

Geometric Graph Filters and Neural Networks: Limit Properties and Discriminability Trade-Offs

Zhiyang Wang , Graduate Student Member, IEEE, Luana Ruiz , Member, IEEE,
and Alejandro Ribeiro , Senior Member, IEEE

Abstract—This paper studies the relationship between a graph neural network (GNN) and a manifold neural network (MNN) when the graph is constructed from a set of points sampled from the manifold, thus encoding geometric information. We consider convolutional MNNs and GNNs where the manifold and the graph convolutions are respectively defined in terms of the Laplace-Beltrami operator and the graph Laplacian. Using the appropriate kernels, we analyze both dense and moderately sparse graphs. We prove non-asymptotic error bounds showing that convolutional filters and neural networks on these graphs converge to convolutional filters and neural networks on the continuous manifold. As a byproduct of this analysis, we observe an important trade-off between the discriminability of graph filters and their ability to approximate the desired behavior of manifold filters. We then discuss how this trade-off is ameliorated in neural networks due to the frequency mixing property of nonlinearities. We further derive a transferability corollary for geometric graphs sampled from the same manifold. We validate our results numerically on a navigation control problem and a point cloud classification task.

Index Terms—Graph neural networks, manifold filters, manifold neural networks, convergence analysis, discriminability.

I. INTRODUCTION

GEOMETRIC data, or data supported in non-Euclidean domains, is the object of much interest in modern information processing. It arises in a number of applications, including protein function prediction [3], [4], robot path planning [5], [6], 3D shape analysis [7], [8], [9] and wireless resource allocation

[10], [11]. Graph convolutional filters [12], [13] and graph neural networks (GNNs) [14], [15], along with manifold convolutional filters [16] and manifold neural networks (MNNs) [17], [18], [19], are the standard tools for invariant information processing on these domains when they are discrete and continuous respectively. The convolution operation is implemented through information diffusion over the geometric structure, thus enabling invariant and stable representations [20], [21], [22], [23] and feature sharing. The cascading neural network architecture interleaves convolutions and nonlinearities, further expanding the model's expressiveness.

Although there is a clear parallel between graphs and manifolds—the former can be seen as discretizations of the latter—, manifolds are infinite-dimensional continuous latent spaces which can only be accessed by discrete point sampling [7], [24], [25], [26]. In general, we have access to a set of sampling points from the manifold, and build a graph model to approximate the underlying continuous manifold while attempting to retain the local and global geometric structure [7], [11], [27]. GNNs have been shown to do well at processing information over the manifold both experimentally and theoretically [16], [25], [28]. Of particular note, conditions that guarantee asymptotic convergence of graph filters and GNNs to manifold filters and MNNs are known [16].

Asymptotic convergence is a minimal guarantee that can be enriched with non-asymptotic approximation error bounds. These bounds are unknown and they are the focus of this paper. These non-asymptotic approximation error bounds relating graph filters and GNNs to manifold filters and MNNs are important because they inform the practical design on graphs of information processing architectures that we want to deploy on manifolds. In addition, explicit finite-sample error bounds often reveal details about the convergence regime (e.g., rates of convergence and discriminability trade-offs) that are not revealed by their asymptotic counterparts. For example, the non-asymptotic convergence analysis of GNNs on graphs sampled from a graphon (also referred to as a *transferability* analysis) gives a more precise characterization of the discriminability-convergence tradeoff that arises in these GNNs [29], which is not elucidated by the corresponding asymptotic convergence result [30].

Contributions. In this paper, we prove and analyze a non-asymptotic approximation error bound for GNNs on graphs sampled from manifold, thus closing the gap between GNNs

Manuscript received 23 June 2023; revised 25 January 2024; accepted 6 April 2024. Date of publication 25 April 2024; date of current version 7 May 2024. This work was supported in part by the NSF-Simons MoDL, in part by the NSF AI Institutes program, and in part by the NSF HDR TRipods. An earlier version of this paper was presented at the 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) [DOI: 10.1109/ICASSP49357.2023.10097267] and the 2023 IEEE 57th Asilomar Conference on Signals, Systems, and Computers [DOI: 10.1109/IEEECONF59524.2023.10476997]. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. David I. Shuman. (Corresponding author: Zhiyang Wang.)

Zhiyang Wang and Alejandro Ribeiro are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: zhiyangw@seas.upenn.edu; aribeiro@seas.upenn.edu).

Luana Ruiz is with the Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD 21211 USA (e-mail: lrubini1@jhu.edu).

Digital Object Identifier 10.1109/TSP.2024.3392360

and MNNs with an explicit numerical relationship. We start by importing the definition of the manifold filter as a convolutional operation where the diffusions are exponentials of the Laplace-Beltrami (LB) operator $\mathcal{L}$ of the manifold $\mathcal{M} \subset \mathbb{R}^N$ [16]. Given a set of discrete sampling points from the manifold, we describe how to construct both dense and relatively sparse geometric graphs that approximate the underlying manifold in both the spatial and the spectral domains.

Next, we import the concept of Frequency Difference Threshold (FDT) filters (Definition 3) [17] to overcome the challenge of dimensionality associated with the infinite-dimensional spectrum of the LB operator. We show that manifold filters exhibit a trade-off between their discriminability and their ability to be approximated by graph filters, which can be observed in the approximation error bound of geometric graph filters in Theorems 1 and 2.

The same analysis is conducted for GNNs by incorporating nonlinearities (Theorem 3), but in GNNs we hypothesize that the trade-off is alleviated, i.e., that we can recover discriminability, through the addition of these nonlinear operations (Section V). In other words, geometric GNNs can be both discriminative and approximative of the underlying MNNs, which we verify empirically through numerical experiments (Section VI). Finally, we show how our approximation result can be extended, by a triangle inequality argument, to a transferability result for geometric GNNs transferred across geometric graphs of different sizes sampled from the same underlying manifold.

Related Works. GNNs have been thoroughly studied and discussed in a number of previous works [14], [15], [31], [32]. Similarly, the convergence and transferability of neural networks on graphs have been studied in many works including [21], [29], [33], [34]. These works however see the graphs as samples from a graphon model. In our paper, we focus on the manifold as the limit model for large-scale graphs, which is more realistic than graphons as it can incorporate the underlying geometric information. Very relevant to this paper, the spectral convergence of relatively sparse graphs sampled from manifolds has been studied at length in [35]. Other settings of transfer learning on GNNs have been studied in [36], [37] with different graph similarity measurement without focus on the graph limits.

Manifold filters and manifold neural networks have been established and discussed in [16], [17], [18], [24] using different features and metrics. The integral of the heat diffusion process is employed in [16], [17] where it is shown that manifold convolutions are consistent with and can recover both the graph convolution and the time convolution. In [18], a local geodesic system capturing local anisotropic structures is used to define the manifold convolution. Taking a different approach, [24] defines a mixture Gaussian kernel over the spatial domain using patch operator that is constructed locally.

Among all these methods, the definition leveraging the heat diffusion process is the only one to build the connection between graph neural networks and manifold neural networks when the graphs are sampled from the manifold [16]. However,

this paper does not provide an explicit convergence rate or approximation error bound between GNNs and MNNs. In [38], a convergence rate is given for GNNs approximating MNNs, but the result is restricted to input signals with a limited bandwidth. Moreover, the graphs constructed in [16] and [38] still need to be dense with a well-defined Gaussian kernel. In contrast, in this paper we extend the analysis of the convergence of GNNs by proving a non-asymptotic approximation error bound for geometric GNNs sampled from the MNNs where the geometric graphs are either dense or relatively sparse; and we eliminate the bandlimited signal assumption by importing frequency-dependent filters.

Organization. Section II introduces preliminary definitions of graph signal processing and graph neural networks along with manifold signal processing and manifold neural networks. Section III introduces the construction of the geometric graphs from the sampling points of the manifold and the convergence of the geometric graph Laplacians. We show the spectrum of the geometric graphs – both dense and relatively sparse – can approximate the spectrum of the LB operator of the underlying manifold with a bounded error. We then go on to study the approximation of the filtering on the geometric graphs which is further extended to the approximation of geometric GNNs to MNNs in Section IV. Section V discusses the implications of the results derived in Section IV. Section VI illustrates the results of Sections IV with numerical examples. Section VII concludes the paper. Proofs are deferred to the appendices and supplementary materials.

II. PRELIMINARY DEFINITIONS

We start by reviewing the basic architecture of graph neural networks and manifold neural networks.

A. Graph Signal Processing and Graph Neural Networks

Let $\mathbf{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ be an undirected graph with node set $\mathcal{V}$, $|\mathcal{V}| = n$, and edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. The edges in $\mathcal{E}$ may be weighted, in which case the edge weights are assigned by a function $\mathcal{W} : \mathcal{E} \rightarrow \mathbb{R}$. In this paper, we are interested in graph signals $\mathbf{x} \in \mathbb{R}^n$ supported on the nodes. For $1 \leq i \leq n$, $[\mathbf{x}]_i$ represents the value of the signal at node i .

Graph shift. Operating on graph signals $\mathbf{x}$, we define the graph shift operator (GSO) $\mathbf{S} \in \mathbb{R}^{n \times n}$. The GSO is any matrix satisfying $[\mathbf{S}]_{ij} \neq 0$ if and only if $(i, j) \in \mathcal{E}$ or $i = j$, e.g., the adjacency matrix $\mathbf{A}$, $[\mathbf{A}]_{ij} = \mathcal{W}(i, j)$, or the graph Laplacian $\mathbf{L} = \text{diag}(\mathbf{A}\mathbf{1}) - \mathbf{A}$ [13], [39]. The GSO is so called because, at each node i , it has the effect of shifting or diffusing the signal values at i 's neighbors to i , where signal values are aggregated. Explicitly, $[\mathbf{S}\mathbf{x}]_i = \sum_{j, (i, j) \in \mathcal{E}} [\mathbf{S}]_{ij} [\mathbf{x}]_j$. In the case of an undirected graph, the GSO is symmetric and can therefore be decomposed as $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^H$. The eigenvalues in the diagonal matrix $\mathbf{\Lambda}$ are seen as spectral frequencies, and the eigenvectors in $\mathbf{V}$ are the corresponding oscillation modes.

Graph convolution. A graph convolution is defined based on the graph diffusion process with $K_t - 1$ consecutive shifts.

More formally, the graph convolutional filter [13], [15], [40] with coefficients $\{h_k\}_{k=0}^{K_t-1}$ is given by

$$\mathbf{h}_G(\mathbf{S})\mathbf{x} = \sum_{k=0}^{K_t-1} h_k \mathbf{S}^k \mathbf{x}. \quad (1)$$

Plugging the spectral decomposition of $\mathbf{S}$ into (1), we see that in the spectral domain this filter can be represented as

$$\mathbf{V}^H \mathbf{h}_G(\mathbf{S})\mathbf{x} = \sum_{k=0}^{K_t-1} h_k \Lambda^k \mathbf{V}^H \mathbf{x} = h(\Lambda) \mathbf{V}^H \mathbf{x}. \quad (2)$$

Hence, the graph frequency response of the graph convolution is given by $h(\lambda) = \sum_{k=0}^{K_t-1} h_k \lambda^k$, which only depends on the weights h_k and on the eigenvalues of $\mathbf{S}$.

Graph neural network. A graph neural network (GNN) is built by cascading layers that each consists of a bank of filters followed by a point-wise nonlinearity $\sigma: \mathbb{R} \rightarrow \mathbb{R}$, $[\sigma(\mathbf{x})]_i = \sigma([\mathbf{x}]_i)$. The l th layer of a GNN produces F_l signals $\mathbf{x}_l^p$, each called a feature, given by

$$\mathbf{x}_l^p = \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_G^{lpq}(\mathbf{S}) \mathbf{x}_{l-1}^q \right), \quad (3)$$

for $1 \leq p \leq F_l$. At each layer $l = 1, 2, \dots, L$, the number of input and output features are F_{l-1} and F_l respectively. The filter $\mathbf{h}_G^{lpq}(\mathbf{S})$ is as in (1) and maps the q -th feature of layer $l-1$ to the p -th feature of layer l . For simplicity, we write the GNN consisting of L layers like (3) as the map $\Phi_G(\mathbf{H}, \mathbf{S}, \mathbf{x})$, where $\mathbf{H}$ denotes a set of the graph filter coefficients at all layers.

B. Manifold Signal Processing and Manifold Neural Networks

Let $\mathcal{M}$ be a d -dimensional embedded manifold in $\mathbb{R}^N$. For simplicity, in this paper whenever we mention the manifold $\mathcal{M}$, we assume that it is a compact, smooth and differentiable d -dimensional submanifold embedded in $\mathbb{R}^N$. Atop $\mathcal{M}$, signals are defined as scalar functions $f: \mathcal{M} \rightarrow \mathbb{R}$ called manifold signals [23]. The inner product of signals $f, g \in L^2(\mathcal{M})$ is defined as

$$\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} f(x)g(x)d\mu(x), \quad (4)$$

where $d\mu(x)$ is the volume element with respect to the measure μ over $\mathcal{M}$. Similarly, the norm of the manifold signal f is

$$\|f\|_{\mathcal{M}}^2 = \langle f, f \rangle_{\mathcal{M}}. \quad (5)$$

Manifold shift. The manifold $\mathcal{M}$ is locally Euclidean, and the local homeomorphic Euclidean space at each point $x \in \mathcal{M}$ is defined as the tangent space $T_x \mathcal{M}$ [41]. The disjoint union of all tangent spaces over $\mathcal{M}$ is called the tangent bundle and denoted $T\mathcal{M}$. The gradient $\nabla: L^2(\mathcal{M}) \rightarrow L^2(T\mathcal{M})$ is the differentiation operator and maps scalar functions to tangent vector functions over $\mathcal{M}$ [25], [42]. The adjoint of ∇ is the divergence, which is defined as $\text{div}: L^2(T\mathcal{M}) \rightarrow L^2(\mathcal{M})$. The Laplace-Beltrami (LB) operator $\mathcal{L}: L^2(\mathcal{M}) \rightarrow L^2(\mathcal{M})$ is defined as the divergence of the gradient [43]. Formally, the operation is written as

$$\mathcal{L}f = -\text{div} \circ \nabla f. \quad (6)$$

This operator measures the difference between the signal value at a point and the average signal value in the point's neighborhood [25]. This is akin to how the graph Laplacian matrix can be used to compute the total variation of a graph signal [44].

On manifolds, the shift operation is defined based on the LB operator and on the solution to the heat equation (see [23] for a detailed exposition). Explicitly, for a manifold signal f the manifold shift is written as $e^{-dt\mathcal{L}}f$. Since the LB operator is self-adjoint and positive-semidefinite and the manifold $\mathcal{M}$ is compact, $\mathcal{L}$ has real, positive and discrete eigenvalues $\{\lambda_i\}_{i=1}^{\infty}$, which can be written as

$$\mathcal{L}\phi_i = \lambda_i \phi_i, \quad (7)$$

where ϕ_i is the eigenfunction associated with eigenvalue λ_i . The eigenvalues, or manifold frequencies, are ordered in increasing order as $0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$, and the eigenfunctions, or manifold oscillation modes, are orthonormal and form an eigenbasis of $L^2(\mathcal{M})$ in the intrinsic sense. The ϕ_i are also eigenfunctions of the shift operator $e^{-dt\mathcal{L}}$, with corresponding eigenvalues $e^{-dt\lambda_i}$.

Manifold convolution. Integrating $e^{-dt\mathcal{L}}f$ over $[0, \infty)$ yields the infinite-horizon diffusion process

$$g(x) = \int_0^\infty e^{-t\mathcal{L}} f(x) dt. \quad (8)$$

Let $\tilde{h}: \mathbb{R}^+ \rightarrow \mathbb{R}$ denote the filter impulse response. Based on this diffusion process, a manifold filter can be defined via the manifold convolution [23],

$$g(x) = \mathbf{h}f(x) := \int_0^\infty \tilde{h}(t)e^{-t\mathcal{L}} f(x) dt = \mathbf{h}(\mathcal{L})f(x). \quad (9)$$

Note that the map $\mathbf{h} := \mathbf{h}(\mathcal{L})$ is parametric on the LB operator, and is a spatial map acting directly on $x \in \mathcal{M}$.

If we write $[\hat{f}]_i = \langle f, \phi_i \rangle_{L^2(\mathcal{M})} = \int_{\mathcal{M}} f(x)\phi_i(x)d\mu(x)$, the manifold convolution can be represented in the manifold frequency domain as

$$[\hat{g}]_i = \int_0^\infty \tilde{h}(t)e^{-t\lambda_i} dt [\hat{f}]_i. \quad (10)$$

Hence, the frequency response of the filter $\mathbf{h}(\mathcal{L})$ is given by $\hat{h}(\lambda) = \int_0^\infty \tilde{h}(t)e^{-t\lambda} dt$, which only depends on the impulse response $\tilde{h}$ and the LB eigenvalues $\lambda = \lambda_i$. Further summing over all i and projecting back onto the spatial domain, we can alternatively represent $\mathbf{h}(\mathcal{L})$ as

$$g = \mathbf{h}(\mathcal{L})f = \sum_{i=1}^{\infty} \hat{h}(\lambda_i) [\hat{f}]_i \phi_i. \quad (11)$$

Manifold neural network. Similarly to the GNN, we can define the manifold neural network (MNN) as a cascade of layers $l = 1, 2, \dots, L$ each of which consists of a bank of manifold filters and a pointwise nonlinearity σ . Layer l can be explicitly written as

$$f_l^p(x) = \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q(x) \right), \quad (12)$$

where each signal f_l^p , $1 \leq p \leq F_l$ is a different feature. Each layer l maps F_{l-1} input features to F_l output features. For a more succinct representation, in the following the MNN will be denoted $\Phi(\mathbf{H}, \mathcal{L}, f)$, where $\mathbf{H}$ is a function set gathering the impulse responses of the manifold filters $\mathbf{h}_l^{pq}$ for all features and all layers.

III. GEOMETRIC GRAPHS AND LAPLACIAN CONVERGENCE

Given discrete set of points sampled from the manifold $\mathcal{M}$, we can construct a discrete approximation of $\mathcal{M}$ by connecting the sample points by means of a graph. The nodes of the graph are the sample points, and the edges (more specifically, their weights) are defined as some function of the Euclidean distance between each node pair to encode geometric information from the manifold. We henceforth refer to graphs carrying topological manifold information as *geometric graphs*, which is slightly different from the graph theory definition that focuses on the geometric properties of the edges of the graph [45].

Let X be a set of n points $\{x_1, x_2, \dots, x_n\}$ sampled uniformly, identically and independently from the manifold $\mathcal{M}$ according to measure μ . The discrete empirical measure associated with these points is defined as $p_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$, where δ_{x_i} represents the Dirac measure at x_i . For signals $u, v \in L^2(\mathcal{M})$, the inner product associated with measure p_n is defined as

$$\langle u, v \rangle_{L^2(\mathbf{G}_n)} = \int u(x)v(x)dp_n = \frac{1}{n} \sum_{i=1}^n u(x_i)v(x_i) \quad (13)$$

and so the norm in $L^2(\mathbf{G}_n)$ is

$$\|u\|_{L^2(\mathbf{G}_n)}^2 = \langle u, u \rangle_{L^2(\mathbf{G}_n)}.$$

For signals $\mathbf{u}, \mathbf{v} \in L^2(\mathbf{G}_n)$, the inner product is

$$\langle \mathbf{u}, \mathbf{v} \rangle_{L^2(\mathbf{G}_n)} = \frac{1}{n} \sum_{i=1}^n [\mathbf{u}]_i [\mathbf{v}]_i$$

and the norm is $\|\mathbf{u}\|_{L^2(\mathbf{G}_n)}^2 = \langle \mathbf{u}, \mathbf{u} \rangle_{L^2(\mathbf{G}_n)}$.

We construct an undirected *geometric graph* $\mathbf{G}_n$ from the sampled points X by seeing each point as a vertex. Every pair of vertices is connected by edges with weight values determined by a function K_ϵ of their Euclidean distance. Explicitly, the weight of edge (i, j) , denoted w_{ij} , is given by

$$w_{ij} = K_\epsilon \left(\frac{\|x_i - x_j\|^2}{\epsilon} \right), \quad (14)$$

where $\|x_i - x_j\|$ is the Euclidean distance between x_i and x_j . The geometric adjacency matrix $\mathbf{A}_n^\epsilon$ and the geometric graph Laplacian $\mathbf{L}_n^\epsilon$ can therefore be expressed as $[\mathbf{A}_n^\epsilon]_{ij} = w_{ij}$ for $1 \leq i, j \leq n$ and $\mathbf{L}_n^\epsilon = \text{diag}(\mathbf{A}_n^\epsilon \mathbf{1}) - \mathbf{A}_n^\epsilon$ [46].

Since it is constructed from points $x_i \in \mathcal{M}$, the geometric graph can be seen as a discrete approximation, or discretization, of the manifold. To analogously obtain the discretization of a manifold signal $f \in L^2(\mathcal{M})$ on this graph, as well as the reconstruction from the graph signal back to the manifold signal, we define a uniform sampling operator $\mathbf{P}_n : L^2(\mathcal{M}) \rightarrow L^2(\mathbf{G}_n)$ and an interpolation operator $\mathbf{I}_n : L^2(\mathbf{G}_n) \rightarrow L^2(\mathcal{M})$.

The application of the operator $\mathbf{P}_n$ to f yields the *geometric graph signal*

$$\mathbf{f} = \mathbf{P}_n f \text{ with } \mathbf{f}(x_i) = f(x_i), \quad x_i \in X. \quad (15)$$

I.e., $\mathbf{f}$ is a signal on the graph $\mathbf{G}_n$ sharing the values of the manifold signal f at the sampled points X .

We put certain restrictions on the sampling and interpolation operators as follows [47].

Definition 1: We call the sampling and interpolation operators $\mathbf{P}_n$ and $\mathbf{I}_n$ *asymptotically reconstructive* if for any manifold signal $f \in L^2(\mathcal{M})$, it holds

$$\lim_{n \rightarrow \infty} \mathbf{I}_n \mathbf{P}_n f = f. \quad (16)$$

Moreover, the sampling and interpolation operators $\mathbf{P}_n$ and $\mathbf{I}_n$ are bounded if there exists a constant D such that

$$\limsup_{n \in \mathbb{N}} \|\mathbf{P}_n\| \leq D, \quad \limsup_{n \in \mathbb{N}} \|\mathbf{I}_n\| \leq D. \quad (17)$$

Seeing the geometric graph Laplacian $\mathbf{L}_n^\epsilon$ as an operator acting on $\mathbf{f} : X \rightarrow \mathbb{R}$, we can write the diffusion operation at each point x_i explicitly as

$$\mathbf{L}_n^\epsilon \mathbf{f}(x_i) = \sum_{j=1}^n K_\epsilon \left(\frac{\|x_i - x_j\|^2}{\epsilon} \right) (\mathbf{f}(x_i) - \mathbf{f}(x_j)) \quad (18)$$

for $i = 1, 2, \dots, n$. This operation can be further extended to continuous manifold signals $f : \mathcal{M} \rightarrow \mathbb{R}$ as

$$\mathbf{L}_n^\epsilon f(x) = \sum_{j=1}^n K_\epsilon \left(\frac{\|x - x_j\|^2}{\epsilon} \right) (f(x) - f(x_j)) \quad (19)$$

where $x \in \mathcal{M}$. If we additionally extend the set of sampled points from X to all of the manifold $\mathcal{M}$, we obtain the functional approximation of the geometric graph Laplacian

$$\mathbf{L}^\epsilon f(x) = \int_{\mathcal{M}} K_\epsilon \left(\frac{\|x - y\|^2}{\epsilon} \right) (f(x) - f(y)) d\mu(y). \quad (20)$$

The definition of the weight function K_ϵ allows the construction of geometric graphs with different levels of sparsity. With different choices of K_ϵ (i.e., whether K_ϵ has a bounded support, the relationship between ϵ and the number of sampling nodes n , etc.), the graphs can be sparse (average node degree $\Theta(1)$), relatively sparse (average node degree $\Theta(\log n)$) or dense (average node degree $\Theta(n)$). In the following, we will focus on two kernel definitions that allow constructing dense and relatively sparse geometric graphs.

A. Laplacian Convergence: Dense Graphs

Dense geometric graphs can be constructed when pairs of points x_i and x_j are connected with a weight function K_ϵ defined on an unbounded support (i.e. $[0, \infty)$), which connects x_i and x_j regardless of the distance between them, but often with the edge weight inversely proportional to this distance. This results in a dense geometric graph with n^2 edges. In particular, the Gaussian kernel has been widely used to define the weight value function due to the good approximation properties

of the corresponding graph Laplacians vis-à-vis the Laplace-Beltrami operator [48], [49], [50]. In the Gaussian case, the weight function K_ϵ is computed explicitly as

$$K_\epsilon \left(\frac{\|x - y\|^2}{\epsilon} \right) = \frac{1}{n} \frac{1}{\epsilon^{d/2+1} (4\pi)^{d/2}} e^{-\frac{\|x - y\|^2}{4\epsilon}}, \quad (21)$$

with d representing the dimension of the manifold. The consistency of the geometric graph Laplacian constructed with this Gaussian kernel is ensured by a non-asymptotic error bound. Explicitly, the following result quantifies the approximation of the dense geometric graph Laplacian in the weak sense.

Proposition 1: Let $\mathcal{M} \subset \mathbb{R}^N$ be equipped with LB operator $\mathcal{L}$, whose eigendecomposition is given by (7). Let $\mathbf{G}_n$ be a dense geometric graph constructed from n points sampled u.i.d. from $\mathcal{M}$ with the edge weights defined as (14) and (21), $\epsilon = \epsilon(n) > n^{-1/(d+4)}$. Then, with probability at least $1 - \delta$, the following holds for each pair of $\{\lambda_i, \phi_i\}$

$$|\mathbf{L}_n^\epsilon \phi_i(x) - \mathcal{L} \phi_i(x)| \leq \left(C_1 \sqrt{\frac{\ln(1/\delta)}{2n\epsilon^{d+2}}} + C_2 \sqrt{\epsilon} \right) \lambda_i^{\frac{d+2}{4}}. \quad (22)$$

The constants C_1, C_2 depend on the volume of the manifold and are defined in Appendix B.

Proof: See Section B in supplemental materials. $\square$

We can see that the quality of the approximation of $\mathcal{L}$ by $\mathbf{L}_n^\epsilon$ relates not only to the number of sampling points n but also grows with the corresponding eigenvalue λ_i . In high frequency domain, a larger eigenvalue leads to a larger difference between the geometric graph Laplacian $\mathbf{L}_n^\epsilon$ and the Laplace-Beltrami operator $\mathcal{L}$ when operating on the corresponding eigenfunction ϕ_i of $\mathcal{L}$. This can be interpreted to mean that eigenfunctions associated with eigenvalues oscillates faster [51] and have slower convergence.

Based on this approximation result and using the Davis-Kahan theorem [52], we can further derive a non-asymptotic approximation bound relating the spectra of the dense geometric graph Laplacian and of the LB operator in a limited spectrum.

Proposition 2: Let $\mathcal{M} \subset \mathbb{R}^N$ be equipped with LB operator $\mathcal{L}$, whose eigendecomposition is given by (7). Let $\mathbf{L}_n^\epsilon$ be the discrete graph Laplacian of the dense geometric graph $\mathbf{G}_n$ defined as in (19) and (21), with spectrum given by $\{\lambda_{i,n}^\epsilon, \phi_{i,n}^\epsilon\}_{i=1}^n$. Fix $K \in \mathbb{N}^+$ and $\epsilon = \epsilon(n) > n^{-1/(d+4)}$. Then, with probability at least $1 - e^{-2n\epsilon^{d+3}}$, we have

$$|\lambda_i - \lambda_{i,n}^\epsilon| \leq \Omega_{1,K} \sqrt{\epsilon}, \quad \|a_i \phi_{i,n}^\epsilon - \phi_i\| \leq \Omega_{2,K} \sqrt{\epsilon} / \theta, \quad (23)$$

with $a_i \in \{-1, 1\}$ for all $i < K$ and θ the eigengap of $\mathcal{L}$, i.e., $\theta = \min_{1 \leq i \leq K} \{\lambda_i - \lambda_{i-1}, \lambda_{i+1} - \lambda_i\}$. The constants $\Omega_{1,K}, \Omega_{2,K}$ depend on λ_K, d and the volume of $\mathcal{M}$.

Proof: See Section C in supplemental materials. $\square$

Looking at Proposition 2, we can see that the element-wise non-asymptotic convergence of the eigenvalues and eigenfunctions is only guaranteed in a limited part of the spectrum ($i < K$), and that the upper bound is related to the upper limit λ_K . This can be interpreted to mean that as eigenfunctions in the high frequency domain oscillate faster, they are harder to approximate. Therefore, when designing filters, we need to be especially careful when trying to discriminate components in the high frequency domain.

B. Laplacian Convergence: Relatively Sparse Graphs

We can construct relatively sparse graphs by setting the weight function K_ϵ with a bounded support, i.e., only nodes that are within a certain distance of one another can be connected by an edge. We consider the following weight function from [35] with normalization, which has been shown to well approximate the LB operator,

$$K_\epsilon \left(\frac{\|x - y\|^2}{\epsilon} \right) = \frac{1}{n} \frac{d+2}{\epsilon^{d/2+1} \alpha_d} \mathbb{1}_{[0,1]} \left(\frac{\|x - y\|^2}{\epsilon} \right), \quad (24)$$

where α_d is the volume of the unit ball in $\mathbb{R}^d$. As shown by the indicator function, edges only connect points whose distance is within the radius $\sqrt{\epsilon}$. From the theory of random geometric graphs [53], the order of the radius ϵ decides the order of the average node degree. I.e., if ϵ is in the order of $\Theta(1)$, the expected node degree is in the order of $\Theta(n)$, which is a dense regime. If ϵ is in the order of $\Theta((\log(n)/n)^{2/d})$, the expected node degree is in the order of $\Theta(\log(n))$, which is a relatively sparse regime. When setting ϵ in the order of $\Theta(n^{-2/d})$, the graph is sparse with average node degree $\Theta(1)$.

The following proposition provides a non-asymptotic error bound (in the weak sense) for the approximation of the LB operator by the relatively sparse graph Laplacian.

Proposition 3 [35, Theorem 3.3]: Let $\mathcal{M} \in \mathbb{R}^N$ be equipped with LB operator $\mathcal{L}$, whose eigendecomposition is given by (7). Let $\mathbf{G}_n$ be a sparse geometric graph constructed from n points sampled u.i.d. from $\mathcal{M}$ with the edge weights defined as (14) and (24), $\epsilon > (\log(n)/2n)^{2/(d+2)}$. Then, with probability at least $1 - \delta$, the following holds for each pair of $\{\lambda_i, \phi_i\}$

$$|\mathbf{L}_n^\epsilon \phi_i(x) - \mathcal{L} \phi_i(x)| \leq \left(C_3 \sqrt{\frac{\ln(2n/\delta)}{cn\epsilon^{d+2}}} + C_4 \sqrt{\epsilon} \right) \lambda_i^{\frac{d+2}{4}}. \quad (25)$$

The constants C_3, C_4 depend on the volume of the manifold.

Comparing Proposition 1 with Proposition 3, we see that for large enough ϵ the difference between the graph Laplacian and the Laplace-Beltrami operator is in the order of $O(\sqrt{\ln(2n/\delta)/(n\epsilon^{d+2})}, \sqrt{\epsilon})$ in the sparse graph setting, which is slightly larger than the order $O(\sqrt{\ln(1/\delta)/(n\epsilon^{d+2})}, \sqrt{\epsilon})$ we observe in the dense graph case (22). This agrees with the intuition that as dense graphs include more information about the manifold—by taking into account all the connections between each pair of sampling points—, their Laplacian provides a better approximation to the Laplace-Beltrami operator of the underlying manifold than those of relatively sparse graphs.

The differences between the eigenvalues and eigenfunctions can be bounded similarly, based on the Davis-Kahan theorem.

Proposition 4 [35, Theorem 2.4, Theorem 2.6]: Let $\mathcal{M} \subset \mathbb{R}^N$ be equipped with LB operator $\mathcal{L}$, whose eigendecomposition is given by (7). Let $\mathbf{L}_n^\epsilon$ be the discrete graph Laplacian of graph $\mathbf{G}_n$ defined as (19) and (24), with spectrum $\{\lambda_{i,n}^\epsilon, \phi_{i,n}^\epsilon\}_{i=1}^n$. Fix $K \in \mathbb{N}^+$ and assume that $\epsilon = \epsilon(n) \geq (\log(n)/2n)^{2/(d+2)}$. Then, with probability at least $1 - 2ne^{-cn\epsilon^{d+3}}$, we have

$$|\lambda_i - \lambda_{i,n}^\epsilon| \leq C_{K,1} \sqrt{\epsilon}, \quad \|a_i \phi_{i,n}^\epsilon - \phi_i\| \leq C_{K,2} \sqrt{\epsilon} / \theta, \quad (26)$$

with $a_i \in \{-1, 1\}$ for all $i < K$ and θ the eigengap of $\mathcal{L}$, i.e., $\theta = \min_{1 \leq i \leq K} \{\lambda_i - \lambda_{i-1}, \lambda_{i+1} - \lambda_i\}$. The constants $C_{K,1}$, $C_{K,2}$ depend on λ_K , d and the volume of $\mathcal{M}$.

Akin to the result described in Proposition 2, the approximation errors of the eigenvalues and eigenfunctions can only be bounded within a certain range of the spectrum ($< \lambda_K$). Comparing with the bounds given in Proposition 2, the orders of the errors both depend on variable ϵ in the weight function. Meanwhile, for the convergence probability there is a higher probability in the dense geometric graph setting compared with the relatively sparse geometric graph setting, which also supports the observation that dense geometric graphs can provide a better approximation of the underlying manifold than sparse ones.

Remark 1: We note that ϵ in (24) controls the connectivity radius of each node to its neighboring nodes, and that the average vertex degree is given by $\alpha_d n \epsilon^{d/2}$ according to the random geometric graph theory [54]. In the case of Proposition 4, where ϵ is fixed in the order of $(\log(n)/2n)^{2/(d+2)}$, the average degree scales with an order between $O((\log(n))^2/n)$ and $O(\log(n))$. Hence, the graphs in Propositions 3 and 4 are relatively sparse.

IV. GEOMETRIC GNN CONVERGENCE

Equipped with Propositions 2 and 4 relating the spectrum of the manifold and graph Laplacian operators, we will next show that convolutional filters operating on dense or sparse geometric graphs constructed by sampling the underlying manifold give good approximations of manifold filters. As neural networks are cascading structures composing convolutional filters, GNNs inherit this approximation property from graph filters, and thus provide good approximations of MNNs. We begin this section by discussing how the definitions of manifold filters and MNNs can be generalized to the sampled geometric graphs.

A. Geometric Graph Convolution

If we fix the impulse response function $\tilde{h}(t)$, the definition of manifold filtering in (9) indicates that the convolution operation on the manifold is parametric with respect to the LB operator. Therefore, we can replace the LB operator, acting on the continuous manifold signal, by the discrete graph Laplacian $\mathbf{L}_n^\epsilon$ acting on a geometric graph signal as defined in (18). Explicitly,

$$\mathbf{g} = \int_0^\infty \tilde{h}(t) e^{-t\mathbf{L}_n^\epsilon} \mathbf{f} dt := \mathbf{h}(\mathbf{L}_n^\epsilon) \mathbf{f}, \quad \mathbf{g}, \mathbf{f} \in \mathbb{R}^n. \quad (27)$$

This can be interpreted as a discrete geometric graph filtering process in *continuous-time*, where the exponential term $e^{-\mathbf{L}_n^\epsilon}$ should be seen as the GSO. This is slightly different than the graph convolutional filtering $\mathbf{h}_G$ defined in (1), where we assumed a *discrete-time* frame.

Leveraging (11), the spectral representation of the above continuous time geometric graph filter can be written as

$$\mathbf{g} = \sum_{i=1}^n \hat{h}(\lambda_{i,n}^\epsilon) \langle \mathbf{f}, \phi_{i,n}^\epsilon \rangle_{L^2(\mathbf{G}_n)} \phi_{i,n}^\epsilon, \quad (28)$$

where $\{\lambda_{i,n}^\epsilon, \phi_{i,n}^\epsilon\}_{i=1}^n$ is the spectrum of $\mathbf{L}_n^\epsilon$. The spectral representation exposes the total dependency of the filter frequency

representation on the eigenspectrum of the Laplacian operator. This indicates that the relationship between geometric graph filters and manifold filters can be established in the spectral domain using Propositions 2 and 4.

B. Geometric Graph Convolution Convergence

Recall that the pointwise convergence of the eigenspectrum in Proposition 2 and 4 is restricted to a certain spectral range. Yet, the frequency representation of the graph filter has a dependency on all spectral components. Hence, the infinite spectrum of the LB operator inevitably presents a challenge in the convergence analysis of geometric graph filters.

To address this issue, we exploit Weyl's law [55]. This classical result states that for any compact region the eigenvalues $\{\lambda_i\}_{i=1}^\infty$ of $\mathcal{L}$ grow proportionally to $i^{2/d}$. This indicates that large eigenvalues tend to accumulate and that the differences between neighboring eigenvalues tend to become smaller as the eigenvalues grow larger. This phenomenon is formally described in the following lemma.

Lemma 1 [17, Proposition 3]: Consider a d -dimensional manifold $\mathcal{M} \subset \mathbb{R}^N$ and let $\mathcal{L}$ be its LB operator with eigenvalues $\{\lambda_k\}_{k=1}^\infty$. Let C_1 be an arbitrary constant and α_d the volume of the d -dimensional unit ball. Let $\text{Vol}(\mathcal{M})$ denote the volume of manifold $\mathcal{M}$. For any $\alpha > 0$ and $d > 2$, there exists N_1 ,

$$N_1 = \lceil (\alpha d / C_1)^{d/(2-d)} (C_d \text{Vol}(\mathcal{M}))^{2/(2-d)} \rceil \quad (29)$$

such that, for all $k > N_1$, $\lambda_{k+1} - \lambda_k \leq \alpha$.

Equipped with this fact, we can employ a partitioning strategy to separate the infinite spectrum into finite intervals as shown in Definition 2.

Definition 2 [17, Definition 4] (α -separated spectrum): The α -separated spectrum of the LB operator $\mathcal{L}$ is defined as a partition $\Lambda_1(\alpha) \cup \dots \cup \Lambda_N(\alpha)$ satisfying $|\lambda_i - \lambda_j| > \alpha$ for $\lambda_i \in \Lambda_k(\alpha)$ and $\lambda_j \in \Lambda_l(\alpha)$, $k \neq l$.

The α -separated spectrum as defined above can be achieved by means of a α -FDT filter, which is defined as follows.

Definition 3 [17, Definition 5] (α -FDT filter): The α -frequency difference threshold (α -FDT) filter is defined as a filter $\mathbf{h}(\mathcal{L})$ whose frequency response satisfies

$$|\hat{h}(\lambda_i) - \hat{h}(\lambda_j)| \leq \gamma_k \text{ for all } \lambda_i, \lambda_j \in \Lambda_k(\alpha) \quad (30)$$

with $\gamma_k \leq \gamma$ for some $\gamma > 0$ and $k = 1, \dots, N$.

To prove convergence, we will also an assumption on the continuity of the manifold filter, which needs to be a *Lipschitz filter* as defined below.

Definition 4 (*Lipschitz filter*): A filter is A_h -Lipschitz if its frequency response is Lipschitz continuous with Lipschitz constant A_h ,

$$|\hat{h}(a) - \hat{h}(b)| \leq A_h |a - b| \text{ for all } a, b \in (0, \infty). \quad (31)$$

Letting the geometric graph filter (27) be a Lipschitz continuous α -FDT filter, we are ready to prove an approximation error bound on the difference between the outputs of a manifold filter and a geometric graph filter on a dense geometric graph.

Theorem 1 (*Convergence of filters on dense geometric graphs*): Let $\mathcal{M} \subset \mathbb{R}^N$ be equipped with LB operator $\mathcal{L}$, whose

eigendecomposition is given by (7). Let $\mathbf{L}_n^\epsilon$ be the discrete graph Laplacian of the dense graph $\mathbf{G}_n$ defined as in (19) and (21), with spectrum given by $\{\lambda_{i,n}^\epsilon, \phi_{i,n}^\epsilon\}_{i=1}^n$. Fix $K \in \mathbb{N}^+$ and assume that $\epsilon = \epsilon(n) > n^{-1/(d+4)}$. Let $\mathbf{h}(\cdot)$ be the convolutional filter. Under the assumption that the frequency response of filter $\mathbf{h}$ is Lipschitz continuous and α -FDT with $\alpha^2 \gg \epsilon$, $\alpha > C_{\mathcal{M},d}K^{2/d-1}$ and $\gamma = \Omega_{2,K}\sqrt{\epsilon}/\alpha$, with probability at least $1 - e^{-2n\epsilon^{d+3}}$ it holds that

$$\|\mathbf{h}(\mathbf{L}_n^\epsilon)\mathbf{P}_n f - \mathbf{P}_n \mathbf{h}(\mathcal{L})f\|_{L^2(\mathbf{G}_n)} \leq \left(\frac{N\Omega_{2,K}}{\alpha} + A_h\Omega_{1,K} \right) \sqrt{\epsilon} + C_{gc} \sqrt{\frac{\log n}{n}} \quad (32)$$

where N is the partition size of the α -FDT filter and C_{gc} depends on both d and the volume of $\mathcal{M}$.

Proof: See Appendix A. $\square$

From this theorem, we can see that, if we take the probability as $\delta \leq e^{-2n\epsilon^{d+3}}$, the order of ϵ can be derived as $O((\ln(1/\delta)/(2n))^{1/(d+3)})$. Therefore, given enough sampling points, one can guarantee good approximation accuracy with high probability. Also note that a higher dimension leads to a larger ϵ , which results in a larger approximation error. This indicates that it is more difficult to approximate manifolds with higher dimension.

Observe that the partition of the spectrum by the α -FDT filter lifts the limitation on the spectrum required in Propositions 2 and 4. This is achieved by setting α large enough that eigenvalues larger than λ_K are grouped. A smaller K leads to a larger α , to ensure that more eigenvalues in the high frequency domain are grouped. Therefore, we can alleviate the divergence of spectral components associated with large eigenvalues by giving them similar frequency responses, which subsequently leads to similar filter outputs. When K gets larger, the constants $\Omega_{1,K}$ and $\Omega'_{2,K}$ scale with λ_K as Propositions 2 and 4 show. Meanwhile, α can be set smaller, and the approximation error becomes larger.

Another observation we can make from this non-asymptotic error bound is that, for a fixed number of sampling points n , an interesting trade-off arises between the approximative and discriminative capabilities of the filter. A larger α (i.e., a smaller K) means a less discriminative filter as more eigenvalues tend to be grouped and treated similarly with an almost constant frequency response. A larger α leads to a smaller number of partitions, as the number of singletons decreases, which results in the decrease of the error bound in (33). This implies that the error bound becomes smaller, as the errors brought by the approximation of the eigenvalues are alleviated by treating different eigenvalues with similar frequency responses.

For relatively sparse graphs, we can also establish an approximation error bound between the outputs of manifold filters and geometric graph filters (27).

Theorem 2 (Convergence of filters on relatively sparse geometric graphs): Let $\mathcal{M} \subset \mathbb{R}^N$ be equipped with LB operator $\mathcal{L}$, whose eigendecomposition is given by (7). Let $\mathbf{L}_n^\epsilon$ be the discrete graph Laplacian of $\mathbf{G}_n$ defined as in (19) and (24). Fix $K \in \mathbb{N}^+$ and assume that $\epsilon = \epsilon(n) > (\log(n)/n)^{1/d}$. Let $\mathbf{h}(\cdot)$ be

the convolutional filter. Under the assumption that the frequency response of the filter $\mathbf{h}$ is Lipschitz continuous and α -FDT with $\alpha^2 \gg \epsilon$, $\alpha > C_{\mathcal{M},d}K^{1-2/d}$ and $\gamma = C_{K,2}\epsilon/\alpha$, with probability at least $1 - 2n\exp(-Cn\epsilon^{d+3})$ it holds that

$$\|\mathbf{h}(\mathbf{L}_n^\epsilon)\mathbf{P}_n f - \mathbf{P}_n \mathbf{h}(\mathcal{L})f\|_{L^2(\mathbf{G}_n)} \leq \left(\frac{NC_{K,2}}{\alpha} + A_h C_{K,1} \right) \sqrt{\epsilon} + C_{gc} \sqrt{\frac{\log n}{n}} \quad (33)$$

where N is the partition size of the α -FDT filter and C_{gc} depends on both d and the volume of $\mathcal{M}$.

Proof: See Appendix A. $\square$

If we take the probability as $\delta \leq 2ne^{-Cn\epsilon^{d+3}}$, the order of ϵ can be derived as $O((\ln(2n/\delta)/n)^{1/(d+3)})$, which is slightly larger than the order derived in the dense geometric graph setting in Proposition 4. Similar comments apply that, when compared with sparse geometric graph filters, dense geometric graph filters provide a better approximation of the underlying manifold filters because dense graphs carry more information about the geometry of the manifold.

C. Geometric Graph Neural Network

As filtering on the constructed geometric graph can successfully approximate filtering on the manifold, we can cascade geometric graph filters and point-wise nonlinearities σ to construct a geometric GNN to provide a good approximation of the manifold neural network. Given the geometric graph filter defined in (27), the geometric GNN on $\mathbf{G}_n$ can be written as

$$\mathbf{x}_l^p = \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{l-1}^q \right), \quad (34)$$

where $\mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon)$ is the filter at the l -th layer of this GNN mapping the q -th feature in the $l-1$ -th layer to the p -th feature in the l -th layer, with $1 \leq q \leq F_{l-1}$ and $1 \leq p \leq F_l$. We denote the number of features in the l -th layer F_l (we have dropped the subscript n in $\mathbf{x}_l^p$ and $\mathbf{x}_{l-1}^q$ for simplicity). Gathering the filter functions in the set $\mathbf{H}$, this geometric GNN on $\mathbf{G}_n$ can be represented more concisely as the map $\Phi(\mathbf{H}, \mathbf{L}_n, \mathbf{x})$.

D. Geometric GNN Convergence

Imposing a continuity assumption on the nonlinearity function, we can prove the following approximation error bound for geometric GNNs and MNNs.

Assumption 1 (Normalized Lipschitz nonlinearity): The nonlinearity σ is normalized Lipschitz continuous, i.e., $|\sigma(a) - \sigma(b)| \leq |a - b|$, with $\sigma(0) = 0$.

We note that this assumption is reasonable, since most common nonlinearity functions (e.g., the ReLU, the modulus and the sigmoid) are normalized Lipschitz.

Theorem 3: With the same hypotheses and definitions as in Theorem 1 and 2 respectively, let $\Phi(\mathbf{H}, \mathcal{L}, \cdot)$ be an L -layer MNN on $\mathcal{M}$ (12) with $F_0 = F_L = 1$ input and output features and $F_l = F$, $l = 1, 2, \dots, L-1$ features per layer. Let $\Phi(\mathbf{H}, \mathbf{L}_n, \cdot)$ be the GNN with the same architecture applied to the geometric graph $\mathbf{G}_n$. If the nonlinearities satisfy

Assumption 1 and the manifold filters satisfy $\|\mathbf{h}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f - \mathbf{P}_n \mathbf{h}(\mathcal{L}) f\|_{L^2(\mathbf{G}_n)} \leq \Delta_{fil,n}$, it holds that

$$\|\Phi(\mathbf{H}, \mathbf{L}_n^\epsilon, \mathbf{P}_n f) - \mathbf{P}_n \Phi(\mathbf{H}, \mathcal{L}, f)\|_{L^2(\mathbf{G}_n)} \leq L F^{L-1} \Delta_{fil,n}$$

with high probability.

Proof: See Section D in supplemental materials. $\square$

We conclude that the geometric GNN converges to the MNN as long as the geometric graph filter components give good approximations of the corresponding manifold filters when restricted to the sampled geometric graphs. Further, we see that this error consists of two terms. The first term, $L F^{L-1}$, depends on the number of filters and layers in the MNN architecture. More specifically, the approximation bound grows linearly with the number of layers L and exponentially with the number of features F where the rate is determined by L . This term appears because the approximation errors propagate across all the manifold filters in all layers of the MNN. The second term is the approximation error constant of the geometric graph filters $\Delta_{fil,n}$, replaced with the error bounds from Theorem 1 and Theorem 2 for dense and sparse graphs respectively.

We note that the approximation error constant $\Delta_{fil,n}$ emphasizes that the trade-off between approximation and discriminability is inherited from the convolutional filters. However, since neural networks also have nonlinearities in their layers, the geometric GNN has a better ability to both discriminate high frequency components and approximate the MNN with satisfying accuracy. This is due to the fact that nonlinearities have a spectral mixing effect. In the neural network architecture, the frequency components in the high frequency domain can be shifted to low frequency domain, where they can then be discriminated by the filters in the following layer. This role of nonlinearities has also been discussed in the stability analysis of GNNs [20] as well as of MNNs [17].

E. From Convergence to Transferability

Leveraging the non-asymptotic convergence results derived in Theorem 3, we can immediately prove a transferability corollary for geometric GNNs constructed from a common underlying MNN. Explicitly, consider two geometric graphs $\mathbf{G}_{n_1}$ and $\mathbf{G}_{n_2}$ that are either dense (21) or relatively sparse graphs (24). The difference between the outputs of the same neural network (i.e., with the same weights) on these two graphs is bounded by the following corollary.

Corollary 1: Let $\mathcal{M} \subset \mathbb{R}^N$ be equipped with LB operator $\mathcal{L}$. Let $\mathbf{L}_{n_1}^\epsilon$ be the discrete graph Laplacian of the graph $\mathbf{G}_{n_1}$ defined as in (19) and (21) or (24). Let $\mathbf{L}_{n_2}^\epsilon$ be the discrete graph Laplacian of the graph $\mathbf{G}_{n_2}$ constructed in the same manner as $\mathbf{G}_{n_1}$. Let $\Phi(\mathbf{H}, \cdot, \cdot)$ be an L -layer GNN with $F_0 = F_L = 1$ input and output features and $F_l = F, l = 1, 2, \dots, L-1$ features per layer. Assume that the filters and nonlinearity functions are as in Definitions 3, 4 and satisfy Assumption 1 respectively. Then, it holds in high probability that

$$\begin{aligned} & \|\mathbf{I}_{n_1} \Phi(\mathbf{H}, \mathbf{L}_{n_1}^\epsilon, \mathbf{P}_{n_1} f) - \mathbf{I}_{n_2} \Phi(\mathbf{H}, \mathbf{L}_{n_2}^\epsilon, \mathbf{P}_{n_2} f)\| \\ & \leq L F^{L-1} (\Delta_{fil,n_1} + \Delta_{fil,n_2}). \end{aligned} \quad (35)$$

From this corollary, we can observe that the geometric GNN trained on one geometric graph can be directly transferred to another geometric graph if they are constructed in the same way from the same underlying manifold. The difference between the outputs of the GNNs depends on the size of the neural network architecture as well as on the approximation capability of the filter functions. Larger numbers of sample points n_1 and n_2 lead to smaller approximation errors Δ_{fil,n_1} and Δ_{fil,n_2} , that is, better approximation leads to better transferability in geometric GNNs.

The main implication of this corollary is that if we have trained a geometric GNN on a relatively small geometric graph, this GNN can be transferred to another, larger geometric graph with performance guarantees, which is helpful because it is costly to train GNNs on large graphs. Using the bound in Corollary 1, we can determine the minimum number of sampled points needed for training a geometric GNN to meet a given approximation error tolerance. Therefore, the geometric GNN is a suitable model to approximate both continuous MNNs and large geometric GNNs.

V. DISCUSSION

Approximation vs. discriminability tradeoff. In the convergence results for geometric graph filters on both dense graphs (Theorem 1) and relatively sparse graphs (Theorem 2), the approximation error bounds depend on the size of the partition (N), the frequency partition threshold (α), the Lipschitz continuity constant (A_h) and the weight function parameter (ϵ), and the number of nodes in the geometric graph (n). The size of the partition and the frequency partition threshold have a concerted effect in the approximation, as a larger frequency partition threshold leads to a larger number of intervals containing more than one eigenvalue and a smaller number of intervals containing only one eigenvalue. The size of the partition however tends to stay the same or decrease because the number of intervals cannot exceed the number of eigenvalues. Therefore, a larger frequency partition threshold results in a smaller partition size. Combined, these lead to a smaller approximation error bound. Meanwhile, a larger frequency threshold also means that a larger number of eigenvalues are grouped and treated with similar frequency responses. This makes the α -FDT filters less discriminative. A similar story holds for the Lipschitz constant A_h , with smaller Lipschitz constant increasing the approximation error bound but decreasing discriminability. A larger ϵ leads to a larger approximation error bound, but at the same time the variation range γ in the α -FDT filter becomes larger which results in better discriminability. In conclusion, there is a trade-off between the approximation and discriminability properties of both geometric graph filters and geometric GNNs, but in GNNs this trade-off is alleviated thanks to the nonlinearities, as we discuss next.

Graph filters vs. GNNs. Due to the aforementioned trade-off between approximation and discriminability arising from the spectrum partition, geometric graph filters cannot simultaneously discriminate high-frequency components and provide good approximations of manifold filters. In geometric GNNs,

though not showing explicitly in Theorem 3, this trade-off is however alleviated by the use of nonlinearities. When processing the outputs of the filters in the previous layer, the nonlinearities in a given layer have the ability to shift some high frequency components to the low frequency domain, where they can then be discriminated by the filters in the following layer. As a consequence of this frequency scattering behavior, geometric GNNs can be simultaneously approximative and discriminative.

Dense graphs vs. sparse graphs. Depending on whether the graphs are dense or sparse, different weight functions are chosen to construct the geometric graph, leading to different convergence regimes. From the results presented in Proposition 1 and Proposition 2 for dense geometric graphs and Proposition 3 and Proposition 2 for sparse geometric graphs, we can see that, given the same number of sample points n and the same weight function parameter ϵ , dense graphs have a slightly smaller approximation error bound when compared with sparse graphs—thus implying that geometric graph filters and geometric GNNs provide better approximations of manifold filters and MNNs when operating on dense graphs. This can be explained by the fact that dense graphs include more information; for one, they are complete. However, the sparse graph setting is a more realistic model in practice, and in particular it has been employed in many areas such as wireless communication networks and sensor networks. The approximation guarantees derived in this paper apply in either sparsity regime.

Discretization over time. Theoretically, we employ the definitions of geometric graph convolution and manifold convolution in a continuous-time domain as (27) and (9) show. However, in practice we need to operate the filters and neural networks in digital systems, which requires the time horizon is discrete. In the following section, we carry out numerical experiments with geometric graph filters defined in a discrete-time domain. Specifically, we discretize the impulse response function with an interval T_s and replace the function with a series of filter coefficients $h_k = \tilde{h}(kT_s)$ [16]. By fixing the time horizon with finite K_t samples, we can write the practical geometric filter as

$$\mathbf{g} = \sum_{k=0}^{K_t-1} h_k e^{-k\mathbf{L}_n^\epsilon} \mathbf{f}, \quad \mathbf{f}, \mathbf{g} \in \mathbb{R}^n, \quad (36)$$

which recovers the graph convolution definition in (1) with $e^{-\mathbf{L}_n^\epsilon}$ seen as the graph shift operator.

VI. NUMERICAL EXPERIMENTS

A. Navigation Control

We consider a problem of automatic navigation control [56]. We intend to navigate an agent starting from a given point to find a path to the goal without colliding into the obstacles. We uniformly sample n grid points over the free space while avoiding the obstacle areas. We first generate several trajectories with Dijkstra's shortest path algorithm that leads the agent from a randomly selected starting point to the pre-selected goal. Every point on a generated trajectory can be labeled with a 2-d direction vector leading to the next point along the trajectory, which finally reaches the goal. This experiment aims

to learn a successful trajectory starting from any given point over the space by predicting the potential directions of all the unlabeled points based on the labeled directions on the generated trajectories. By connecting every pair of points that are not intercepted by the obstacles, we can construct a geometric graph structure involving the topological space information. The geometric Laplacian matrix of the constructed graph $\mathbf{L}_n$ is calculated based on the weight function defined in (14) with (21). Specifically, the underlying topology information is captured by setting the Euclidean distance between two points as infinity (i.e. weight value as zero) if there is no direct path connecting these two points without colliding into the obstacles.

The input graph signal is the position coordinates of some point over the space and the output is the direction vector pointing to the next point that can lead to the goal point. The learned architectures are tested via randomly generating 100 starting points and computing the trajectory by predicting directions of the points along the trajectory. If the trajectory can reach the goal point without colliding into the obstacles, the trajectory is marked as "success". The performances are measured by calculating the successful rates of all the learned architectures.

Learning architectures and experiment settings. We build dense geometric graphs as the approximations to the models as the regular grid graph structure brings little difference between the dense and sparse graph setting. We set the position of each point as input signals on the graphs. We control the number of labeled points 30% of the total sampled points. The weights of the edges are calculated based on the Euclidean distance between the nodes and the weight function is determined as (21) with $\epsilon = 0.2$. We use the Laplacian matrix as the input graph shift operator. We train and test three architectures, including 2-layer Graph Filters (GF), 2-layer Graph Neural Network (GNN) and 2-layer Lipschitz Graph Neural Network (Lipschitz GNN) which all contain $F_0 = 2$ input features, $F_1 = 128$ and $F_2 = 64$ features with $K_t = 10$ filters in each layer. ReLU is used as the nonlinearity function in GNN and Lipschitz GNN. In Lipschitz GNN architecture, we add a penalty term which is the scaled derivative of the filter function $C_L |\hat{h}'(\lambda)|$ to the loss function similar to the Lipschitz regularization scheme in [57]. A larger penalty constant C_L indicates a smoother filter function. All architectures also include a linear readout layer to map the 64 features on each node to a 2-d direction vector. The loss function is MSE loss and the optimizer is SGD with the learning rate set as 0.0002. The architectures are trained over 30,000 epochs. We run 5 independent realizations and show the average successful rates and the standard deviation across these realizations.

Convergence analysis. The convergence is tested by calculating the norm of the differences between the outputs of the final graph filter layer for each architecture. By training the graph filters and GNNs $\Phi(\mathbf{H}_n, \cdot, \cdot)$ on geometric graphs with $n = 190, 276, 435, 630, 780$ and graph filters and GNNs $\Phi(\mathbf{H}_{n_0}, \cdot, \cdot)$ on graphs with $n_0 = 1225$, we plot the differences between the output on the relatively large enough geometric graphs with size $n_0 = 1225$ measured as $\|\Phi(\mathbf{H}_{n_0}, \mathbf{L}_{n_0}, \mathbf{x}_{n_0}) - \Phi(\mathbf{H}_n, \mathbf{L}_{n_0}, \mathbf{x}_{n_0})\|$. From Figure 2 we can see that the output




Figure 1. Lipschitz GNN trained on geometric graph with $n = 435$ and tested on geometric graphs with $n = 435, 630, 780, 1225$. The red arrows in the figures show the predicted directions for all sampled points over the space.




Figure 2. Differences of the outputs of 2-layer graph filters, 2-layer GNN and 2-layer Lipschitz GNN.




Figure 3. Differences of the outputs of Lipschitz GNN with penalty terms.

differences decrease and converge as the trained geometric graph size grows, which is accordant with Corollary 1. This also verifies the convergence of geometric GNNs as presented in Theorem 3 if we see the large geometric graph as an approximation of the underlying manifold. We can conclude that Lipschitz GNNs have better approximations while GNNs outperform Graph Filters.

We further study the effect of Lipschitz continuity of the graph filter functions by changing the penalty constant C_L . From Figure 3 we can see as the constant grows, the output differences become smaller, which attests our claim that smoother filter functions provide better approximations.

TABLE I
SUCCESSFUL RATES WHEN TESTING THE ARCHITECTURES
TRAINED ON GEOMETRIC GRAPHS WITH $n = 435, 630,$
780, 1225 ON THE GEOMETRIC GRAPH WITH $n_0 = 1225$

	Graph Filter	GNN	Lipschitz GNN
$n = 435$	0.73 ± 0.12	0.71 ± 0.15	0.77 ± 0.14
$n = 630$	0.72 ± 0.10	0.76 ± 0.11	0.80 ± 0.14
$n = 780$	0.80 ± 0.12	0.83 ± 0.09	0.82 ± 0.11
$n = 1225$	0.84 ± 0.08	0.82 ± 0.07	0.84 ± 0.09

Transferability verification. We verify the transferability result presented in Corollary 1 by training the architectures on a sampled geometric graph with size $n = 435, 630, 780, 1225$ and testing these architectures on the large geometric graphs with $n_0 = 1225$. Note that due to the non-scalability of the linear readout layer, we keep the graph filter functions unchanged and only retrain the final linear layer to map the outputs to a 2-d direction vector. The successful rates are shown in Table I. We can observe that the architectures trained on smaller sampled geometric graphs can still perform well when implemented on larger graphs. Moreover, Lipschitz GNN performs slightly better than GNN which outperforms Graph Filters. This can be understood as the Lipschitz continuity of the filter functions and the nonlinearity functions help with the transferability as we have discussed in Section 3.

To make the results more explicit, we show the predicted directions for each point over the space tested with a Lipschitz GNN trained on graph with $n = 435$ in Figure 1. The red arrows point to the learned directions starting from each unlabeled points while the blue arrows are the learned directions for labeled points. We can see that Lipschitz GNN can efficiently learn the successful directions for unlabeled points based on these geometric graphs and can transfer to larger graphs.

B. Point Cloud Model Classification

We further evaluate the convergence results on the ModelNet10 [58] classification problem. The dataset includes 3,991 meshed CAD models from 10 categories for training and 908 models for testing. In each model, n points are uniformly randomly selected to construct geometric graphs to approximate the underlying model, such as chairs, tables. Figure 4 shows the point cloud model with different sampling points. Our goal is to identify the models for chairs from other models.




Figure 4. Point cloud models with 50 and 300 sampling points in each model. Our goal is to identify chair models from other models such as toilet and table.

Learning architectures and experiment settings. We build both dense and sparse geometric graphs as the approximations to the models. We set the coordinates of each point as input signals on the graphs. The weights of the edges are calculated based on the Euclidean distance between the nodes. For the dense geometric graphs, the weight function is determined as (21) with $\epsilon = 0.1$. Similarly, the weight function of a sparse geometric graph is calculated as (24) with $\epsilon = 0.001$ as the threshold. We calculate the Laplacian matrix for each graph as the input graph shift operator. In this experiment, we implement and compare three different architectures, including 2-layer Graph Filters (GF), 2-layer Graph Neural Network (GNN) and 2-layer Lipschitz Graph Neural Network (Lipschitz GNN). The architectures contain $F_0 = 3$ input features which are the 3-d coordinates of each point, $F_1 = 64$ and $F_2 = 32$ features with $K_t = 5$ filters in each layer. In GNN and Lipschitz GNN, ReLU is used as the nonlinearity function. The filters in Lipschitz GNN is regularized as Lipschitz continuous by imposing a penalty term $C_L |\hat{h}'(\lambda)|$ to the loss function with C_L set as 0.3. All architectures include a linear readout layer to map the final classification outputs.

All the architectures are trained by minimizing the cross-entropy loss. We implement an ADAM optimizer with the learning rate set as 0.005 along with the forgetting factors 0.9 and 0.999. We carry out the training for 40 epochs with the size of batches set as 10. We run 5 random dataset partitions and show the average estimation error rates and the standard deviation across these partitions.

Convergence Verification. We verify our convergence results by training the graph filters and GNNs on geometric graphs with $n = 300, 400, 500, 600, 700, 800, 900$ sampled points, and plotting the norm of differences between the final graph filter layer outputs on geometric graphs with size n and on the relatively large enough geometric graphs with size 1000. Figure 5 shows the convergence results for dense graphs and Figure 6 shows the results for sparse graphs.

From the figures we can see that the differences between the outputs on the trained smaller geometric graphs and the relatively large geometric graphs decrease and converge as the trained graph size grows, which verifies Corollary 1. This also proves what we state in Theorem 3 that the approximation errors of geometric GNNs to the MNNs decrease as the number of nodes grows if we see the large enough graph with 1000 nodes as the manifold. We can also observe that Lipschitz GNNs have better approximations than GNNs because of the continuous filter function while GNNs outperform graph filters due to the nonlinearity function.




Figure 5. Differences of the outputs of Lipschitz GNN, GNN and graph filter trained on dense geometric graphs.




Figure 6. Differences of the outputs of Lipschitz GNN, GNN and graph filter trained on sparse geometric graphs.

TABLE II
CLASSIFICATION ERROR RATES FOR MODEL ‘CHAIR’ WHEN TESTING THE ARCHITECTURES TRAINED ON DENSE GEOMETRIC GRAPHS WITH $n = 300, 500, 700, 900$ TO DENSE GEOMETRIC GRAPHS WITH 1000 NODES. AVERAGE OVER 5 DATA REALIZATIONS

	Graph Filters	GNN	Lipschitz GNN
$n = 300$	$21.15\% \pm 3.48\%$	$9.35\% \pm 2.46\%$	$7.63\% \pm 3.36\%$
$n = 500$	$18.09\% \pm 6.28\%$	$7.80\% \pm 3.50\%$	$7.54\% \pm 4.01\%$
$n = 700$	$17.31\% \pm 6.59\%$	$8.16\% \pm 2.95\%$	$7.97\% \pm 2.45\%$
$n = 900$	$15.58\% \pm 4.54\%$	$7.20\% \pm 3.77\%$	$6.68\% \pm 3.94\%$

Transferability verification. We train the architectures on both dense and sparse geometric graphs with $n = 300, 500, 700, 900$ sampled points from each CAD model. To justify the theoretical results that we have stated in Section IV, we test these trained architectures on a relatively large graph containing 1000 sampled points to verify that geometric GNNs have the transferability to large graphs, i.e. the trained geometric graph filters and GNNs can be directly implemented on larger geometric graphs as long as the geometric graphs are constructed in the same manner. The classification error rates are shown in Table II for the dense graph setting and in Table III for the sparse graph setting.

We can see from the results shown in the tables that the trained architectures can still perform well on the relatively large geometric graphs, both for dense and sparse graph settings. We can see that Lipschitz GNNs outperform GNNs, which perform better than graph filters. This attests the effects of the filter function continuity and nonlinearity that we have

TABLE III
CLASSIFICATION ERROR RATES FOR MODEL ‘CHAIR’ WHEN TESTING THE
ARCHITECTURES TRAINED ON SPARSE GEOMETRIC GRAPHS WITH
 $n = 300, 500, 700, 900$ TO SPARSE GEOMETRIC GRAPHS
WITH SIZE 1000. AVERAGE OVER 5 DATA REALIZATIONS

	Graph Filters	GNN	Lipschitz GNN
$n = 300$	$19.83\% \pm 5.94\%$	$7.74\% \pm 4.05\%$	$7.68\% \pm 3.75\%$
$n = 500$	$21.97\% \pm 4.17\%$	$10.10\% \pm 1.40\%$	$8.60\% \pm 2.95\%$
$n = 700$	$13.85\% \pm 3.81\%$	$7.45\% \pm 4.03\%$	$8.02\% \pm 2.77\%$
$n = 900$	$16.62\% \pm 2.38\%$	$7.92\% \pm 3.14\%$	$7.44\% \pm 3.30\%$

discussed in Section IV, which is continuous filter function and nonlinearity function both help with improving transferability. We can also observe that architectures trained on geometric graphs with more number of nodes can achieve better performances on the relatively large graphs. This can be understood when we see this large enough geometric graph with 1000 nodes approximately as the underlying manifold and as Theorem 3 shows the geometric GNNs can give better approximations to the MNNs as the number of nodes grow. This can also be understood as the transferability property presented in Corollary 1, which also indicates that graph filters and GNNs trained on larger geometric graphs have better performance approximations. Furthermore, the results also show that architectures trained on dense geometric graphs transfer better than the ones trained on sparse geometric graphs, which is also accordant with what we have claimed in Section IV.

VII. CONCLUSION

In this paper, we import the definition of manifold convolutional filters with an exponential Laplace-Beltrami operator to process manifold signals. The manifold model is accessible with a set of i.i.d. uniformly sampled points over the manifold. We construct both dense and sparse graph models to approximate the underlying manifold. We first prove the approximation error bounds of discrete graph Laplacians to the LB operator in the spectral domain. We transfer the definition of manifold filter to the constructed graph and prove the graph filter can approximate the manifold filter with a non-asymptotic error bound. The constructed graph filters need to trade-off between the discriminative and approximative powers while GNNs composed with graph filters and nonlinearities can alleviate the trade-off with the frequency mixing by nonlinearities. We conclude that the GNNs are thus both good approximations of MNNs and discriminative. We finally verified our results numerically with a navigation control problem over a manifold.

APPENDIX

A. Proof of Theorem 1 and Theorem 2

We first write out the filter representation as

$$\begin{aligned} \|\mathbf{h}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f - \mathbf{P}_n \mathbf{h}(\mathcal{L}) f\| \leq & \left\| \sum_{i=1}^{\infty} \hat{h}(\lambda_{i,n}^\epsilon) \langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon \right. \\ & \left. - \sum_{i=1}^{\infty} \hat{h}(\lambda_i) \langle f, \phi_i \rangle_{\mathcal{M}} \mathbf{P}_n \phi_i \right\| \end{aligned} \quad (37)$$

We denote the index of partitions that contain a single eigenvalue as a set $\mathcal{K}_s$ ($|\mathcal{K}_s| = N_s$) and the rest as a set $\mathcal{K}_m$ ($|\mathcal{K}_m| = N_m$). We decompose the α -FDT filter function as $\hat{h}(\lambda) = h^{(0)}(\lambda) + \sum_{l \in \mathcal{K}_m} h^{(l)}(\lambda)$ as

$$h^{(0)}(\lambda) = \begin{cases} \hat{h}(\lambda) - \sum_{l \in \mathcal{K}_m} \hat{h}(C_l) & \lambda \in [\Lambda_k(\alpha)]_{k \in \mathcal{K}_s} \\ 0 & \text{otherwise} \end{cases} \quad (38)$$

$$h^{(l)}(\lambda) = \begin{cases} \hat{h}(C_l) & \lambda \in [\Lambda_k(\alpha)]_{k \in \mathcal{K}_s} \\ \hat{h}(\lambda) & \lambda \in \Lambda_l(\alpha) \\ 0 & \text{otherwise} \end{cases} \quad (39)$$

with C_l some constant in $\Lambda_l(\alpha)$. With the triangle inequality, we start by analyzing the output difference of $h^{(0)}(\lambda)$ as

$$\begin{aligned} & \left\| \sum_{i=1}^n h^{(0)}(\lambda_{i,n}^\epsilon) \langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon \right. \\ & \quad \left. - \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle_{\mathcal{M}} \mathbf{P}_n \phi_i \right\| \\ & \leq \left\| \sum_{i=1}^n \left(h^{(0)}(\lambda_{i,n}^\epsilon) - h^{(0)}(\lambda_i) \right) \langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon \right\| \\ & \quad + \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \left(\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon - \langle f, \phi_i \rangle_{\mathcal{M}} \mathbf{P}_n \phi_i \right) \right\|. \end{aligned} \quad (40)$$

The first term in (40) can be bounded by leveraging the A_h -Lipschitz continuity of the frequency response. From the eigenvalue difference in Proposition 4, we can claim that for each eigenvalue $\lambda_i \leq \lambda_K$, we have

$$|\lambda_{i,n}^\epsilon - \lambda_i| \leq \Omega_{1,K} \sqrt{\epsilon}. \quad (41)$$

The square of the first term is bounded as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} \left(h^{(0)}(\lambda_{i,n}^\epsilon) - h^{(0)}(\lambda_i) \right) \langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon \right\|^2 \\ & \leq \sum_{i=1}^{\infty} |h^{(0)}(\lambda_{i,n}^\epsilon) - h^{(0)}(\lambda_i)|^2 |\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n}|^2 \end{aligned} \quad (42)$$

$$\leq \sum_{i=1}^{\infty} A_h^2 |\lambda_{i,n}^\epsilon - \lambda_i|^2 \|\mathbf{P}_n f\|^2 \leq A_h^2 \Omega_{1,K}^2 \epsilon. \quad (43)$$

The second term in (40) can be bounded combined with the convergence of eigenfunctions in (45) as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \left(\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon - \langle f, \phi_i \rangle_{\mathcal{M}} \mathbf{P}_n \phi_i \right) \right\| \\ & \leq \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \left(\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon \right. \right. \\ & \quad \left. \left. - \langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \mathbf{P}_n \phi_i \right) \right\| \\ & \quad + \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \left(\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \mathbf{P}_n \phi_i - \langle f, \phi_i \rangle_{\mathcal{M}} \mathbf{P}_n \phi_i \right) \right\| \end{aligned} \quad (44)$$

From the convergence stated in Theorem 4, we have

$$\|a_i \phi_{i,n}^\epsilon - \phi_i\| \leq \Omega_{2,K} \sqrt{\epsilon}/\theta, \quad (45)$$

with the eigengap $\theta \geq \alpha$ under the α -FDT filter. Therefore, the first term in (44) can be bounded as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) (\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \phi_{i,n}^\epsilon - \langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathcal{M}} \mathbf{P}_n \phi_i) \right\| \\ & \leq \sum_{i=1}^{N_s} \|\mathbf{P}_n f\| \|\phi_{i,n}^\epsilon - \mathbf{P}_n \phi_i\| \leq \frac{N_s \Omega_{2,K}}{\alpha} \sqrt{\epsilon}. \end{aligned} \quad (46)$$

The last equation comes from the definition of norm in $L^2(\mathbf{G}_n)$. The second term in (44) can be written as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) (\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} \mathbf{P}_n \phi_i - \langle f, \phi_i \rangle_{\mathcal{M}} \mathbf{P}_n \phi_i) \right\| \\ & \leq \sum_{i=1}^{\infty} |h^{(0)}(\lambda_i)| |\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} - \langle f, \phi_i \rangle_{\mathcal{M}}| \|\mathbf{P}_n \phi_i\|. \end{aligned} \quad (47)$$

Because $\{x_1, x_2, \dots, x_n\}$ is a set of uniform sampled points from $\mathcal{M}$, based on Theorem 19 in [59] we can claim that

$$|\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} - \langle f, \phi_i \rangle_{\mathcal{M}}| = O\left(\sqrt{\frac{\log n}{n}}\right). \quad (48)$$

Consider the boundedness of frequency response $|h^{(0)}(\lambda)| \leq 1$ and the bounded energy $\|\mathbf{P}_n \phi_i\|$, we have

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} \hat{h}(\lambda_i) (\langle \mathbf{P}_n f, \phi_{i,n}^\epsilon \rangle_{\mathbf{G}_n} - \langle f, \phi_i \rangle_{\mathcal{M}}) \mathbf{P}_n \phi_i \right\| \\ & = O\left(\sqrt{\frac{\log n}{n}}\right). \end{aligned}$$

Combining the above results, we can bound the output difference of $h^{(0)}$. Then we need to analyze the output difference of $h^{(l)}(\lambda)$ and bound this as

$$\begin{aligned} & \left\| \mathbf{P}_n \mathbf{h}^{(l)}(\mathcal{L}) f - \mathbf{h}^{(l)}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f \right\| \\ & \leq \left\| (\hat{h}(C_l) + \gamma) \mathbf{P}_n f - (\hat{h}(C_l) - \gamma) \mathbf{P}_n f \right\| \leq 2\gamma \|\mathbf{P}_n f\|, \end{aligned} \quad (49)$$

where $\mathbf{h}^{(l)}(\mathcal{L})$ and $\mathbf{h}^{(l)}(\mathbf{L}_n^\epsilon)$ are filters with filter function $h^{(l)}(\lambda)$ on the LB operator $\mathcal{L}$ and graph Laplacian $\mathbf{L}_n^\epsilon$ respectively. Combining the filter functions, we can write

$$\begin{aligned} & \left\| \mathbf{P}_n \mathbf{h}(\mathcal{L}) f - \mathbf{h}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f \right\| \\ & = \left\| \mathbf{P}_n \mathbf{h}^{(0)}(\mathcal{L}) f + \mathbf{P}_n \sum_{l \in \mathcal{K}_m} \mathbf{h}^{(l)}(\mathcal{L}) f - \mathbf{h}^{(0)}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f - \sum_{l \in \mathcal{K}_m} \mathbf{h}^{(l)}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f \right\| \\ & \leq \left\| \mathbf{P}_n \mathbf{h}^{(0)}(\mathcal{L}) f - \mathbf{h}^{(0)}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f \right\| + \sum_{l \in \mathcal{K}_m} \left\| \mathbf{P}_n \mathbf{h}^{(l)}(\mathcal{L}) f - \mathbf{h}^{(l)}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f \right\| \\ & \leq A_h \Omega_{1,K} \sqrt{\epsilon} + \frac{N_s \Omega_{2,K}}{\alpha} \sqrt{\epsilon} + N_m \gamma + C_{gc} \sqrt{\frac{\log(n)}{n}} \end{aligned} \quad (50)$$

With $\gamma = \Omega_{2,K} \sqrt{\epsilon}/\alpha$, we have

$$\begin{aligned} & \|\mathbf{h}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f - \mathbf{P}_n \mathbf{h}(\mathcal{L}) f\| \\ & \leq \left(\frac{N \Omega_{2,K}}{\alpha} + A_h \Omega_{1,K} \right) \sqrt{\epsilon} + C_{gc} \sqrt{\frac{\log n}{n}} \end{aligned} \quad (52)$$

We can prove Theorem 2 similarly by importing Proposition 4 into the eigenvalue and eigenfunction differences.

B. Proof of Proposition 1

We decompose the operator difference between the graph Laplacian and the LB operator with an intermediate term $\mathbf{L}^\epsilon$, which is the functional approximation defined in (20). We first focus on the operator difference between $\mathbf{L}^\epsilon$ and $\mathcal{L}$. From [49], we can get the bound as

$$\|\mathbf{L}^\epsilon \phi_i - \mathcal{L} \phi_i\| \leq C \sqrt{\epsilon} \|\phi_i\|_{H^{d/2+1}}, \quad (53)$$

For the Sobolev norm of eigenfunction ϕ_i , according to [60, Lemma 4.4] we have

$$\|\phi_i\|_{H^{d/2+1}} \leq C \lambda_i^{\frac{d+2}{4}}, \quad (54)$$

which leads to

$$\|\mathbf{L}^\epsilon \phi_i - \mathcal{L} \phi_i\| \leq C_1 \sqrt{\epsilon} \lambda_i^{\frac{d+2}{4}}. \quad (55)$$

For the operator difference between $\mathbf{L}_n^\epsilon$ and $\mathbf{L}^\epsilon$ with Hoeffding's inequality as

$$\mathbb{P}(|\mathbf{L}_n^\epsilon \phi_i(x) - \mathbf{L}^\epsilon \phi_i(x)| > \epsilon_1) \leq \exp\left(-\frac{2n\epsilon_1^2}{\|\phi_i\|_{H^{d/2+1}}^2}\right). \quad (56)$$

Therefore, we can claim that with probability at least $1 - \delta$, we have

$$|\mathbf{L}_n^\epsilon \phi_i(x) - \mathbf{L}^\epsilon \phi_i(x)| \leq \sqrt{\frac{\ln 1/\delta}{2n}} \|\phi_i\|_{H^{d/2+1}}. \quad (57)$$

Combining (55) and (57) with triangle inequality, we can get the conclusion in Theorem 1.

C. Proof of Proposition 2

We first import two lemmas to help prove the spectral properties.

Lemma 2: Let $\mathbf{A}, \mathbf{B}$ be self-adjoint operators with $\{\lambda_i(\mathbf{A}), \mathbf{u}_i\}_{i=1}^\infty$ and $\{\lambda_i(\mathbf{B}), \mathbf{w}_i\}_{i=1}^\infty$ as the corresponding spectrum. Let $Pr_{\mathbf{w}_i}$ be the orthogonal projection operation onto the subspace generated by $\mathbf{w}_i$. Then we have

$$\|a_i \mathbf{u}_i - \mathbf{w}_i\| \leq 2 \|\mathbf{u}_i - Pr_{\mathbf{w}_i} \mathbf{u}_i\| \leq \frac{2 \|\mathbf{B} \mathbf{u}_i - \mathbf{A} \mathbf{u}_i\|}{\min_{j \neq i} |\lambda_j(\mathbf{B}) - \lambda_i(\mathbf{A})|}. \quad (58)$$

Proof: The first inequality is directly from [59, Proposition 18]. Let $Pr_{\mathbf{w}_i}^\perp$ be the orthogonal projection onto the complement of the subspace generated by $\mathbf{w}_i$. Then we have

$$\|\mathbf{u}_i - Pr_{\mathbf{w}_i} \mathbf{u}_i\| = \|Pr_{\mathbf{w}_i}^\perp \mathbf{u}_i\| = \left\| \sum_{j \neq i} \langle \mathbf{u}_i, \mathbf{w}_j \rangle \mathbf{w}_j \right\|. \quad (59)$$

Therefore, we have

$$\begin{aligned} & \|Pr_{\mathbf{w}_i}^\perp \mathbf{B}\mathbf{u}_i - Pr_{\mathbf{w}_i}^\perp \mathbf{A}\mathbf{u}_i\| \\ &= \left\| \sum_{j \neq i} \langle \mathbf{B}\mathbf{u}_i, \mathbf{w}_j \rangle \mathbf{w}_j - \sum_{j \neq i} \langle \mathbf{A}\mathbf{u}_i, \mathbf{w}_j \rangle \mathbf{w}_j \right\| \end{aligned} \quad (60)$$

$$= \left\| \langle \mathbf{u}_i, \mathbf{B}\mathbf{w}_j \rangle \mathbf{w}_j - \sum_{j \neq i} \lambda_i(\mathbf{A}) \langle \mathbf{u}_i, \mathbf{w}_j \rangle \mathbf{w}_j \right\| \quad (61)$$

$$= \left\| \sum_{j \neq i} (\lambda_i(\mathbf{B}) - \lambda_i(\mathbf{A})) \langle \mathbf{u}_i, \mathbf{w}_j \rangle \mathbf{w}_j \right\| \quad (62)$$

$$\geq \min_{j \neq i} |\lambda_i(\mathbf{B}) - \lambda_i(\mathbf{A})| \left\| \sum_{j \neq i} \langle \mathbf{u}_i, \mathbf{w}_j \rangle \mathbf{w}_j \right\| \quad (63)$$

$$= \min_{j \neq i} |\lambda_i(\mathbf{B}) - \lambda_i(\mathbf{A})| \|\mathbf{u}_i - Pr_{\mathbf{w}_i} \mathbf{u}_i\|, \quad (64)$$

together with $\|\mathbf{B}\mathbf{u}_i - \mathbf{A}\mathbf{u}_i\| \geq \|Pr_{\mathbf{w}_i}^\perp \mathbf{B}\mathbf{u}_i - Pr_{\mathbf{w}_i}^\perp \mathbf{A}\mathbf{u}_i\|$. We can conclude the proof. $\square$

The following lemma is adapted from [48, Lemma 5c]

Lemma 3: Let $\mathbf{A}, \mathbf{B}$ be self-adjoint operators with $\{\lambda_i(\mathbf{A}), \mathbf{u}_i\}_{i=1}^\infty$ and $\{\lambda_i(\mathbf{B}), \mathbf{w}_i\}_{i=1}^\infty$ as the corresponding spectrum. Then we have

$$|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{B})| = \frac{\langle (\mathbf{A} - \mathbf{B})\mathbf{u}_i, \mathbf{w}_i \rangle}{|\langle \mathbf{u}_i, \mathbf{w}_i \rangle|} \leq \frac{\|(\mathbf{A} - \mathbf{B})\mathbf{u}_i\|}{|\langle \mathbf{u}_i, \mathbf{w}_i \rangle|} \quad (65)$$

With the above lemmas and our proposed Theorem 1, which includes the operator difference, we can prove Theorem 2. We first fix some $K \in \mathbb{N}$, which provides an upper bound for $\lambda_i \leq \lambda_K$ for all $1 \leq i \leq K$. By taking the probability $1 - n^{-2}$ and $\epsilon = n^{-1/(d+4)}$, we can conclude that the operator difference in Theorem 1 can be bounded with order $O(\sqrt{\epsilon})$, with the constant scaling with $\lambda_K^{\frac{d+2}{4}}$. Combine with Lemma 2 and $\theta = \min_{1 \leq j \neq i \leq K} |\lambda_j - \lambda_{i,n}^\epsilon|$, we can get

$$\|a_i \phi_{i,n}^\epsilon - \phi_i\| \leq \frac{C_k}{\theta} \sqrt{\epsilon}, \quad (66)$$

where we denote the constant as Ω_1 to include the effects of K , the eigengap and the volume of $\mathcal{M}$.

This upper bound of the eigenfunction difference leads to $|\langle \mathbf{u}_i, \mathbf{w}_i \rangle| \geq 1 - \Omega_1/2\sqrt{\epsilon} \geq 1$. Combining with Lemma 3, the difference of the eigenvalues can also be bounded in the order of $O(\sqrt{\epsilon})$.

D. Proof of Theorem 3

To bound the output difference of MNNs, we need to write in the form of features of the final layer

$$\begin{aligned} & \|\Phi(\mathbf{H}, \mathbf{L}_n^\epsilon, \mathbf{P}_n f) - \mathbf{P}_n \Phi(\mathbf{H}, \mathcal{L}, f)\| \\ &= \left\| \sum_{q=1}^{F_L} \mathbf{x}_{n,L}^q - \sum_{q=1}^{F_L} \mathbf{P}_n f_L^q \right\| \leq \sum_{q=1}^{F_L} \|\mathbf{x}_{n,L}^q - \mathbf{P}_n f_L^q\|. \end{aligned} \quad (67)$$

By inserting the definitions, we have

$$\begin{aligned} & \|\mathbf{x}_{n,l}^p - \mathbf{P}_n f_l^p\| \\ &= \left\| \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{n,l-1}^q \right) - \mathbf{P}_n \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q \right) \right\| \end{aligned} \quad (68)$$

with $\mathbf{x}_{n,0} = \mathbf{P}_n f$ as the input of the first layer. With a normalized point-wise Lipschitz nonlinearity, we have

$$\begin{aligned} & \|\mathbf{x}_{n,l}^p - \mathbf{P}_n f_l^p\| \\ & \leq \left\| \sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{n,l-1}^q - \mathbf{P}_n \sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q \right\| \end{aligned} \quad (69)$$

$$\leq \sum_{q=1}^{F_{l-1}} \left\| \mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{n,l-1}^q - \mathbf{P}_n \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q \right\| \quad (70)$$

The difference can be further decomposed as

$$\begin{aligned} & \|\mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{n,l-1}^q - \mathbf{P}_n \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q\| \\ & \leq \|\mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{n,l-1}^q - \mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f_{l-1}^q\| \\ & \quad + \|\mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f_{l-1}^q - \mathbf{P}_n \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q\| \end{aligned} \quad (71)$$

$$\begin{aligned} & \leq \|\mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{n,l-1}^q - \mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f_{l-1}^q\| \\ & \quad + \|\mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f_{l-1}^q - \mathbf{P}_n \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q\| \end{aligned} \quad (72)$$

The second term can be bounded with $\|\mathbf{h}(\mathbf{L}_n^\epsilon) \mathbf{P}_n f - \mathbf{P}_n \mathbf{h}(\mathcal{L}) f\|_{L^2(\mathbf{G}_n)} \leq \Delta_{fil,n}$. The first term can be decomposed by Cauchy-Schwartz inequality and non-amplifying of the filter functions as

$$\begin{aligned} \|\mathbf{x}_{n,l}^p - \mathbf{P}_n f_l^p\| & \leq \sum_{q=1}^{F_{l-1}} \Delta_{fil,n} \|\mathbf{x}_{n,l-1}^q\| \\ & \quad + \sum_{q=1}^{F_{l-1}} \|\mathbf{x}_{l-1}^q - \mathbf{P}_n f_{l-1}^q\|, \end{aligned} \quad (73)$$

where C_{per} representing the constant in the error bound of manifold filters in (33). To solve this recursion, we need to compute the bound for $\|\mathbf{x}_l^p\|$. By normalized Lipschitz continuity of σ and the fact that $\sigma(0) = 0$, we can get

$$\begin{aligned} \|\mathbf{x}_l^p\| & \leq \left\| \sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon) \mathbf{x}_{l-1}^q \right\| \leq \sum_{q=1}^{F_{l-1}} \|\mathbf{h}_l^{pq}(\mathbf{L}_n^\epsilon)\| \|\mathbf{x}_{l-1}^q\| \\ & \leq \sum_{q=1}^{F_{l-1}} \|\mathbf{x}_{l-1}^q\| \leq \prod_{l'=1}^{l-1} F_{l'} \sum_{q=1}^{F_0} \|\mathbf{x}^q\|. \end{aligned} \quad (74)$$

Insert this conclusion back to solve the recursion, we can get

$$\|\mathbf{x}_{n,l}^p - \mathbf{P}_n f_l^p\| \leq l \Delta_{fil,n} \left(\prod_{l'=1}^{l-1} F_{l'} \right) \sum_{q=1}^{F_0} \|\mathbf{x}^q\|. \quad (75)$$

Replace l with L we can obtain

$$\begin{aligned} & \|\Phi(\mathbf{H}, \mathbf{L}_n^\epsilon, \mathbf{P}_n f) - \mathbf{P}_n \Phi(\mathbf{H}, \mathcal{L}, f)\| \\ & \leq \sum_{q=1}^{F_L} \left(L \Delta_{fil,n} \left(\prod_{l'=1}^{L-1} F_{l'} \right) \sum_{q=1}^{F_0} \|\mathbf{x}^q\| \right). \end{aligned} \quad (76)$$

With $F_0 = F_L = 1$ and $F_l = F$ for $1 \leq l \leq L-1$, then we have

$$\|\Phi(\mathbf{H}, \mathbf{L}_n^\epsilon, \mathbf{P}_n f) - \mathbf{P}_n \Phi(\mathbf{H}, \mathcal{L}, f)\| \leq L F^{L-1} \Delta_{fil,n}, \quad (77)$$

which concludes the proof.

REFERENCES

- [1] Z. Wang, L. Ruiz, and A. Ribeiro, "Convolutional filtering on sampled manifolds," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 1–5.
- [2] Z. Wang, L. Ruiz, and A. Ribeiro, "Convergence of graph neural networks on relatively sparse graphs," in *Proc. 57th Asilomar Conf. Signals, Syst., Comput.*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 566–572.
- [3] V. N. Ioannidis, A. G. Marques, and G. B. Giannakis, "Graph neural networks for predicting protein functions," in *Proc. IEEE 8th Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process. (CAMSAP)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 221–225.
- [4] V. Gligorijević et al., "Structure-based protein function prediction using graph convolutional networks," *Nature Commun.*, vol. 12, no. 1, pp. 1–14, 2021.
- [5] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Piscataway, NJ, USA: IEEE Press, 2020, pp. 11785–11792.
- [6] E. Tolstaya, J. Paulos, V. Kumar, and A. Ribeiro, "Multi-robot coverage and exploration using spatial graph neural networks," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 8944–8950.
- [7] J. Zeng, G. Cheung, M. Ng, J. Pang, and C. Yang, "3D point cloud denoising using graph Laplacian regularization of a low dimensional manifold model," *IEEE Trans. Image Process.*, vol. 29, pp. 3474–3489, 2019.
- [8] M. Devanne, H. Wannous, S. Berretti, P. Pala, M. Daoudi, and A. Del Bimbo, "3-D human action recognition by shape analysis of motion trajectories on Riemannian manifold," *IEEE Trans. Cybern.*, vol. 45, no. 7, pp. 1340–1352, Jul. 2015.
- [9] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph CNN for learning on point clouds," *ACM Trans. Graph. (ToG)*, vol. 38, no. 5, pp. 1–12, 2019.
- [10] Z. He, L. Wang, H. Ye, G. Y. Li, and B.-H. F. Juang, "Resource allocation based on graph neural networks in vehicular communications," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Piscataway, NJ, USA: IEEE Press, 2020, pp. 1–5.
- [11] Z. Wang, L. Ruiz, M. Eisen, and A. Ribeiro, "Stable and transferable wireless resource allocation policies via manifold neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 8912–8916.
- [12] F. Gama, E. Isufi, G. Leus, and A. Ribeiro, "Graphs, convolutions, and neural networks: From graph filters to graph neural networks," *IEEE Signal Process. Mag.*, vol. 37, no. 6, pp. 128–138, Nov. 2020.
- [13] A. Ortega, P. Frossard, J. Kovacević, J. M. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proc. IEEE*, vol. 106, no. 5, pp. 808–828, May 2018.
- [14] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE Trans. Neural Netw.*, vol. 20, no. 1, pp. 61–80, Jan. 2009.
- [15] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, "Convolutional neural network architectures for signals supported on graphs," *IEEE Trans. Signal Process.*, vol. 67, no. 4, pp. 1034–1049, Feb. 2019.
- [16] Z. Wang, L. Ruiz, and A. Ribeiro, "Convolutional neural networks on manifolds: From graphs and back," in *Proc. 56th Asilomar Conf. Signals, Syst., Comput.*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 356–360.
- [17] Z. Wang, L. Ruiz, and A. Ribeiro, "Stability of neural networks on manifolds to relative perturbations," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 5473–5477.
- [18] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst, "Geodesic convolutional neural networks on Riemannian manifolds," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, 2015, pp. 37–45.
- [19] R. Chakraborty, J. Bouza, J. H. Manton, and B. C. Vemuri, "Manifold-Net: A deep neural network for manifold-valued data with applications," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 799–810, Feb. 2022.
- [20] F. Gama, J. Bruna, and A. Ribeiro, "Stability properties of graph neural networks," *IEEE Trans. Signal Process.*, vol. 68, pp. 5680–5695, 2020.
- [21] N. Keriven, A. Bietti, and S. Vaiter, "Convergence and stability of graph convolutional networks on large random graphs," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 21512–21523.
- [22] D. Zou and G. Lerman, "Graph convolutional neural networks via scattering," *Appl. Comput. Harmon. Anal.*, vol. 49, no. 3, pp. 1046–1074, 2020.
- [23] Z. Wang, L. Ruiz, and A. Ribeiro, "Stability to deformations of manifold filters and manifold neural networks," *IEEE Trans. Signal Process.*, early access, Mar. 19, 2024.
- [24] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein, "Geometric deep learning on graphs and manifolds using mixture model CNNs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5115–5124.
- [25] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: Going beyond Euclidean data," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 18–42, Jul. 2017.
- [26] M. Eliasof, E. Haber, and E. Treister, "PDE-GCN: Novel architectures for graph neural networks motivated by partial differential equations," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 3836–3849.
- [27] W. Shi and R. Rajkumar, "Point-GNN: Graph neural network for 3D object detection in a point cloud," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 1711–1719.
- [28] M. T. Kejani, F. Dornaika, and H. Talebi, "Graph convolution networks with manifold regularization for semi-supervised learning," *Neural Netw.*, vol. 127, pp. 160–167, July 2020.
- [29] L. Ruiz, L. F. Chamon, and A. Ribeiro, "Transferability properties of graph neural networks," *IEEE Trans. Signal Process.*, vol. 71, pp. 3474–3489, 2023.
- [30] L. Ruiz, L. F. Chamon, and A. Ribeiro, "Graphon signal processing," *IEEE Trans. Signal Process.*, vol. 69, pp. 4961–4976, 2021.
- [31] J. Zhou et al., "Graph neural networks: A review of methods and applications," *AI Open*, vol. 1, pp. 57–81, 2020.
- [32] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. Int. Conf. Learn. Representations*, 2016.
- [33] L. Ruiz, L. Chamon, and A. Ribeiro, "Graphon neural networks and the transferability of graph neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 1702–1712.
- [34] S. Maskey, R. Levie, and G. Kutyniok, "Transferability of graph neural networks: An extended graphon approach," *Appl. Comput. Harmon. Anal.*, vol. 63, pp. 48–83, Mar. 2023.
- [35] J. Calder and N. G. Trillos, "Improved spectral convergence rates for graph Laplacians on ε -graphs and k-NN graphs," *Appl. Comput. Harmon. Anal.*, vol. 60, pp. 123–175, Sep. 2022.
- [36] J. Lee, H. Kim, J. Lee, and S. Yoon, "Transfer learning for deep learning on graph-structured data," in *Proc. AAAI Conf. Artif. Intell.*, vol. 31, no. 1, 2017, doi: 10.1609/aaai.v31i1.10904.
- [37] Q. Zhu, C. Yang, Y. Xu, H. Wang, C. Zhang, and J. Han, "Transfer learning of graph neural networks with ego-graph information maximization," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 1766–1779.
- [38] J. A. Chew, D. Needell, and M. Perlmutter, "A convergence rate for manifold neural networks," in *Proc. Int. Conf. Sampling Theory Appl. (SampTA)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 1–5.
- [39] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE signal Process. Mag.*, vol. 30, no. 3, pp. 83–98, May 2013.
- [40] A. Sandryhaila and J. M. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.
- [41] E. Grong and S. Sommer, "Most probable paths for anisotropic Brownian motions on manifolds," *Found. Comput. Math.*, vol. 24, pp. 313–345, 2024.
- [42] N. Boumal, *An Introduction to Optimization on Smooth Manifolds*. Cambridge, U.K.: Cambridge Univ. Press, 2023.
- [43] S. Rosenberg and R. Steven, *The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds*. Cambridge, U.K.: Cambridge Univ. Press, 1997.
- [44] F. R. Chung, *Spectral Graph Theory*, vol. 92. Providence, RI, USA: AMS, 1997.
- [45] J. Pach, "The beginnings of geometric graph theory," in *Erdos Centennial*, 2013, Berlin, Heidelberg: Springer, pp. 465–484.
- [46] R. Merris, "A survey of graph Laplacians," *Linear Multilinear Algebra*, vol. 39, nos. 1–2, pp. 19–31, 1995.

- [47] R. Levie, W. Huang, L. Bucci, M. Bronstein, and G. Kutyniok, "Transferability of spectral graph convolutional neural networks," *J. Mach. Learn. Res.*, vol. 22, no. 272, pp. 1–59, 2021.
- [48] D. B. Dunson, H.-T. Wu, and N. Wu, "Spectral convergence of graph Laplacian and heat kernel reconstruction in l^∞ from random samples," *Appl. Comput. Harmon. Anal.*, vol. 55, pp. 282–336, Nov. 2021.
- [49] M. Belkin and P. Niyogi, "Towards a theoretical foundation for Laplacian-based manifold methods," *J. Comput. Syst. Sci.*, vol. 74, no. 8, pp. 1289–1308, 2008.
- [50] Y. Xie, J. Ho, and B. C. Vemuri, "Multiple atlas construction from a heterogeneous brain MR image collection," *IEEE Trans. Med. Imag.*, vol. 32, no. 3, pp. 628–635, Mar. 2013.
- [51] Y. Shi and B. Xu, "Gradient estimate of an eigenfunction on a compact Riemannian manifold without boundary," *Ann. Global Anal. Geometry*, vol. 38, no. 1, pp. 21–26, 2010.
- [52] A. Seelmann, "Notes on the $\sin 2\theta$ theorem," *Integral Equ. Operator Theory*, vol. 79, no. 4, pp. 579–597, 2014.
- [53] M. Penrose, *Random Geometric Graphs*, vol. 5. London, U.K.: Oxford Univ. Press, 2003.
- [54] M. Hamidouche, "Spectral analysis of random geometric graphs," Ph.D. dissertation, Université Cote d'Azur, Nice, France, 2020.
- [55] W. Arendt and W. P. Schleich, *Mathematical Analysis of Evolution, Information, and Complexity*. Hoboken, NJ, USA: Wiley, 2009.
- [56] J. Cervino, L. F. Chamon, B. D. Haeffele, R. Vidal, and A. Ribeiro, "Learning globally smooth functions on manifolds," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2023, pp. 3815–3854.
- [57] Y. Jia, D. Zou, H. Wang, and H. Jin, "Enhancing node-level adversarial defenses by Lipschitz regularization of graph neural networks," in *Proc. 29th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, 2023, pp. 951–963.
- [58] Z. Wu et al., "3D ShapeNets: A deep representation for volumetric shapes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1912–1920.
- [59] U. Von Luxburg, M. Belkin, and O. Bousquet, "Consistency of spectral clustering," *Ann. Statist.*, pp. 555–586, 2008.
- [60] M. Belkin and P. Niyogi, "Convergence of Laplacian eigenmaps," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 19, 2006.