Phi(\tilde{r}^2) \tilde{r}^2 d\tilde{s}^1 \cdots d\tilde{s}^d = m_2 d.$$

Let $A(x)$ be the average of the length of the second fundamental form over the standard unit sphere, i.e.,

$$A(x) = \frac{1}{\omega_{d-1}} \int_{S_1(0)} \|\Pi(\theta, \theta)\|^2 d\theta.$$

Integrating the term involving $\Phi'(\tilde{r}^2)$ by parts to get

$$\epsilon^{\frac{d}{2}} \cdot \epsilon^2 g(x) \int_{S_1(0)} \|\Pi(\theta, \theta)\|^2 d\theta \int_0^1 \Phi'(\tilde{r}^2) \tilde{r}^{3+d} d\tilde{r} = -\epsilon^{\frac{d}{2}} \cdot \epsilon^2 \frac{m_2}{2} d(d+2) g(x) A(x). \quad (\text{D.9})$$

Therefore,

$$\begin{aligned} & \int_M \left[\Phi\left(\frac{\|x-y\|^2}{\epsilon}\right) - \Phi\left(\frac{d_M^2(x,y)}{\epsilon}\right) \right] g(y) d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \left[\epsilon \frac{m_2}{24} d(d+2) A(x) g(x) + O(\epsilon^2) \right], \end{aligned}$$

and thus

$$(\hat{\Phi}_\epsilon g)(x) = \epsilon^{\frac{d}{2}} \left[m_0 g(x) + \epsilon \frac{m_2}{2} (\Delta_M g(x) + E(x) g(x)) + O(\epsilon^2) \right]$$

where

$$E(x) := -\frac{1}{3} \text{Scal}(x) + \frac{1}{12} d(d+2) A(x). \quad \square$$

Remark 7.5. The only difference between the conclusions in Lemma 17 and Lemma 13 is that the scalar function $E(x)$ takes the place of the scalar curvature $\text{Scal}^M(x)$; one can check, essentially by going through the proof of Theorem 4, that the proof still works through, due to the cancellation of the terms involving $E(x)$. In fact, applying Lemma 17 repeatedly, one has

$$\begin{aligned} & \int_M \hat{\Phi}_\epsilon(x, y) f(y, P_{y,x} v) d\text{vol}_M(y) \\ &= \epsilon^{\frac{d}{2}} \left\{ m_0 f(x, v) + \epsilon \frac{m_2}{2} [\Delta_H f(x, v) + E_1(x) f(x, v)] + O(\epsilon^2) \right\}, \end{aligned} \quad (\text{D.10})$$

and

$$\begin{aligned} & \int_{UTM} \tilde{K}_{\epsilon, \delta}(x, v; y, w) g(y, w) d\Theta(y, w) \\ &= \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \left\{ m_0 g(x, v) + \epsilon \frac{m_{21}}{2} [\Delta_{UTM}^H g(x, v) + E_1(x) g(x, v)] \right. \\ & \quad \left. + \delta \frac{m_{22}}{2} [\Delta_{UTM}^V g(x, v) + E_2 \cdot g(x, v)] + O(\epsilon^2 + \delta^2) \right\}, \end{aligned} \quad (\text{D.11})$$

where

$$E_1(\xi_i) = -\frac{1}{3}\text{Scal}^M(\xi_i) + \frac{d(d+2)}{12} \cdot \frac{1}{\omega_{d-1}} \int_{S_1(0)} \|\Pi_M(\theta, \theta)\|^2 d\theta$$

only depends on the scalar curvature Scal^M and the second fundamental form Π_M of the base manifold M at ξ , and

$$E_2 = -\frac{1}{3}\text{Scal}^S + \frac{(d-1)(d+1)}{12} \cdot \frac{1}{\omega_{d-2}} \int_{S_1(0)} \|\Pi_S(\theta, \theta)\|^2 d\theta$$

is a constant because

$$\text{Scal}^S \equiv (d-1)(d-2), \quad \|\Pi_S(\theta, \theta)\|^2 \equiv 1 \quad \text{for any unit tangent vector } \theta.$$

These expansions are essentially the equivalents of Lemma 14 and Lemma 15 for $\tilde{K}_{\epsilon, \delta}$. Using (D.10), (D.11) and picking $\delta = O(\epsilon)$ as $\epsilon \rightarrow 0$, a version of Theorem 4 holds true when $K_{\epsilon, \delta}^{(\alpha)}$ is replaced with $\tilde{K}_{\epsilon, \delta}^{(\alpha)}$, i.e., as $\epsilon \rightarrow 0$ (and thus $\delta \rightarrow 0$),

$$\begin{aligned} \tilde{H}_{\epsilon, \delta}^{(\alpha)} f(x, v) &= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_H [fp^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} - f(x, v) \frac{\Delta_H p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} \right] \\ &\quad + \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_E [fp^{1-\alpha}](x, v)}{p^{1-\alpha}(x, v)} - f(x, v) \frac{\Delta_E p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} \right] + O(\epsilon^2 + \epsilon\delta + \delta^2). \end{aligned} \quad (\text{D.12})$$

As we shall see below, this observation is the key to establishing estimates for the bias error in the proof of Theorem 9.

Before we present the proof of Theorem 9, we establish a large deviation bound for our two-step sampling strategy. Recall from Assumption 8 that we first sample N_B points ξ_1, \dots, ξ_{N_B} i.i.d. with respect to $\langle p \rangle$ on the base manifold M , then sample N_F points on each fibre S_{ξ_j} i.i.d. with respect to $p(\cdot | \xi_j)$. The resulting $N_B \times N_F$ points on UTM

$$\begin{matrix} x_{1,1}, & x_{1,2}, & \cdots, & x_{1,N_F} \\ x_{2,1}, & x_{2,2}, & \cdots, & x_{2,N_F} \\ \vdots & \vdots & \cdots & \vdots \\ x_{N_B,1}, & x_{N_B,2}, & \cdots, & x_{N_B,N_F} \end{matrix}$$

are generally not i.i.d. sampled from UTM . This forbids applying the *Law of Large Numbers* directly to quantities that take the form of an average over the entire unit tangent bundle, such as

$$\frac{1}{N_B N_F} \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) f(x_{j,s}).$$

However, due to the conditional i.i.d. fibrewise sampling, it makes sense to apply the law of large numbers to average quantities on a fixed fibre, e.g.,

$$\frac{1}{N_F} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) f(x_{j,s}) \longrightarrow \mathbb{E}_Z [\tilde{K}_{\epsilon, \delta}(x_{i,r}, (\xi_j, Z)) f(\xi_j, Z)],$$

where \mathbb{E}_Z stands for the expectation with respect to the “fibre component” of the coordinates of the points on S_{ξ_j} . Explicitly,

$$\mathbb{E}_Z [\tilde{K}_{\epsilon,\delta}(x_{i,r}, (\xi_j, \cdot)) f(\xi_j, \cdot)] = \int_{S_{\xi_j}} \tilde{K}_{\epsilon,\delta}(x_{i,r}, (\xi_j, w)) f(\xi_j, w) p(w \mid \xi_j) d\sigma_{\xi_j}(w).$$

Next, note that ξ_1, \dots, ξ_{N_B} are i.i.d. sampled from the base manifold M , the partial expectations

$$\{\mathbb{E}_Z [\tilde{K}_{\epsilon,\delta}(x_{i,r}, (\xi_j, Z)) f(\xi_j, Z)]\}_{j=1}^{N_B}$$

are i.i.d. random variables on M with respect to $\langle p \rangle$. Thus

$$\frac{1}{N_B} \sum_{j=1}^{N_B} \mathbb{E}_Z [\tilde{K}_{\epsilon,\delta}(x_{i,r}, (\xi_j, Z)) f(\xi_j, Z)] \longrightarrow \mathbb{E}_Y [\mathbb{E}_Z [\tilde{K}_{\epsilon,\delta}(x_{i,r}, (Y, Z)) f(Y, Z)]] ,$$

where

$$\begin{aligned} & \mathbb{E}_Y [\mathbb{E}_Z [\tilde{K}_{\epsilon,\delta}(x_{i,r}, (Y, Z)) f(Y, Z)]] \\ &= \int_M \langle p \rangle(y) \int_{S_{\xi_j}} \tilde{K}_{\epsilon,\delta}(x_{i,r}, (\xi_j, w)) f(y, w) p(w \mid y) d\sigma_y(w) d\text{vol}_M(y) \\ &= \int_M \int_{S_{\xi_j}} \tilde{K}_{\epsilon,\delta}(x_{i,r}, (\xi_j, w)) f(y, w) p(y, w) d\sigma_y(w) d\text{vol}_M(y). \end{aligned}$$

This gives

$$\begin{aligned} & \lim_{N_B \rightarrow \infty} \lim_{N_F \rightarrow \infty} \frac{1}{N_B N_F} \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) f(x_{j,s}) \\ &= \int_M \int_{S_{\xi_j}} \tilde{K}_{\epsilon,\delta}(x_{i,r}, (\xi_j, w)) f(y, w) p(y, w) d\sigma_y(w) d\text{vol}_M(y), \end{aligned}$$

in which the two limits on the left hand side do not commute in general. Thus it is natural to consider iterated partial expectations rather than expectation on the entire UTM . From now on, we denote $\mathbb{E}_Y, \mathbb{E}_Z$ as $\mathbb{E}_1, \mathbb{E}_2$ for simplicity.

Definition 7.2. Let p be a probability density function on UTM , and

$$\langle p \rangle(x) = \int_{S_x} p(x, w) d\sigma_x(w), \quad p(v \mid x) = \frac{p(x, v)}{\langle p \rangle(x)}$$

be as defined in (4.6), (4.7). For any $f \in C^\infty(M)$ and $g \in C^\infty(S_\xi)$, $\xi \in M$, define

$$\begin{aligned} \mathbb{E}_1 f &:= \int_M f(y) p(y) d\text{vol}_M(y), \\ \mathbb{E}_2^\xi g &:= \int_{S_\xi} g(\xi, w) p(w \mid \xi) d\sigma_\xi(w). \end{aligned}$$

Definition 7.3. Let p be a probability density function on UTM . We call a collection of $N_B \times N_F$ real-valued random functions

$$\{X_{j,s} \mid 1 \leq j \leq N_B, 1 \leq s \leq N_F\}$$

Procrustean with respect to p on UTM, if

- (i) For each $1 \leq j \leq N_B$, the subcollection $\{X_{j,s} \mid 1 \leq s \leq N_F\}$ are i.i.d. on S_{ξ_j} for some $\xi_j \in M$, with respect to the conditional probability density $p(\cdot \mid \xi_j)$;
- (ii) The points $\{\xi_j \mid 1 \leq j \leq N_B\}$ are i.i.d. on M with respect to the fibre average density $\langle p \rangle(\cdot)$.

Due to (7.3), we can drop the dependency of $X_{j,s}$ with respect to s and simply write

$$\mathbb{E}_2^{\xi_j} X_j := \mathbb{E}_2^{\xi_j} X_{j,s}, \quad \mathbb{E}_2^{\xi_j} X_j^2 := \mathbb{E}_2^{\xi_j} X_{j,s}^2.$$

Similarly, because of (7.3) we can write

$$\mathbb{E}_1 \mathbb{E}_2 X := \mathbb{E}_1 \mathbb{E}_2^{\xi_j} X_j, \quad \mathbb{E}_1 (\mathbb{E}_2 X)^2 := \mathbb{E}_1 \left(\mathbb{E}_2^{\xi_j} X_j \right).$$

Lemma 18. Let $\{X_{j,s} \mid 1 \leq j \leq N_B, 1 \leq s \leq N_F\}$ be a collection of Procrustean random functions with respect to some density function p on UTM. If

$$|X_{j,s}| \leq M_0, \quad \left| \mathbb{E}_2^{\xi_j} X_j \right| \leq M_1, \quad |\mathbb{E}_1 \mathbb{E}_2 X| \leq M_2 \quad \text{a.s. for all } 1 \leq j \leq N_B, 1 \leq s \leq N_F,$$

then for any $t > 0$ and $0 < \theta < 1$,

$$\begin{aligned} & \mathbb{P} \left\{ \frac{1}{N_B N_F} \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} X_{j,s} - \mathbb{E}_1 \mathbb{E}_2 X > t \right\} \\ & \leq \sum_{j=1}^{N_B} \exp \left\{ - \frac{\frac{1}{2} (1-\theta)^2 N_F t^2}{\left[\mathbb{E}_2^{\xi_j} X_j^2 - \left(\mathbb{E}_2^{\xi_j} X_j \right)^2 \right] + \frac{1}{3} (M_0 + M_1) (1-\theta) t} \right\} \\ & \quad + \exp \left\{ - \frac{\frac{1}{2} \theta^2 N_B t^2}{\left[\mathbb{E}_1 (\mathbb{E}_2 X)^2 - (\mathbb{E}_1 \mathbb{E}_2 X)^2 \right] + \frac{1}{3} (M_1 + M_2) \theta t} \right\}. \end{aligned}$$

Proof. Note that for any $\theta \in (0, 1)$

$$\begin{aligned} & \mathbb{P} \left\{ \frac{1}{N_B N_F} \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} X_{j,s} - \mathbb{E}_1 \mathbb{E}_2 X > t \right\} \\ & \leq \mathbb{P} \left\{ \frac{1}{N_B N_F} \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} X_{j,s} - \frac{1}{N_B} \sum_{j=1}^{N_B} \mathbb{E}_2^{\xi_j} X_j > (1-\theta) t \right\} \\ & \quad + \mathbb{P} \left\{ \frac{1}{N_B} \sum_{j=1}^{N_B} \mathbb{E}_2^{\xi_j} X_j - \mathbb{E}_1 \mathbb{E}_2 X > \theta t \right\} =: (\text{I}) + (\text{II}). \end{aligned}$$

Since

$$\left| \mathbb{E}_2^{\xi_j} X_j - \mathbb{E}_1 \mathbb{E}_2 X \right| \leq M_1 + M_2,$$

by Bernstein's Inequality [108, §2.2],

$$\begin{aligned}
 \text{(II)} &= \mathbb{P} \left\{ \sum_{j=1}^{N_B} \left(\mathbb{E}_2^{\xi_j} X_j - \mathbb{E}_1 \mathbb{E}_2 X \right) > \theta N_B t \right\} \\
 &\leq \exp \left\{ - \frac{\frac{1}{2} \theta^2 N_B t}{\left[\mathbb{E}_1 (\mathbb{E}_2 X)^2 - (\mathbb{E}_1 \mathbb{E}_2 X)^2 \right] + \frac{1}{3} (M_1 + M_2) \theta t} \right\}.
 \end{aligned}$$

For (I), a union bound plus Bernstein's Inequality gives

$$\begin{aligned}
 \text{(I)} &= \mathbb{P} \left\{ \sum_{j=1}^{N_B} \left(\frac{1}{N_F} \sum_{s=1}^{N_F} X_{j,s} - \mathbb{E}_2^{\xi_j} X_j \right) > (1-\theta) N_B t \right\} \\
 &\leq \sum_{j=1}^{N_B} \mathbb{P} \left\{ \sum_{s=1}^{N_F} \left(X_{j,s} - \mathbb{E}_2^{\xi_j} X_j \right) > (1-\theta) N_F t \right\} \\
 &\leq \sum_{j=1}^{N_B} \exp \left\{ - \frac{\frac{1}{2} (1-\theta)^2 N_F t^2}{\left[\mathbb{E}_2^{\xi_j} X_j^2 - (\mathbb{E}_2^{\xi_j} X_j)^2 \right] + \frac{1}{3} (M_0 + M_1) (1-\theta) t} \right\}.
 \end{aligned}$$

The conclusion follows from combining these two bounds. \square

Remark 7.6. Intuitively, the second term in the bound comes from the sampling error on the base manifold, and is thus independent of δ and N_F ; the first term in the bound comes from accumulating fibrewise sampling error across all N_B fibres.

Proof of Theorem 9. We first establish the result for $\alpha = 0$. In this case, $\hat{K}_{\epsilon,\delta}^{(0)}(\cdot, \cdot) = \hat{K}_{\epsilon,\delta}(\cdot, \cdot)$, and

$$\begin{aligned}
 \hat{H}_{\epsilon,\delta}^{(0)} f(x_{i,r}) &= \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) f(x_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s})} \\
 &= \frac{\frac{1}{N_B N_F} \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} K \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} x_{i,r} - x_{j,s}\|^2}{\delta} \right) f(x_{j,s})}{\frac{1}{N_B N_F} \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} K \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} x_{i,r} - x_{j,s}\|^2}{\delta} \right)}.
 \end{aligned}$$

Since $\{x_{j,s}\}_{s=1}^{N_F}$ are i.i.d. with respect to $p(\cdot | \xi_j)$, by the law of large numbers, for each fixed $j = 1, \dots, N_B$, as $N_F \rightarrow \infty$,

$$\begin{aligned}
 &\lim_{N_F \rightarrow \infty} \frac{1}{N_F} \sum_{s=1}^{N_F} K \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} x_{i,r} - x_{j,s}\|^2}{\delta} \right) f(x_{j,s}) \\
 &= \int_{S_{\xi_j}} K \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} x_{i,r} - w\|^2}{\delta} \right) f(\xi_j, w) p(w | \xi_j) d\sigma_{\xi_j}(w).
 \end{aligned}$$

Note that $\{\xi_j\}_{j=1}^{N_B}$ are i.i.d. with respect to $\langle p \rangle$, it follows again from the law of large numbers that

$$\begin{aligned} & \lim_{N_B \rightarrow \infty} \frac{1}{N_B} \sum_{j=1}^{N_B} \lim_{N_F \rightarrow \infty} \frac{1}{N_F} \sum_{s=1}^{N_F} K \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} x_{i,r} - x_{j,s}\|^2}{\delta} \right) f(x_{j,s}) \\ &= \int_{UTM} K \left(\frac{\|\xi_i - y\|^2}{\epsilon}, \frac{\|P_{y, \xi_i} x_{i,r} - w\|^2}{\delta} \right) f(y, w) p(y, w) d\Theta(y, w), \end{aligned}$$

where we used $p(y, w) = \langle p \rangle(y) p(w | y)$. For $f \equiv 1$,

$$\begin{aligned} & \lim_{N_B \rightarrow \infty} \frac{1}{N_B} \sum_{j=1}^{N_B} \lim_{N_F \rightarrow \infty} \frac{1}{N_F} \sum_{s=1}^{N_F} K \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} x_{i,r} - x_{j,s}\|^2}{\delta} \right) \\ &= \int_{UTM} K \left(\frac{\|\xi_i - y\|^2}{\epsilon}, \frac{\|P_{y, \xi_i} x_{i,r} - w\|^2}{\delta} \right) p(y, w) d\Theta(y, w). \end{aligned}$$

Therefore,

$$\begin{aligned} & \lim_{N_B \rightarrow \infty} \lim_{N_F \rightarrow \infty} \hat{H}_{\epsilon, \delta}^{(0)} f(x_{i,r}) = \tilde{H}_{\epsilon, \delta}^{(0)} f(x_{i,r}) \\ &= f(x_{i,r}) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_{UTM}^H [fp](x_{i,r})}{p(x_{i,r})} - f(x_{i,r}) \frac{\Delta_{UTM}^H p(x_{i,r})}{p(x_{i,r})} \right] \\ &+ \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_{UTM}^V [fp](x_{i,r})}{p(x_{i,r})} - f(x_{i,r}) \frac{\Delta_{UTM}^V p(x_{i,r})}{p(x_{i,r})} \right] + O(\epsilon^2 + \epsilon\delta + \delta^2). \end{aligned}$$

The last equality makes use of the assumption $\delta = O(\epsilon)$ as $\epsilon \rightarrow 0$ and Remark 7.5. This establishes the bias error for the special case $\alpha = 0$ and it remains to estimate the variance error. To this end, denote

$$F_{j,s} = \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) f(x_{j,s}), \quad G_{j,s} = \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s})$$

for any fixed $x_{i,r} \in UTM$. Note that $F_{i,s} = 0, G_{i,s} = 0$ for all $s = 1, \dots, N_F$, by Definition 5.1 (5.1); by the compactness of UTM we have the following trivial bounds uniform in j, s :

$$|F_{j,s}| \leq \|K\|_\infty \|f\|_\infty, \quad |G_{j,s}| \leq \|K\|_\infty.$$

Thus we already have

$$\lim_{N_B \rightarrow \infty} \lim_{N_F \rightarrow \infty} \hat{H}_{\epsilon, \delta}^{(0)} f(x_{i,r}) = \frac{\mathbb{E}_1 \mathbb{E}_2 F}{\mathbb{E}_1 \mathbb{E}_2 G},$$

and would like to estimate

$$p(N_B, N_F, \beta) := \mathbb{P} \left\{ \frac{\sum_j \sum_s F_{j,s}}{\sum_j \sum_s G_{j,s}} - \frac{\mathbb{E}_1 \mathbb{E}_2 F}{\mathbb{E}_1 \mathbb{E}_2 G} > \beta \right\}$$

for sufficiently small $\beta > 0$. An upper bound for

$$\mathbb{P} \left\{ \frac{\sum_j \sum_s F_{j,s}}{\sum_j \sum_s G_{j,s}} - \frac{\mathbb{E}_1 \mathbb{E}_2 F}{\mathbb{E}_1 \mathbb{E}_2 G} < -\beta \right\}$$

can be obtained in a similar manner. Since $G_{j,s} > 0$,

$$\begin{aligned} p(N_B, N_F, \beta) &= \mathbb{P} \left\{ \frac{\left(\sum_j \sum_s F_{j,s} \right) \mathbb{E}_1 \mathbb{E}_2 G - \left(\sum_j \sum_s G_{j,s} \right) \mathbb{E}_1 \mathbb{E}_2 F}{\left(\sum_j \sum_s G_{j,s} \right) \mathbb{E}_1 \mathbb{E}_2 G} > \beta \right\} \\ &= \mathbb{P} \left\{ \left(\sum_j \sum_s F_{j,s} \right) \mathbb{E}_1 \mathbb{E}_2 G - \left(\sum_j \sum_s G_{j,s} \right) \mathbb{E}_1 \mathbb{E}_2 F > \beta \left(\sum_j \sum_s G_{j,s} \right) \mathbb{E}_1 \mathbb{E}_2 G \right\}. \end{aligned}$$

Denote

$$Y_{j,s} := F_{j,s} \mathbb{E}_1 \mathbb{E}_2 G - G_{j,s} \mathbb{E}_1 \mathbb{E}_2 F + \beta (\mathbb{E}_1 \mathbb{E}_2 G - G_{j,s}) \mathbb{E}_1 \mathbb{E}_2 G,$$

then it is easily verifiable that $\mathbb{E}_1 \mathbb{E}_2 Y_{j,s} = 0$ for all $1 \leq j \leq N_B$, $1 \leq s \leq N_F$, and

$$p(N_B, N_F, \beta) = \mathbb{P} \left\{ \frac{1}{N_B N_F} \sum_j \sum_s Y_{j,s} > \beta (\mathbb{E}_1 \mathbb{E}_2 G)^2 \right\}.$$

By Lemma 18, bounding this quantity reduces to computing various moments. Define

$$X_j := \mathbb{E}_2 Y_j,$$

then X_1, \dots, X_{N_B} are i.i.d. on M with respect to $\langle p \rangle$, and $\mathbb{E}_1 X_j = 0$ for $1 \leq j \leq N_B$. Furthermore, X_1, \dots, X_{N_B} are uniformly bounded. To find this bound explicitly, note that

$$\begin{aligned} |X_j| &= |\mathbb{E}_2 Y_j| = |(\mathbb{E}_2 F_j) \mathbb{E}_1 \mathbb{E}_2 G - (\mathbb{E}_2 G_j) \mathbb{E}_1 \mathbb{E}_2 F + \beta (\mathbb{E}_1 \mathbb{E}_2 G - \mathbb{E}_2 G_j) \mathbb{E}_1 \mathbb{E}_2 G| \\ &\leq |(\mathbb{E}_2 F_j) \mathbb{E}_1 \mathbb{E}_2 G| + |(\mathbb{E}_2 G_j) \mathbb{E}_1 \mathbb{E}_2 F| + \beta (\mathbb{E}_1 \mathbb{E}_2 G)^2 + \beta |\mathbb{E}_2 G_j| |\mathbb{E}_1 \mathbb{E}_2 G|, \end{aligned}$$

and recall from Lemma 17 and Remark 7.5 that

$$\begin{aligned} \mathbb{E}_1 \mathbb{E}_2 F &= O\left(\epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}\right), \quad \mathbb{E}_1 \mathbb{E}_2 G = O\left(\epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}\right), \\ \mathbb{E}_2 F_j &= O\left(\delta^{\frac{d-1}{2}}\right), \quad \mathbb{E}_2 G_j = O\left(\delta^{\frac{d-1}{2}}\right), \end{aligned}$$

thus

$$|X_j| \leq \tilde{C} \epsilon^{\frac{d}{2}} \delta^{d-1} + \beta \left(\epsilon^d \delta^{d-1} + \epsilon^{\frac{d}{2}} \delta^{d-1} \right)$$

where \tilde{C} is some positive constant depending on the pointwise bounds of K , p , and f . Since we will be mostly interested in small $\beta > 0$, let us pick $\beta = O(\epsilon^2 + \epsilon \delta + \delta^2)$ and rewrite the upper bound as

$$|X_j| \leq C \epsilon^{\frac{d}{2}} \delta^{d-1}, \quad C = C(\|K\|_\infty, \|f\|_\infty, p_m, p_M) > 0. \quad (\text{D.13})$$

We then need to bound $\mathbb{E}_1 X_j^2$. Since

$$\begin{aligned} \mathbb{E}_1 X_j^2 &= \left[\mathbb{E}_1 (\mathbb{E}_2 F_j)^2 \right] (\mathbb{E}_1 \mathbb{E}_2 G)^2 + \left[\mathbb{E}_1 (\mathbb{E}_2 G_j)^2 \right] (\mathbb{E}_1 \mathbb{E}_2 F)^2 \\ &\quad - 2 \mathbb{E}_1 [(\mathbb{E}_2 F_j) (\mathbb{E}_2 G_j)] (\mathbb{E}_1 \mathbb{E}_2 F) (\mathbb{E}_1 \mathbb{E}_2 G) + \beta^2 (\mathbb{E}_1 \mathbb{E}_2 G)^2 \left[\mathbb{E}_1 (\mathbb{E}_2 G)^2 - (\mathbb{E}_1 \mathbb{E}_2 G)^2 \right] \\ &\quad + 2\beta (\mathbb{E}_1 \mathbb{E}_2 G) \left[\mathbb{E}_1 (\mathbb{E}_2 G)^2 \mathbb{E}_1 \mathbb{E}_2 F - (\mathbb{E}_1 \mathbb{E}_2 G) \mathbb{E}_1 (\mathbb{E}_2 F_j \mathbb{E}_2 G_j) \right], \end{aligned}$$

it suffices to compute the first and second moments of $\mathbb{E}_2 F_j$, $\mathbb{E}_2 G_j$ for $1 \leq j \leq N_B$. By (D.11),

$$\begin{aligned}\mathbb{E}_1 \mathbb{E}_2 F &= \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \left\{ m_0 [fp](x_{i,r}) + \epsilon \frac{m_{21}}{2} (\Delta_{UTM}^H [fp](x_{i,r}) + E_1(\xi_i) [fp](x_{i,r})) \right. \\ &\quad \left. + \delta \frac{m_{22}}{2} (\Delta_{UTM}^V [fp](x_{i,r}) + E_2 \cdot [fp](x_{i,r})) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}, \\ \mathbb{E}_1 \mathbb{E}_2 G &= \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \left\{ m_0 p(x_{i,r}) + \epsilon \frac{m_{21}}{2} (\Delta_{UTM}^H p(x_{i,r}) + E_1(\xi_i) p(x_{i,r})) \right. \\ &\quad \left. + \delta \frac{m_{22}}{2} (\Delta_{UTM}^V p(x_{i,r}) + E_2 p(x_{i,r})) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}.\end{aligned}$$

Using the notation and applying Lemma 17 once

$$p(v | x) = \frac{p(x, v)}{\langle p \rangle(x)} = \frac{p(x, v)}{\langle p \rangle \circ \pi(x, v)} = \left[\frac{p}{\langle p \rangle \circ \pi} \right](x, v),$$

we have

$$\begin{aligned}\mathbb{E}_2 F_j &= \delta^{\frac{d-1}{2}} \left\{ M_0 \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon} \right) \left[\frac{fp}{\langle p \rangle \circ \pi} \right] (P_{\xi_j, \xi_i} x_{i,r}) \right. \\ &\quad \left. + \frac{\delta}{2} M_2 \left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon} \right) \left(\Delta_{UTM}^V \left[\frac{fp}{\langle p \rangle \circ \pi} \right] + E_2 \cdot \left[\frac{fp}{\langle p \rangle \circ \pi} \right] \right) (P_{\xi_j, \xi_i} x_{i,r}) + O(\delta^2) \right\},\end{aligned}$$

where $M_0(\cdot)$, $M_2(\cdot)$ are functions depending only on the kernel K , as in the proof of Lemma 15. By a direct computation using Lemma 17,

$$\begin{aligned}\mathbb{E}_1 (\mathbb{E}_2 F_j)^2 &= \epsilon^{\frac{d}{2}} \delta^{d-1} \left\{ m'_0 \left[\frac{(fp)^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + \epsilon \frac{m'_{21}}{2} \left(\Delta_{UTM}^H \left[\frac{(fp)^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + E_1(\xi_i) \left[\frac{(fp)^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \right) \right. \\ &\quad \left. + \delta m'_{22} [fp](x_{i,r}) \left(\Delta_{UTM}^V \left[\frac{fp}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + E_2 \cdot \left[\frac{fp}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \right) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\},\end{aligned}$$

where m'_0 , m'_{21} , m'_{22} are positive constants determined by the kernel function K and dimension d :

$$\begin{aligned}m'_0 &= \int_{B_1^d(0)} M_0^2(r^2) ds^1 \cdots ds^d, \quad r^2 = \sum_{k=1}^d (s^k)^2, \\ m'_{21} &= \int_{B_1^d(0)} M_0(r^2) (s^1)^2 ds^1 \cdots ds^d, \quad m'_{22} = \int_{B_1^d(0)} M_0(r^2) M_2(r^2) ds^1 \cdots ds^d.\end{aligned}$$

Setting $f \equiv 1$,

$$\begin{aligned}\mathbb{E}_1 (\mathbb{E}_2 G_j)^2 &= \epsilon^{\frac{d}{2}} \delta^{d-1} \left\{ m'_0 \left[\frac{p^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + \epsilon \frac{m'_{21}}{2} \left(\Delta_{UTM}^H \left[\frac{p^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + E_1(\xi_i) \left[\frac{p^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \right) \right. \\ &\quad \left. + \delta m'_{22} p(x_{i,r}) \left(\Delta_{UTM}^V \left[\frac{p}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + E_2 \cdot \left[\frac{p}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \right) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}.\end{aligned}$$

Similarly,

$$\begin{aligned} \mathbb{E}_1 [(\mathbb{E}_2 F_j) (\mathbb{E}_2 G_j)] &= \epsilon^{\frac{d}{2}} \delta^{d-1} \left\{ m'_0 \left[\frac{fp^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \right. \\ &+ \epsilon \frac{m'_{21}}{2} \left(\Delta_{UTM}^H \left[\frac{fp^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + E_1 (\xi_i) \left[\frac{fp^2}{\langle p \rangle \circ \pi} \right] \right) \\ &+ \delta \frac{m'_{22}}{2} \left(p(x_{i,r}) \Delta_{UTM}^V \left[\frac{fp}{\langle p \rangle \circ \pi} \right] (x_{i,r}) + [fp](x_{i,r}) \Delta_{UTM}^V \left[\frac{p}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \right. \\ &\left. \left. + 2E_2 \cdot \left[\frac{fp^2}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \right) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\}. \end{aligned}$$

Take $\beta = O(\epsilon^2 + \epsilon\delta + \delta^2)$ so that $O(\beta)$ and $O(\beta^2)$ terms are absorbed into

$$O \left[\epsilon^{\frac{3d}{2}} \delta^{2(d-1)} (\epsilon^2 + \epsilon\delta + \delta^2) \right].$$

Direct computation using

$$\Delta_{UTM}^H (f^2 g) + f^2 \Delta_{UTM}^H g - 2f \Delta_{UTM}^H (fg) = 2 \|\nabla_{UTM}^H f\|^2 g$$

gives

$$\begin{aligned} \mathbb{E}_1 X_j^2 &= \left[\mathbb{E}_1 (\mathbb{E}_2 F_j)^2 \right] (\mathbb{E}_1 \mathbb{E}_2 G)^2 + \left[\mathbb{E}_1 (\mathbb{E}_2 G_j)^2 \right] (\mathbb{E}_1 \mathbb{E}_2 F)^2 \\ &- 2\mathbb{E}_1 [(\mathbb{E}_2 F_j) (\mathbb{E}_2 G_j)] (\mathbb{E}_1 \mathbb{E}_2 F) (\mathbb{E}_1 \mathbb{E}_2 G) + O(\epsilon^2 + \delta^2) \\ &= \epsilon^{\frac{3d}{2}} \delta^{2(d-1)} \left\{ \epsilon m_0^2 m'_{21} \left[\frac{p^4}{\langle p \rangle \circ \pi} \right] (x_{i,r}) \|\nabla_{UTM}^H f\|^2 (x_{i,r}) + O(\epsilon^2 + \epsilon\delta + \delta^2) \right\} \\ &\leq \epsilon^{\frac{3d}{2}} \delta^{2(d-1)} (C' \epsilon + O(\epsilon^2 + \epsilon\delta + \delta^2)) \end{aligned}$$

where

$$C' = \frac{m_0^2 m'_{21} p_M^4 \|\nabla_{UTM}^H f\|_\infty^2}{\omega_{d-1}^4 p_m} > 0.$$

Note that $O(\delta)$ terms do not show up in this bound, intuitively because $X_j = \mathbb{E}_2 Y_{j,s}$ is already the expectation along the fibre direction, which “freezes” the variability controlled by the fibrewise bandwidth δ .

It remains to bound

$$\mathbb{E}_2^{\xi_j} Y_j^2 - \left(\mathbb{E}_2^{\xi_j} Y_j \right)^2$$

for each $1 \leq j \leq N_B$. Since we picked $\beta = O(\epsilon^2 + \epsilon\delta + \delta^2)$,

$$|Y_{j,s}| = |F_{j,s} \mathbb{E}_1 \mathbb{E}_2 G - G_{j,s} \mathbb{E}_1 \mathbb{E}_2 F + \beta (\mathbb{E}_1 \mathbb{E}_2 G - G_{j,s}) \mathbb{E}_1 \mathbb{E}_2 G| \leq C \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}$$

where

$$C = C(\|K\|_\infty, \|f\|_\infty, p_m, p_M)$$

is a positive constant. Again taking advantage of $\beta = O(\epsilon^2 + \epsilon\delta + \delta^2)$, we have

$$\begin{aligned} \mathbb{E}_2 Y_j^2 - (\mathbb{E}_2 Y_j)^2 &= \left[\mathbb{E}_2 F_{j,s}^2 - (\mathbb{E}_2 F_j)^2 \right] (\mathbb{E}_1 \mathbb{E}_2 G)^2 + \left[\mathbb{E}_2 G_{j,s}^2 - (\mathbb{E}_2 G_j)^2 \right] (\mathbb{E}_1 \mathbb{E}_2 F)^2 \\ &\quad + 2 [(\mathbb{E}_2 F_j) (\mathbb{E}_2 G_j) - \mathbb{E}_2 (F_{j,s} G_{j,s})] (\mathbb{E}_1 \mathbb{E}_2 F) (\mathbb{E}_1 \mathbb{E}_2 G) \\ &\quad + O \left[\epsilon^d \delta^{2(d-1)} (\epsilon^2 + \epsilon\delta + \delta^2) \right]. \end{aligned}$$

Note that

$$\mathbb{E}_2 F_{j,s}^2 = O \left(\delta^{\frac{d-1}{2}} \right), \quad \mathbb{E}_2 G_{j,s}^2 = O \left(\delta^{\frac{d-1}{2}} \right), \quad \mathbb{E}_2 [F_{j,s} G_{j,s}] = O \left(\delta^{\frac{d-1}{2}} \right),$$

but

$$(\mathbb{E}_2 F_{j,s})^2 = O(\delta^{d-1}), \quad (\mathbb{E}_2 G_{j,s})^2 = O(\delta^{d-1}), \quad (\mathbb{E}_2 F_j) (\mathbb{E}_2 G_j) = O(\delta^{d-1}),$$

the leading order error term in $\mathbb{E}_2 Y_j^2 - (\mathbb{E}_2 Y_j)^2$ is

$$(\mathbb{E}_2 F_{j,s}^2) (\mathbb{E}_1 \mathbb{E}_2 G)^2 + (\mathbb{E}_2 G_{j,s}^2) (\mathbb{E}_1 \mathbb{E}_2 F)^2 - 2\mathbb{E}_2 (F_{j,s} G_{j,s}) (\mathbb{E}_1 \mathbb{E}_2 F) (\mathbb{E}_1 \mathbb{E}_2 G).$$

By Lemma 17,

$$\begin{aligned} \mathbb{E}_2 F_{j,s}^2 &= \delta^{\frac{d-1}{2}} \left\{ \widetilde{M}_0 \left[\frac{f^2 p}{\langle p \rangle \circ \pi} \right] (P_{\xi_j, \xi_i} x_{i,r}) \right. \\ &\quad \left. + \frac{\delta}{2} \widetilde{M}_2 \left(\Delta_{UTM}^V \left[\frac{f^2 p}{\langle p \rangle \circ \pi} \right] (P_{\xi_j, \xi_i} x_{i,r}) + E_2 \cdot \left[\frac{f^2 p}{\langle p \rangle \circ \pi} \right] \right) (P_{\xi_j, \xi_i} x_{i,r}) + O(\delta^2) \right\}. \end{aligned}$$

Similarly,

$$\begin{aligned} \mathbb{E}_2 G_{j,s}^2 &= \delta^{\frac{d-1}{2}} \left\{ \widetilde{M}_0 \left[\frac{p}{\langle p \rangle \circ \pi} \right] (P_{\xi_j, \xi_i} x_{i,r}) \right. \\ &\quad \left. + \frac{\delta}{2} \widetilde{M}_2 \left(\Delta_{UTM}^V \left[\frac{p}{\langle p \rangle \circ \pi} \right] + E_2 \cdot \left(\frac{p}{\langle p \rangle \circ \pi} \right) \right) (P_{\xi_j, \xi_i} x_{i,r}) + O(\delta^2) \right\}, \\ \mathbb{E}_2 (F_{j,s} G_{j,s}) &= \delta^{\frac{d-1}{2}} \left\{ \widetilde{M}_0 \left[\frac{f p}{\langle p \rangle \circ \pi} \right] (P_{\xi_j, \xi_i} x_{i,r}) \right. \\ &\quad \left. + \frac{\delta}{2} \widetilde{M}_2 \left(\Delta_{UTM}^V \left[\frac{f p}{\langle p \rangle \circ \pi} \right] + E_2 \cdot \left[\frac{f p}{\langle p \rangle \circ \pi} \right] \right) (P_{\xi_j, \xi_i} x_{i,r}) + O(\delta^2) \right\}. \end{aligned}$$

Since the kernel K is compactly supported and f is Lipschitz (UTM compact), the difference $f(P_{\xi_j, \xi_i} x_{i,r}) - f(x_{i,r})$ is of order $O(d_M(\xi_j, \xi_i)) = O(\epsilon^{\frac{1}{2}})$. Thus

$$\begin{aligned} &|(\mathbb{E}_2 F_{j,s}^2) (\mathbb{E}_1 \mathbb{E}_2 G)^2 + (\mathbb{E}_2 G_{j,s}^2) (\mathbb{E}_1 \mathbb{E}_2 F)^2 - 2\mathbb{E}_2 (F_{j,s} G_{j,s}) (\mathbb{E}_1 \mathbb{E}_2 F) (\mathbb{E}_1 \mathbb{E}_2 G)| \\ &\leq \epsilon^d \delta^{\frac{3(d-1)}{2}} (C' \epsilon + C'' \delta), \quad C' > 0, C'' > 0 \end{aligned}$$

and

$$\mathbb{E}_2^{\xi_j} Y_j^2 - \left(\mathbb{E}_2^{\xi_j} Y_j \right)^2 = O \left(\epsilon^d \delta^{\frac{3(d-1)}{2}} (\epsilon + \delta) \right).$$

If we let C'_1, C''_2 be constants such that

$$C_1'' \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \leq |\mathbb{E}_1 \mathbb{E}_2 G| \leq C_2'' \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}},$$

then for any $\theta \in (0, 1)$, by $\beta = O(\epsilon^2 + \epsilon\delta + \delta^2)$ and Lemma 18,

$$p(N_B, N_F, \beta) \leq N_B \exp \left\{ -\frac{(1-\theta)^2 N_F \epsilon^d \delta^{\frac{d-1}{2}} \beta^2}{C_1(\epsilon + \delta) + O(\epsilon^{\frac{d}{2}} (\epsilon^2 + \delta^2))} \right\} + \exp \left\{ -\frac{\theta^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{C_2 \epsilon + O(\epsilon^2 + \delta^2)} \right\}. \quad (\text{D.14})$$

As pointed out in Remark 7.6, the second term in this bound is the sampling error on the base manifold; the noise error resulted from this term is of the order

$$O \left[\left(N_B \epsilon^{\frac{d}{2}-1} \right)^{-\frac{1}{2}} \right] = O \left(N_B^{-\frac{1}{2}} \epsilon^{\frac{1}{2}-\frac{d}{4}} \right),$$

which is in accordance with the convergence rate obtained in [72]. The first term in the bound reflects the accumulated fibrewise sampling error and grows linearly with respect to the number of fibres sampled, but can be reduced as one increases N_F accordingly (which has an effect of reducing fibrewise sampling errors). The choice of θ is important: as θ increases from 0 to 1, the first term in the bound decreases but the second term increases. One may wish to pick an “optimal” $\theta \in (0, 1)$, but this does not make sense unless one chooses ϵ, δ, N_F appropriately so as to make the sum of the two terms smaller than 1. Let us consider $\theta_* \in (0, 1)$ satisfying

$$(1 - \theta_*)^2 N_F \epsilon^d \delta^{\frac{d-1}{2}} = \theta_*^2 N_B \epsilon^{\frac{d}{2}}, \quad (\text{D.15})$$

or equivalently

$$\epsilon^{\frac{d}{4}} \delta^{\frac{d-1}{4}} \sqrt{\frac{N_F}{N_B}} = \frac{\theta_*}{1 - \theta_*} \Leftrightarrow \theta_* = \frac{\epsilon^{\frac{d}{4}} \delta^{\frac{d-1}{4}} \sqrt{\frac{N_F}{N_B}}}{1 + \epsilon^{\frac{d}{4}} \delta^{\frac{d-1}{4}} \sqrt{\frac{N_F}{N_B}}}. \quad (\text{D.16})$$

Setting $\theta = \theta_*$ in (D.14), we have for some $C > 0$

$$p(N_B, N_F, \beta) \leq (N_B + 1) \exp \left\{ -\frac{\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{C(\epsilon + \delta)} \right\} = \exp \left(-\frac{\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{C(\epsilon + \delta)} + \log(N_B + 1) \right). \quad (\text{D.17})$$

Since

$$\lim_{N_B \rightarrow \infty} \frac{N_B}{\log N_B} = \infty,$$

for any fixed ϵ, δ we have $p(N_B, N_F, \beta) \rightarrow 0$ as $N_B \rightarrow \infty$, as long as one increases N_F accordingly so as to prevent θ_* from approaching 0 or 1; for instance, this is the case if the assumption (9) in Theorem 9 is satisfied:

$$\lim_{\substack{N_B \rightarrow \infty \\ N_F \rightarrow \infty}} \frac{N_F}{N_B} = \beta \in (0, \infty). \quad (\text{D.18})$$

This completes the proof for the pointwise convergence of $\hat{H}_{\epsilon, \delta}^{(0)} f$ in probability.

We now turn to the general case $\alpha \neq 0$. Recall that

$$\hat{H}_{\epsilon,\delta}^{(\alpha)} f(x_{i,r}) = \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}^{(\alpha)}(x_{i,r}, x_{j,s}) f(x_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}^{(\alpha)}(x_{i,r}, x_{j,s})} = \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \frac{\hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) f(x_{j,s})}{\hat{p}_{\epsilon,\delta}^{\alpha}(x_{i,r}) \hat{p}_{\epsilon,\delta}^{\alpha}(x_{j,s})}}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \frac{\hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s})}{\hat{p}_{\epsilon,\delta}^{\alpha}(x_{i,r}) \hat{p}_{\epsilon,\delta}^{\alpha}(x_{j,s})}}$$

where

$$\hat{p}(x_{j,s}) = \sum_{k=1}^{N_B} \sum_{t=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{j,s}, x_{k,t}).$$

By the law of large numbers,

$$\begin{aligned} \lim_{N_B \rightarrow \infty} \frac{1}{N_B} \lim_{N_F \rightarrow \infty} \frac{1}{N_F} \hat{p}(x_{j,s}) &= \int_{UTM} \tilde{K}_{\epsilon,\delta}(x_{i,r}, \eta) p(\eta) d\Theta(\eta) \\ &= \tilde{p}(x_{i,r}) = \mathbb{E}_1 \mathbb{E}_2 [\tilde{K}_{\epsilon,\delta}(x_{j,s}, \cdot)]. \end{aligned}$$

Therefore, as $N_B \rightarrow \infty, N_F \rightarrow \infty$, we expect $\hat{H}_{\epsilon,\delta}^{(\alpha)} f(x_{i,r})$ to converge to

$$\begin{aligned} &\frac{\int_{UTM} \tilde{K}_{\epsilon,\delta}^{(\alpha)}(x_{i,r}, \eta) f(\eta) p(\eta) d\Theta(y, w)}{\int_{UTM} \tilde{K}_{\epsilon,\delta}^{(\alpha)}(x_{i,r}, \eta) p(\eta) d\Theta(y, w)} = \tilde{H}_{\epsilon,\delta}^{(\alpha)} f(x_{i,r}) \\ &= f(x_{i,r}) + \epsilon \frac{m_{21}}{2m_0} \left(\frac{\Delta_{UTM}^H [fp^{1-\alpha}](x_{i,r})}{p^{1-\alpha}(x_{i,r})} - f(x_{i,r}) \frac{\Delta_{UTM}^H p^{1-\alpha}(x_{i,r})}{p^{1-\alpha}(x_{i,r})} \right) \\ &\quad + \delta \frac{m_{22}}{2m_0} \left(\frac{\Delta_{UTM}^V [fp^{1-\alpha}](x_{i,r})}{p^{1-\alpha}(x_{i,r})} - f(x_{i,r}) \frac{\Delta_{UTM}^V p^{1-\alpha}(x_{i,r})}{p^{1-\alpha}(x_{i,r})} \right) + O(\epsilon^2 + \epsilon\delta + \delta^2), \end{aligned}$$

which gives the same bias error $O(\epsilon^2 + \epsilon\delta + \delta^2)$ as in the $\alpha = 0$ case.

It remains to estimate the variance error. By

$$\hat{H}_{\epsilon,\delta}^{\alpha} f(x_{i,r}) = \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) \hat{p}_{\epsilon,\delta}^{-\alpha}(x_{j,s}) f(x_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) \hat{p}_{\epsilon,\delta}^{-\alpha}(x_{j,s})},$$

and

$$\frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) \hat{p}_{\epsilon,\delta}^{-\alpha}(x_{j,s}) f(x_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) \hat{p}_{\epsilon,\delta}^{-\alpha}(x_{j,s})} - \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) \tilde{p}_{\epsilon,\delta}^{-\alpha}(x_{j,s}) f(x_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) \hat{p}_{\epsilon,\delta}^{-\alpha}(x_{j,s})}$$

$$\begin{aligned}
& \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) \left[N_B^{(\alpha)} N_F^{(\alpha)} \hat{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) - \tilde{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) \right] f(x_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) N_B^\alpha N_F^\alpha \hat{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s})} \\
& + \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) \tilde{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) f(x_{j,s}) \times \\
& \left[\frac{- \sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) \left[N_B^\alpha N_F^\alpha \hat{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) - \tilde{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) \right]}{\left(\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) N_B^\alpha N_F^\alpha \hat{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) \right) \left(\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) \tilde{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) \right)} \right] \\
& =: (A) + (B),
\end{aligned}$$

thus if we estimate (A) , (B) by controlling the error

$$\left[N_B^\alpha N_F^\alpha \hat{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) - \tilde{p}_{\epsilon, \delta}^{-\alpha}(x_{j,s}) \right]$$

then it suffices to estimate the variance error caused by

$$\left[\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \frac{\hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s}) f(x_{j,s})}{\tilde{p}_{\epsilon, \delta}^{(\alpha)}(x_{i,r}) \tilde{p}_{\epsilon, \delta}^{(\alpha)}(x_{j,s})} \right] \left[\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \frac{\hat{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s})}{\tilde{p}_{\epsilon, \delta}^{(\alpha)}(x_{i,r}) \tilde{p}_{\epsilon, \delta}^{(\alpha)}(x_{j,s})} \right]^{-1}. \quad (\text{D.19})$$

Our previous proof for the special case $\alpha = 0$ can then be applied to (D.19): the only adjustment is to replace the kernel $\hat{K}_{\epsilon, \delta}(x, y)$ in that proof with the α -normalized kernel

$$\frac{\tilde{K}_{\epsilon, \delta}(x, y)}{\tilde{p}_{\epsilon, \delta}^\alpha(x) \tilde{p}_{\epsilon, \delta}^\alpha(y)}.$$

We would like to estimate the tail probability

$$\mathbb{P} \left\{ \frac{1}{N_B N_F} \hat{p}_{\epsilon, \delta}(x_{j,s}) - \tilde{p}_{\epsilon, \delta}(x_{j,s}) > \beta \right\},$$

but since $\tilde{p}_{\epsilon, \delta}(x_{j,s}) = O\left(\epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}\right)$, it is not lower bounded away from 0 as $\epsilon, \delta \rightarrow 0$. We thus estimate the following tail probability instead:

$$\begin{aligned}
q(N_B, N_F, \beta) &:= \mathbb{P} \left\{ \frac{1}{N_B N_F} \epsilon^{-\frac{d}{2}} \delta^{-\frac{d-1}{2}} \tilde{p}_{\epsilon, \delta}(x_{j,s}) - \epsilon^{-\frac{d}{2}} \delta^{-\frac{d-1}{2}} \tilde{p}_{\epsilon, \delta}(x_{j,s}) > \beta \right\} \\
&= \mathbb{P} \left\{ \frac{1}{N_B N_F} \tilde{p}_{\epsilon, \delta}(x_{j,s}) - \tilde{p}_{\epsilon, \delta}(x_{j,s}) > \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \beta \right\},
\end{aligned}$$

where

$$\hat{p}_{\epsilon, \delta}(x_{j,s}) = \sum_{k=1}^{N_B} \sum_{t=1}^{N_F} \hat{K}_{\epsilon, \delta}(x_{j,s}, x_{k,t}), \quad \tilde{p}_{\epsilon, \delta}(x_{j,s}) = \mathbb{E}_1 \mathbb{E}_2 [\tilde{K}_{\epsilon, \delta}(x_{j,s}, \cdot)].$$

Noting that for some positive constant $C = C(\|K\|_\infty, p_M, p_m, d)$

$$\begin{aligned} |\tilde{K}_{\epsilon, \delta}(x_{i,r}, x_{j,s})| &\leq \|K\|_\infty, \quad |\mathbb{E}_2[\tilde{K}_{\epsilon, \delta}(x_{i,r}, \cdot)]| \leq C\delta^{\frac{d-1}{2}}, \\ |\mathbb{E}_1\mathbb{E}_2[\tilde{K}_{\epsilon, \delta}(x_{i,r}, \cdot)]| &\leq C\epsilon^{\frac{d}{2}}\delta^{\frac{d-1}{2}}, \end{aligned}$$

and by direct computation

$$\mathbb{E}_2^{\xi_j}[\tilde{K}_{\epsilon, \delta}(x_{i,r}, \cdot)]^2 = O\left(\delta^{\frac{d-1}{2}}\right), \quad \mathbb{E}_1[\mathbb{E}_2\tilde{K}_{\epsilon, \delta}(x_{i,r}, \cdot)]^2 = O\left(\epsilon^{\frac{d}{2}}\delta^{d-1}\right),$$

Lemma 18 and $\beta = O(\epsilon^2 + \epsilon\delta + \delta^2)$ gives

$$\begin{aligned} q(N_B, N_F, \beta) &\leq N_B \exp\left\{-\frac{(1-\theta)^2 N_F \epsilon^d \delta^{d-1} \beta^2}{2C_1 \delta^{\frac{d-1}{2}}}\right\} + \exp\left\{-\frac{\theta^2 N_B \epsilon^d \delta^{d-1} \beta^2}{2C_1 \epsilon^{\frac{d}{2}} \delta^{d-1}}\right\} \\ &= N_B \exp\left\{-\frac{(1-\theta)^2 N_F \epsilon^d \delta^{\frac{d-1}{2}} \beta^2}{2C_1}\right\} + \exp\left\{-\frac{\theta^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{2C_1}\right\} \end{aligned}$$

for $C_1 > 0$ some constant. A simple union bound gives

$$\begin{aligned} &\mathbb{P}\left(\bigcup_{j,s} \left\{\left|\frac{1}{N_B N_F} \hat{p}_{\epsilon, \delta}(x_{j,s}) - \tilde{p}_{\epsilon, \delta}(x_{j,s})\right| > \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \beta\right\}\right) \\ &\leq N_B N_F \left[N_B \exp\left\{-\frac{(1-\theta)^2 N_F \epsilon^d \delta^{\frac{d-1}{2}} \beta^2}{2C_1}\right\} + \exp\left\{-\frac{\theta^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{2C_1}\right\} \right] \quad (\text{D.20}) \\ &= N_B (N_B + 1) N_F \exp\left\{-\frac{\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{2C_1}\right\} \quad \text{setting } \theta = \theta_* \text{ as in (D.15).} \end{aligned}$$

We are interested in seeing how this bound compares with the bound in (D.17). As $N_B, N_F \rightarrow \infty$, as long as (D.18) holds,

$$\begin{aligned} &\frac{N_B (N_B + 1) N_F \exp\left\{-\frac{\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{2C_1}\right\}}{(N_B + 1) \exp\left\{-\frac{\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{C(\epsilon + \delta)}\right\}} \\ &= N_B N_F \exp\left\{-\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2 \left[\frac{1}{2C_1} - \frac{1}{C(\epsilon + \delta)}\right]\right\} \rightarrow \infty \quad \text{for small } \epsilon, \delta, \end{aligned}$$

thus the bound in (D.17) is asymptotically negligible compared to the bound in (D.20). This means that when $\alpha \neq 0$ the density estimation in general slows down the convergence rate by a factor $(\epsilon + \delta)^{\frac{1}{2}}$, which is consistent with the conclusion for standard diffusion maps on manifolds [106,9]. Therefore, for probability at least

$$1 - N_B (N_B + 1) N_F \exp\left\{-\frac{\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{2C_1}\right\}$$

we have

$$\left| \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \frac{\hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) f(x_{j,s})}{\tilde{p}_{\epsilon,\delta}^\alpha(x_{i,r}) \tilde{p}_{\epsilon,\delta}^\alpha(x_{j,s})}}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \frac{\hat{K}_{\epsilon,\delta}(x_{i,r}, x_{j,s}) \tilde{p}_{\epsilon,\delta}^{-\alpha}(x_{j,s})}{\tilde{p}_{\epsilon,\delta}^\alpha(x_{i,r}) \tilde{p}_{\epsilon,\delta}^\alpha(x_{j,s})}} - \tilde{H}_{\epsilon,\delta}^{(\alpha)} f(x_{i,r}) \right| \leq \beta$$

as well as

$$\left| \frac{1}{N_B N_F} \hat{p}_{\epsilon,\delta}(x_{j,s}) - \tilde{p}_{\epsilon,\delta}(x_{j,s}) \right| \leq \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \beta \quad \text{for all } 1 \leq j \leq N_B, 1 \leq s \leq N_F.$$

Note that by our assumption

$$0 < p_m \leq p(x, v) \leq p_M < \infty \quad \text{for all } (x, v) \in UTM$$

there exist constants C_1, C_2 such that

$$0 < C_1 < \epsilon^{-\frac{d}{2}} \delta^{-\frac{d-1}{2}} \tilde{p}_{\epsilon,\delta}(x_{j,s}) < C_2 < \infty.$$

For sufficiently small β , we also have

$$0 < C_1 < \frac{1}{N_B N_F} \epsilon^{-\frac{d}{2}} \delta^{-\frac{d-1}{2}} \hat{p}_{\epsilon,\delta}(x_{j,s}) < C_2 < \infty.$$

Thus

$$\left| N_B N_F \hat{p}_{\epsilon,\delta}^{-1}(x_{j,s}) - \tilde{p}_{\epsilon,\delta}^{-1}(x_{j,s}) \right| \leq \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} \beta \cdot \frac{1}{C_1^2 \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}} = \frac{\beta}{C_1^2 \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}}$$

and (A), (B) can be bounded as

$$\begin{aligned} |(A)| &\leq C_2^{(\alpha)} \epsilon^{\frac{\alpha d}{2}} \delta^{\frac{\alpha(d-1)}{2}} \|f\|_\infty \cdot \alpha \left(\frac{2}{C_2 \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}} \right)^{\alpha-1} \frac{\beta}{C_1^2 \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}} = \frac{2^{\alpha-1} \alpha C_2 \|f\|_\infty}{C_1^2} \beta, \\ |(B)| &\leq \frac{C_2^{2\alpha} \epsilon^{\alpha d} \delta^{\alpha(d-1)}}{C_1^{(\alpha)} \epsilon^{\frac{\alpha d}{2}} \delta^{\frac{\alpha(d-1)}{2}}} \|f\|_\infty \cdot \alpha \left(\frac{2}{C_2 \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}} \right)^{\alpha-1} \frac{\beta}{C_1^2 \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}} = \frac{2^{\alpha-1} \alpha C_2^{\alpha+1} \|f\|_\infty}{C_1^{\alpha+2}} \beta. \end{aligned}$$

Since C_1, C_2 only depend on the kernel function K , the dimension d , and p_m, p_M , these bounds ensure that

$$\left| \hat{H}_{\epsilon,\delta}^{(\alpha)} f(x_{i,r}) - \tilde{H}_{\epsilon,\delta}^{(\alpha)} f(x_{i,r}) \right| < C \beta$$

with probability at least

$$1 - N_B (N_B + 1) N_F \exp \left\{ - \frac{\theta_*^2 N_B \epsilon^{\frac{d}{2}} \beta^2}{2 C_1} \right\},$$

where constants C, C_1 only depend on the kernel function K , the dimension d , and p_m, p_M . This establishes the conclusion for all $\alpha \in [0, 1]$. \square

D.2. Sampling from empirical tangent spaces

The following two lemmas from [9] provide estimates for the error of approximating parallel-transport from local PCA. We adapted these lemmas to our notation; note that the statements are more compact than their original form since we assume M is closed.

Lemma 19. Suppose $K_{\text{PCA}} \in C^2([0, 1])$. If $\epsilon_{\text{PCA}} = O\left(N_B^{-\frac{2}{d+2}}\right)$, then, with high probability, the columns of the $D \times d$ matrix O_i determined by local PCA form an orthonormal basis to a d -dimensional subspace of \mathbb{R}^D that deviates from $\iota_* T_{x_i} M$ by $O\left(\epsilon_{\text{PCA}}^{\frac{3}{2}}\right)$, in the following sense:

$$\min_{O \in O(d)} \|O_i^\top \Theta_i - O\|_{\text{HS}} = O\left(\epsilon_{\text{PCA}}^{\frac{3}{2}}\right) = O\left(N_B^{-\frac{3}{d+2}}\right), \quad (\text{D.21})$$

where Θ_i is a $D \times d$ matrix whose columns form an orthonormal basis to $\iota_* T_{x_i} M$. Let the minimizer if (D.21) be

$$\hat{O}_i = \arg \min_{O \in O(d)} \|O_i^\top \Theta_i - O\|_{\text{F}}, \quad (\text{D.22})$$

and denote by Q_i the $D \times d$ matrix

$$Q_i := \Theta_i \hat{O}_i^\top. \quad (\text{D.23})$$

The columns of Q_i form an orthonormal basis to $\iota_* T_{x_i} M$, and

$$\|O_i - Q_i\|_{\text{F}} = O(\epsilon_{\text{PCA}}), \quad (\text{D.24})$$

where $\|\cdot\|_{\text{F}}$ is the matrix Frobenius norm.

Proof. See [9, Lemma B.1]. \square

Lemma 20. Consider points $x_i, x_j \in M$ such that the geodesic distance between them is $O\left(\epsilon^{\frac{1}{2}}\right)$. For $\epsilon_{\text{PCA}} = O\left(N_B^{-\frac{2}{d+2}}\right)$, with high probability, O_{ij} approximates P_{x_i, x_j} in the following sense:

$$O_{ij} \bar{X}_j = (\langle \iota_* P_{x_i, x_j} X(x_j), u_l(x_i) \rangle)_{l=1}^d + O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right), \quad \text{for all } X \in \Gamma(M, TM), \quad (\text{D.25})$$

where $\{u_l(x_i)\}_{l=1}^d$ is an orthonormal set determined by local PCA, and

$$\bar{X}_i \equiv (\langle \iota_* X(x_i), u_l(x_i) \rangle)_{l=1}^d \in \mathbb{R}^d.$$

Proof. See [9, Theorem B.2]. \square

Proof of Theorem 10. By Definition 5.2 (5.2),

$$O_{ji} c_{i,r} = \frac{O_{ji} B_i^\top \bar{\tau}_{i,r}}{\|B_i^\top \bar{\tau}_{i,r}\|}.$$

By Lemma 20,

$$O_{ji}B_i^\top \bar{\tau}_{i,r} = B_j^\top (P_{\xi_j, \xi_i} \bar{\tau}_{i,r}) + O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right),$$

thus

$$\frac{O_{ji}B_i^\top \bar{\tau}_{i,r}}{\|B_i^\top \bar{\tau}_{i,r}\|} = \frac{B_j^\top (P_{\xi_j, \xi_i} \bar{\tau}_{i,r})}{\|B_i^\top \bar{\tau}_{i,r}\|} + O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right),$$

where we used $\|B_j^\top (P_{\xi_j, \xi_i} \bar{\tau}_{i,r})\| \leq \|P_{\xi_j, \xi_i} \bar{\tau}_{i,r}\| = 1$ and

$$\begin{aligned} |\|B_i^\top \bar{\tau}_{i,r}\|_F - 1| &= |\|B_i^\top \bar{\tau}_{i,r}\|_F - \|Q_i^\top \bar{\tau}_{i,r}\|_F| \leq \|B_i^\top \bar{\tau}_{i,r} - Q_i^\top \bar{\tau}_{i,r}\|_F \\ &\leq \|B_i^\top - Q_i^\top\|_F = O(\epsilon_{\text{PCA}}). \end{aligned}$$

Thus

$$\begin{aligned} O_{ji}c_{i,r} - c_{j,s} &= \frac{B_j^\top (P_{\xi_j, \xi_i} \bar{\tau}_{i,r})}{\|B_i^\top \bar{\tau}_{i,r}\|} + O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right) - \frac{B_j^\top \bar{\tau}_{j,s}}{\|B_j^\top \bar{\tau}_{j,s}\|} \\ &= P_{\xi_j, \xi_i} \bar{\tau}_{i,r} - \bar{\tau}_{j,s} + O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right), \\ |\|O_{ji}c_{i,r} - c_{j,s}\|^2 - \|P_{\xi_j, \xi_i} \bar{\tau}_{i,r} - \bar{\tau}_{j,s}\|^2| &= O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right), \end{aligned}$$

and

$$\begin{aligned} K\left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|O_{ji}c_{i,r} - c_{j,s}\|^2}{\delta}\right) &= K\left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} \bar{\tau}_{i,r} - \bar{\tau}_{j,s}\|^2}{\delta}\right) \\ &+ \partial_2 K\left(\frac{\|\xi_i - \xi_j\|^2}{\epsilon}, \frac{\|P_{\xi_j, \xi_i} \bar{\tau}_{i,r} - \bar{\tau}_{j,s}\|^2}{\delta}\right) \cdot \frac{O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right)}{\delta}. \end{aligned}$$

Thus for any function $g \in C^\infty(UTM)$ we have

$$\begin{aligned} &\int_{UTM} \mathcal{K}_{\epsilon, \delta}(\bar{\tau}_{i,r}, \eta) g(\eta) d\Theta(\eta) \\ &= \int_{UTM} \tilde{K}_{\epsilon, \delta}(\bar{\tau}_{i,r}, \eta) g(\eta) d\Theta(\eta) + \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}-1} O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right). \end{aligned}$$

Following the notation used in the proof of Theorem 9, by the law of large numbers

$$\begin{aligned} \lim_{N_B \rightarrow \infty} \lim_{N_F \rightarrow \infty} \frac{1}{N_B N_F} \hat{q}_{\epsilon, \delta}(\bar{\tau}_{i,r}) &= \mathbb{E}_1 \mathbb{E}_2 [\mathcal{K}_{\epsilon, \delta}(\bar{\tau}_{i,r}, \cdot)] \\ &= \mathbb{E}_1 \mathbb{E}_2 [\tilde{K}_{\epsilon, \delta}(\bar{\tau}_{i,r}, \cdot)] + \epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}-1} O\left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right), \end{aligned}$$

hence we expect $\mathcal{H}_{\epsilon, \delta}^{(\alpha)} f(\bar{\tau}_{i,r})$ to converge to

$$\begin{aligned} &\tilde{H}_{\epsilon, \delta}^{(\alpha)} f(\bar{\tau}_{i,r}) + O\left(\delta^{-1} \left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}}\right)\right) \\ &= f(\bar{\tau}_{i,r}) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_{UTM}^H [fp^{1-\alpha}](\bar{\tau}_{i,r})}{p^{1-\alpha}(\bar{\tau}_{i,r})} - f(\bar{\tau}_{i,r}) \frac{\Delta_{UTM}^H p^{1-\alpha}(\bar{\tau}_{i,r})}{p^{1-\alpha}(\bar{\tau}_{i,r})} \right] \end{aligned}$$

$$\begin{aligned}
& + \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_{UTM}^V [fp^{1-\alpha}] (\bar{\tau}_{i,r})}{p^{1-\alpha} (\bar{\tau}_{i,r})} - f (\bar{\tau}_{i,r}) \frac{\Delta_{UTM}^V p^{1-\alpha} (\bar{\tau}_{i,r})}{p^{1-\alpha} (\bar{\tau}_{i,r})} \right] \\
& + O(\epsilon^2 + \epsilon\delta + \delta^2) + O\left(\delta^{-1} \left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}} \right)\right).
\end{aligned}$$

In fact, noting that

$$\frac{1}{\epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} N_B N_F} \hat{q}_{\epsilon,\delta} (\bar{\tau}_{i,r}) = \frac{1}{\epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}} N_B N_F} \hat{p}_{\epsilon,\delta} (\bar{\tau}_{i,r}, \bar{\tau}_{j,s}) + \frac{O\left(\delta^{-1} \left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}} \right)\right)}{\epsilon^{\frac{d}{2}} \delta^{\frac{d-1}{2}}},$$

we have

$$\begin{aligned}
& \epsilon^{\alpha d} \delta^{\alpha(d-1)} N_B^{2\alpha} N_F^{2\alpha} \mathcal{K}_{\epsilon,\delta}^{(\alpha)} (\bar{\tau}_{i,r}, \bar{\tau}_{j,s}) \\
& = \epsilon^{\alpha d} \delta^{\alpha(d-1)} N_B^{2\alpha} N_F^{2\alpha} K_{\epsilon,\delta}^{(\alpha)} (\bar{\tau}_{i,r}, \bar{\tau}_{j,s}) + O\left(\delta^{-1} \left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}} \right)\right).
\end{aligned}$$

Consequently,

$$\begin{aligned}
\mathcal{K}_{\epsilon,\delta}^{(\alpha)} f (\bar{\tau}_{i,r}) & = \frac{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \tilde{K}_{\epsilon,\delta}^{(\alpha)} (\bar{\tau}_{i,r}, \bar{\tau}_{j,s}) f (\bar{\tau}_{j,s})}{\sum_{j=1}^{N_B} \sum_{s=1}^{N_F} \tilde{K}_{\epsilon,\delta}^{(\alpha)} (\bar{\tau}_{i,r}, \bar{\tau}_{j,s})} + O\left(\delta^{-1} \left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}} \right)\right) \\
& = \hat{H}_{\epsilon,\delta}^{(\alpha)} f (\bar{\tau}_{i,r}) + O\left(\delta^{-1} \left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}} \right)\right).
\end{aligned}$$

Under the assumption that

$$\delta^{-1} \left(\epsilon_{\text{PCA}}^{\frac{1}{2}} + \epsilon^{\frac{3}{2}} \right) \rightarrow 0 \quad \text{as } \epsilon \rightarrow 0,$$

we can apply Theorem 9. This completes the proof of Theorem 10. \square

References

- [1] F.R. Chung, Spectral Graph Theory, CBMS Regional Conference Series in Mathematics, vol. 92, American Mathematical Society, 1997.
- [2] S. Rosenberg, The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds, London Mathematical Society Student Texts, vol. 31, Cambridge University Press, 1997.
- [3] S.T. Roweis, L.K. Saul, Nonlinear dimensionality reduction by locally linear embedding, *Science* 290 (2000) 2323–2326.
- [4] J.B. Tenenbaum, V.d. Silva, J.C. Langford, A global geometric framework for nonlinear dimensionality reduction, *Science* 290 (2000) 2319–2323.
- [5] D.L. Donoho, C. Grimes, Hessian eigenmaps: locally linear embedding techniques for high-dimensional data, *Proc. Natl. Acad. Sci. USA* 100 (2003) 5591–5596.
- [6] Z. Zhang, H. Zha, Principal manifolds and nonlinear dimensionality reduction via tangent space alignment, *SIAM J. Sci. Comput.* 26 (2005) 313–338.
- [7] R.R. Coifman, S. Lafon, Diffusion maps, *Appl. Comput. Harmon. Anal.* 21 (2006) 5–30; Special Issue: Diffusion Maps and Wavelets.
- [8] A. Singer, H.-T. Wu, Orientability and diffusion maps, *Appl. Comput. Harmon. Anal.* 31 (2011) 44–58.
- [9] A. Singer, H.-T. Wu, Vector diffusion maps and the connection Laplacian, *Comm. Pure Appl. Math.* 65 (2012) 1067–1144.
- [10] F.A. Velez-Reyes, Miguel Kruse (Eds.), Schrödinger Eigenmaps with Nondiagonal Potentials for Spatial-Spectral Clustering of Hyperspectral Imagery, vol. 9088, 2014.
- [11] E. Goffman, Frame Analysis: An Essay on the Organization of Experience, Harvard University Press, 1974.
- [12] S. Wasserman, K. Faust, Social Network Analysis: Methods and Applications, vol. 8, Cambridge University Press, 1994.
- [13] S.P. Borgatti, A. Mehra, D.J. Brass, G. Labianca, Network analysis in the social sciences, *Science* 323 (2009) 892–895.
- [14] C. Bothorel, J.D. Cruz, M. Magnani, B. Micenková, Clustering attributed graphs: models, measures and methods, *Netw. Sci.* 3 (2015) 408–444.

- [15] G. D'Agostino, A. Scala, Networks of Networks: The Last Frontier of Complexity, vol. 340, Springer, 2014.
- [16] D.Y. Kenett, M. Perc, S. Boccaletti, Networks of networks—an introduction, *Chaos Solitons Fractals* 80 (2015) 1–6.
- [17] S. Boccaletti, G. Bianconi, R. Criado, C.I. Del Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, M. Zanin, The structure and dynamics of multilayer networks, *Phys. Rep.* 544 (2014) 1–122.
- [18] M. Kivelä, A. Arenas, M. Barthelemy, J.P. Gleeson, Y. Moreno, M.A. Porter, Multilayer networks, *J. Complex Netw.* 2 (2014) 203–271.
- [19] G. Bianconi, Multilayer networks: structure and dynamics, *Bull. Amer. Phys. Soc.* (2019).
- [20] R. Al-Aifari, I. Daubechies, Y. Lipman, Continuous Procrustes distance between two surfaces, *Comm. Pure Appl. Math.* 66 (2013) 934–964.
- [21] S.S. Lafon, Diffusion Maps and Geometric Harmonics, Ph.D. thesis, Yale University, 2004.
- [22] R.R. Coifman, S. Lafon, A.B. Lee, M. Maggioni, B. Nadler, F. Warner, S.W. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: diffusion maps, *Proc. Natl. Acad. Sci. USA* 102 (2005) 7426–7431.
- [23] R.R. Coifman, S. Lafon, A.B. Lee, M. Maggioni, B. Nadler, F. Warner, S.W. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: multiscale methods, *Proc. Natl. Acad. Sci. USA* 102 (2005) 7432–7437.
- [24] R.R. Coifman, M. Maggioni, Diffusion wavelets, *Appl. Comput. Harmon. Anal.* 21 (2006) 53–94; Special Issue: Diffusion Maps and Wavelets.
- [25] S. Watanabe, N. Ikeda, Stochastic Differential Equations and Diffusion Processes, Elsevier, 1981.
- [26] R. Durrett, Stochastic Calculus: A Practical Introduction, vol. 6, CRC Press, 1996.
- [27] M. Belkin, P. Niyogi, Towards a theoretical foundation for Laplacian-based manifold methods, in: *Learning Theory*, Springer, 2005, pp. 486–500.
- [28] M. Belkin, P. Niyogi, Convergence of Laplacian eigenmaps, *Adv. Neural Inf. Process. Syst.* 19 (2007) 129.
- [29] P. Bérard, G. Besson, S. Gallot, Embedding Riemannian manifolds by their heat kernel, *Geom. Funct. Anal.* 4 (1994) 373–398.
- [30] R. Bott, L.W. Tu, Differential Forms in Algebraic Topology, Graduate Texts in Mathematics, vol. 82, Springer-Verlag, New York, 1982.
- [31] H.-T. Wu, Embedding Riemannian manifolds by the heat kernel of the connection Laplacian, *Adv. Math.* 304 (2017) 1055–1079.
- [32] N. El Karoui, H.-t. Wu, Graph connection Laplacian and random matrices with random blocks, *Inf. Inference* 4 (2015) 1–44.
- [33] N. El Karoui, H.-T. Wu, Graph connection Laplacian methods can be made robust to noise, *Ann. Statist.* 44 (2016) 346–372.
- [34] K. Turner, S. Mukherjee, D.M. Boyer, Persistent homology transform for modeling shapes and surfaces, *Inf. Inference* (2014).
- [35] A. Singer, H.-T. Wu, Spectral convergence of the connection Laplacian from random samples, *Inf. Inference* 6 (2016) 58–123.
- [36] D. Giannakis, Dynamics-adapted cone kernels, *SIAM J. Appl. Dyn. Syst.* 14 (2015) 556–608.
- [37] Z. Zhao, D. Giannakis, Analog forecasting with dynamics-adapted kernels, *Nonlinearity* 29 (2016) 2888.
- [38] T. Berry, T. Sauer, Local kernels and the geometric structure of data, *Appl. Comput. Harmon. Anal.* 40 (2016) 439–469.
- [39] L. Hörmander, Hypoelliptic second order differential equations, *Acta Math.* 119 (1967) 147–171.
- [40] A.L. Besse, Einstein Manifolds, Springer, 2007.
- [41] V.G. Kim, W. Li, N.J. Mitra, S. DiVerdi, T. Funkhouser, Exploring collections of 3D models using fuzzy correspondences, *ACM Trans. Graph.* 31 (2012) 54:1–54:11.
- [42] C. Wang, S. Mahadevan, Manifold alignment without correspondence, in: *IJCAI*, vol. 2, 2009, p. 3.
- [43] U. Boscain, R.A. Chertovskih, J.P. Gauthier, A.O. Remizov, Hypoelliptic diffusion and human vision: a semidiscrete new twist, *SIAM J. Imaging Sci.* 7 (2014) 669–695.
- [44] N. Berline, E. Getzler, M. Vergne, Heat Kernels and Dirac Operators, 1992 ed., Grundlehren Text Editions, Springer, 2003.
- [45] P.W. Michor, Topics in Differential Geometry, vol. 93, American Mathematical Soc., 2008.
- [46] C. Taubes, Differential Geometry: Bundles, Connections, Metrics and Curvature, vol. 23, Oxford University Press, 2011.
- [47] T. Gao, J. Brodzki, S. Mukherjee, The geometry of synchronization problems and learning group actions, arXiv preprint, arXiv:1610.09051, 2016.
- [48] J. Hansen, R. Ghrist, Toward a spectral theory of cellular sheaves, arXiv preprint, arXiv:1808.01513, 2018.
- [49] A.M. Bronstein, M.M. Bronstein, R. Kimmel, Numerical Geometry of Non-rigid Shapes, Springer Science & Business Media, 2008.
- [50] Q. Huang, G. Zhang, L. Gao, S. Hu, A. Bustcher, L. Guibas, An optimization approach for extracting and encoding consistent maps in a shape collection, *ACM Trans. Graph.* 31 (2012) 125:1–125:11.
- [51] Q. Huang, L. Guibas, Consistent shape maps via semidefinite programming, in: *Computer Graphics Forum*, in: *Proc. Eurographics Symposium on Geometry Processing (SGP)*, vol. 32, 2013, pp. 177–186.
- [52] R. Lai, H. Zhao, Multiscale nonrigid point cloud registration using rotation-invariant sliced-Wasserstein distance via Laplace–Beltrami eigenmap, *SIAM J. Imaging Sci.* 10 (2017) 449–483.
- [53] K. Grove, H. Karcher, E.A. Ruh, Jacobi fields and Finsler metrics on compact Lie groups with an application to differentiable pinching problems, *Math. Ann.* 211 (1974) 7–21.
- [54] D. Kazhdan, On ϵ -representations, *Israel J. Math.* 43 (1982) 315–323.
- [55] M. De Chiffre, L. Glebsky, A. Lubotzky, A. Thom, Stability, cohomology vanishing, and non-approximable groups, arXiv preprint, arXiv:1711.10238, 2017.
- [56] R. Hermann, A sufficient condition that a mapping of Riemannian manifolds be a fibre bundle, *Proc. Amer. Math. Soc.* 11 (1960) 236–242.

- [57] U. Boscain, R. Neel, L. Rizzi, Intrinsic random walks and sub-Laplacians in sub-Riemannian geometry, *Adv. Math.* 314 (2017) 124–184.
- [58] K. Itô, The Brownian motion and tensor fields on Riemannian manifold, in: *Proc. Int. Congr. Math., Stockholm, 1962*.
- [59] E.P. Hsu, *Stochastic Analysis on Manifolds*, vol. 38, American Mathematical Soc., 2002.
- [60] A. Rangarajan, H. Chui, F.L. Bookstein, The softassign Procrustes matching algorithm, in: *Information Processing in Medical Imaging*, Springer, 1997, pp. 29–42.
- [61] J. Solomon, A. Nguyen, A. Butscher, M. Ben-Chen, L. Guibas, Soft maps between surfaces, in: *Computer Graphics Forum*, vol. 31, Wiley Online Library, 2012, pp. 1617–1626.
- [62] C. Villani, *Topics in Optimal Transportation*, Graduate Studies in Mathematics, American Mathematical Society, 2003.
- [63] C. Villani, *Optimal Transport: Old and New*, 2009 ed., Grundlehren der mathematischen Wissenschaften, Springer, 2008.
- [64] P.B. Gilkey, J.V. Leahy, J. Park, *Spectral Geometry, Riemannian Submersions, and the Gromov-Lawson Conjecture*, vol. 30, CRC Press, 1999.
- [65] K. Liu, W. Zhang, Adiabatic limits and foliations, arXiv preprint, arXiv:math/9912223, 1999.
- [66] J.-M. Bismut, *Hypoelliptic Laplacian and Bott-Chern Cohomology: A Theorem of Riemann-Roch-Grothendieck in Complex Geometry*, Progress in Mathematics, vol. 305, Birkhäuser, Basel, 2013.
- [67] A. Singer, R. Erban, I.G. Kevrekidis, R.R. Coifman, Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps, *Proc. Natl. Acad. Sci. USA* 106 (2009) 16090–16095.
- [68] P. Bühlmann, S. Van De Geer, *Statistics for High-Dimensional Data: Methods, Theory and Applications*, Springer Science & Business Media, 2011.
- [69] P. Rigollet, J.-C. Hütter, *High Dimensional Statistics*, Lecture Notes, MIT, 2017.
- [70] I. Chavel, *Riemannian Geometry: A Modern Introduction*, 2nd ed., Cambridge Studies in Advanced Mathematics, vol. 98, Cambridge University Press, 2006.
- [71] T. Gao, *Hypoelliptic Diffusion Maps and Their Applications in Automated Geometric Morphometrics*, Ph.D. thesis, Duke University, 2015.
- [72] A. Singer, From graph to manifold Laplacian: the convergence rate, *Appl. Comput. Harmon. Anal.* 21 (2006) 128–134.
- [73] V.A. Epanechnikov, Non-parametric estimation of a multivariate probability density, *Theory Probab. Appl.* 14 (1969) 153–158.
- [74] A.V. Little, M. Maggioni, L. Rosasco, Multiscale geometric methods for estimating intrinsic dimension, in: *Proc. SampTA*, 2011.
- [75] M.E. Taylor, *Noncommutative Harmonic Analysis*, vol. 22, American Mathematical Soc., 1990.
- [76] P. Mitteroecker, P. Gunz, Advances in geometric morphometrics, *Evol. Biol.* 36 (2009) 235–247.
- [77] I.L. Dryden, K.V. Mardia, *Statistical Shape Analysis*, vol. 4, John Wiley & Sons, New York, 1998.
- [78] D.G. Kendall, Shape manifolds, Procrustean metrics, and complex projective spaces, *Bull. Lond. Math. Soc.* 16 (1984) 81–121.
- [79] D.S. Fry, *Shape Recognition Using Metrics on the Space of Shapes*, Ph.D. thesis, Harvard University, Cambridge, MA, USA, 1993, UMI Order No. GAX94-12337.
- [80] D.M. Boyer, Y. Lipman, E.St. Clair, J. Puente, B.A. Patel, T. Funkhouser, J. Jernvall, I. Daubechies, Algorithms to automatically quantify the geometric similarity of anatomical surfaces, *Proc. Natl. Acad. Sci. USA* 108 (2011) 18221–18226.
- [81] J. Puente, *Distances and Algorithms to Compare Sets of Shapes for Automated Biological Morphometrics*, Ph.D. thesis, Princeton University, 2013.
- [82] Y. Lipman, I. Daubechies, Conformal Wasserstein distances: comparing surfaces in polynomial time, *Adv. Math.* 227 (2011) 1047–1077.
- [83] D.M. Boyer, L. Costeur, Y. Lipman, Earliest record of *Platychoerops* (Primates, Plesiadapidae), a new species from Mouras Quarry, Mont de Berre, France, *Amer. J. Phys. Anthropol.* 149 (2012) 329–346.
- [84] Y. Lipman, J. Puente, I. Daubechies, Conformal Wasserstein distance: II. Computational aspects and extensions, *Math. Comput.* 82 (2013).
- [85] D.M. Boyer, J. Puente, J.T. Gladman, C. Glynn, S. Mukherjee, G.S. Yapuncich, I. Daubechies, A new fully automated approach for aligning and comparing shapes, *Anat. Rec.* 298 (2015) 249–276.
- [86] P. Koehl, J. Hass, Landmark-free geometric methods in biological shape analysis, *J. R. Soc. Interface* 12 (2015) 20150795.
- [87] N.S. Vitek, C.L. Manz, T. Gao, J.I. Bloch, S.G. Strait, D.M. Boyer, Semi-supervised determination of pseudocryptic morphotypes using observer-free characterizations of anatomical alignment and shape, *Ecol. Evol.* 7 (2017) 5041–5055.
- [88] B.R. Hassett, T. Lewis-Bale, Comparison of 3D landmark and 3D dense cloud approaches to hominin mandible morphometrics using structure-from-motion, *Archaeometry* 59 (2017) 191–203, ARCH-05-0070-2015.R2.
- [89] T. Gao, G.S. Yapuncich, I. Daubechies, S. Mukherjee, D.M. Boyer, Development and assessment of fully automated and globally transitive geometric morphometric methods, with application to a biological comparative dataset with high interspecific variation, *Anat. Rec.* 301 (2018) 636–658.
- [90] T. Gao, S.Z. Kovalsky, I. Daubechies, Gaussian process landmarking on manifolds, *SIAM J. Math. Data Sci.* 1 (2019) 208–236.
- [91] T. Gao, S.Z. Kovalsky, D.M. Boyer, I. Daubechies, Gaussian process landmarking for three-dimensional geometric morphometrics, *SIAM J. Math. Data Sci.* 1 (2019) 237–267.
- [92] M.L. Zelditch, D.L. Swiderski, D.H. Sheets, W.L. Fink, *Geometric Morphometrics for Biologists*, Elsevier Academic Press, San Diego, 2004.
- [93] D.F. Wiley, N. Amenta, D.A. Alcantara, D. Ghosh, Y.J. Kil, E. Delson, W. Harcourt-Smith, F.J. Rohlf, K. St John, B. Hamann, Evolutionary morphing, in: *Visualization, 2005, VIS 05, IEEE, IEEE, 2005*, pp. 431–438.
- [94] P.D. Polly, N. MacLeod, Locomotion in fossil Carnivora: an application of eigensurface analysis for morphometric comparison of 3D surfaces, *Palaeontol. Electronica* 11 (2008) 10–13.
- [95] A.S. Bandeira, A. Singer, D.A. Spielman, A Cheeger inequality for the graph connection Laplacian, *SIAM J. Matrix Anal. Appl.* 34 (2013) 1611–1630.

- [96] S. Lafon, A.B. Lee, Diffusion maps and coarse-graining: a unified framework for dimensionality reduction, graph partitioning, and data set parameterization, *IEEE Trans. Pattern Anal. Mach. Intell.* 28 (2006) 1393–1403.
- [97] C. Ehresmann, *Les Connexions Infinitésimales dans un Espace Fibré Différentiable*, Sémin. Bourbaki 1 (1950) 153–168.
- [98] D.W. Stroock, *An Introduction to the Analysis of Paths on a Riemannian Manifold*, Mathematical Surveys and Monographs, vol. 74, American Mathematical Soc., 2005.
- [99] F. Baudoin, Sub-Laplacians and hypoelliptic operators on totally geodesic Riemannian foliations, arXiv preprint, arXiv: 1410.3268, 2014.
- [100] L. Bérard, J.-P. Bourguignon, Laplacians and Riemannian submersions with totally geodesic fibers, *Illinois J. Math.* 26 (1982) 181–200.
- [101] B. O'Neill, The fundamental equations of a submersion, *Michigan Math. J.* 13 (1966) 459–469.
- [102] M.P. Do Carmo, *Riemannian Geometry*, Springer, 1992.
- [103] S. Sasaki, On the differential geometry of tangent bundles of Riemannian manifolds, *Tohoku Math. J. (2)* 10 (1958) 338–354.
- [104] S. Sasaki, On the differential geometry of tangent bundles of Riemannian manifolds, II, *Tohoku Math. J. (2)* 14 (1962) 146–155.
- [105] P. Petersen, *Riemannian Geometry*, Graduate Texts in Mathematics, vol. 171, Springer Science & Business Media, 2006.
- [106] M. Hein, J.-Y. Audibert, U. Von Luxburg, Graph Laplacians and their convergence on random neighborhood graphs, *J. Mach. Learn. Res.* 8 (2007) 1325–1368.
- [107] O.G. Smolyanov, H.v. Weizsäcker, O. Wittich, Chernoff's theorem and discrete time approximations of Brownian motion on manifolds, *Potential Anal.* 26 (2007) 1–29.
- [108] F.R. Chung, L. Lu, *Complex Graphs and Networks*, vol. 107, American Mathematical Society, Providence, 2006.