

Stability to Deformations of Manifold Filters and Manifold Neural Networks

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

Abstract—The paper defines and studies manifold (M) convolutional filters and neural networks (NNs). *Manifold filters* and MNNs are defined in terms of the Laplace-Beltrami operator exponential and are such that *graph* (G) filters and neural networks (NNs) are recovered as discrete approximations when the manifold is sampled. These filters admit a spectral representation which is a generalization of both the spectral representation of graph filters and the frequency response of standard convolutional filters in continuous time. The main technical contribution of the paper is to analyze the stability of manifold filters and MNNs to smooth deformations of the manifold. This analysis generalizes known stability properties of graph filters and GNNs and it is also a generalization of known stability properties of standard convolutional filters and neural networks in continuous time. The most important observation that follows from this analysis is that manifold filters, same as graph filters and standard continuous time filters, have difficulty discriminating high frequency components in the presence of deformations. This is a challenge that can be ameliorated with the use of manifold, graph, or continuous time neural networks. The most important practical consequence of this analysis is to shed light on the behavior of graph filters and GNNs in large scale graphs.

Index Terms—Graph signal processing, graph neural networks, manifolds, manifold filters, manifold neural networks, manifold deformations, operator stability.

Manuscript received 28 October 2022; revised 1 August 2023; accepted 29 February 2024. Date of publication 19 March 2024; date of current version 30 April 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 29th European Signal Processing Conference (EUSIPCO), 2021 [DOI: 10.23919/EUSIPCO54536.2021.9616024], the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2022 [DOI: 10.1109/ICASSP43922.2022.9747865], and the 56th Asilomar Conference on Signals, Systems, and Computers, 2022 [DOI: 10.1109/IEEECONF56349.2022.10051964]. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Ivan Dokmanić. (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).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TSP.2024.3378379>, provided by the authors.

Digital Object Identifier 10.1109/TSP.2024.3378379

I. INTRODUCTION

GRAPH convolutional filters [4], [5], [6] and graph neural networks (GNNs) [7], [8], [9] have become the tool of choice for signal and information processing on graphs, e.g., [10], [11], [12], [13]. In several applications, graphs can be considered as samples of a manifold. This is sometimes quite explicit as in the case of, e.g., point clouds [14], [15] and sometimes somewhat implicit as in the case of, e.g., wireless communication networks [10], [11]. In this context we can think of convolutions on graphs and GNNs as discretizations of their manifold counterparts. This is important because one can often gain valuable insights about discrete (graph) signal processing by studying continuous (manifold) signal processing.

The main technical contribution of this paper is analyzing the stability of manifold filters and manifold neural networks (MNNs) to smooth deformations of the manifold. Prior to developing stability analyses we define *manifold* convolutional filters and frequency responses that are consistent with *graph* convolutional filters and frequency responses.

Manifold Filters. Consider a manifold with Laplace-Beltrami (LB) operator $\mathcal{L}$ and formulate the corresponding manifold diffusion equation with respect to an auxiliary time variable t . For initial condition $f(x)$, manifold diffusions generate the time varying manifold function $u(x, t) = e^{-t\mathcal{L}} f(x)$ in which $e^{-t\mathcal{L}}$ is the LB operator exponential (Section II-A). A manifold convolutional filter $\tilde{h}(t)$ then acts on the function $f(x)$ as the integral over time of the product of $\tilde{h}(t)$ with the diffusion sequence $u(x, t)$ (Section II-B),

$$g(x) = (\mathbf{h}f)(x) = \int_0^\infty \tilde{h}(t) e^{-t\mathcal{L}} f(x) dt. \quad (1)$$

A manifold convolutional filter is such that we can recover graph filters through discretization of the manifold and discretization of the auxiliary time variable t (Section VI). The reason for this connection is that graph filters are linear combinations of the elements of the graph diffusion sequence [4], [16], [17] and (1) defines manifold filters as linear combinations of the elements of the manifold diffusion sequence. Manifold filters are also generalizations of standard time convolutions. This requires consideration of the wave equation on the real line so that the exponential of the derivative operator $e^{-t\partial/\partial x}$ appears in (1) (Appendix F).

Frequency Response of Manifold Filters We define a filter's frequency response as

$$\hat{h}(\lambda) = \int_0^\infty \tilde{h}(t) e^{-t\lambda} dt. \quad (2)$$

This definition is motivated by the fact that manifold convolutions can be decomposed on separate spectral components (Section II-C). Indeed, let ϕ_i be an eigenfunction of the LB operator associated with eigenvalue λ_i . If $[\hat{f}]_i = \langle f, \phi_i \rangle_{L^2(\mathcal{M})}$ and $[\hat{g}]_i = \langle g, \phi_i \rangle_{L^2(\mathcal{M})}$ are projections of the functions f and g in (1) on this eigenfunction, $[\hat{g}]_i$ depends only on $[\hat{f}]_i$ and can be written as $[\hat{g}]_i = \hat{h}(\lambda_i)[\hat{f}]_i$ (Proposition 1).

This definition of the frequency response of a manifold filter generalizes the frequency response of a standard time filter. This is because the definition in (2) is a Laplace transform, which reduces to the Fourier transform of a filter's impulse response when restricted to the imaginary axis $\lambda = j\omega$. It is also a generalization of the frequency response of a graph filter, which is a z-transform [18].

Stability to Deformations. We establish stability of manifold filters and MNNs with respect to domain deformations formally defined as manifold diffeomorphisms. We consider a signal f defined on a manifold $\mathcal{M}$ and the signal $f \circ \tau$ made up of the composition of f with a diffeomorphism. The signals f and $f \circ \tau$ are passed through the same MNN Φ . In this paper we prove that if τ is ϵ -small and ϵ -smooth the respective MNN outputs satisfy

$$\|\Phi(f \circ \tau) - \Phi(f)\|_{L^2(\mathcal{M})} = O(\epsilon)\|f\|_{L^2(\mathcal{M})}, \quad (3)$$

provided that the manifold filters in the layers of the MNN satisfy certain spectral properties (Sections III and IV).

The bound in (3) is a generalization of the standard convolutional neural network (CNN) bound in [19], which studies diffeomorphisms of the real line and its effect on convolutional filter banks and CNNs. The bound is also a limit version of the GNN bounds of [20] and related literature [21], [22], [23], [24], [25]. As is the case of [19], [20], [21], [22], [23], [24], [25] the bound in (3) holds when the frequency response $\hat{h}(\lambda)$ of the filters that make up the layers of the MNN have decreasing variability with increasing λ (Section III). Thus, stability requires *layers* with limited ability to discriminate high frequency components and implies that one may expect *multi-layered* MNNs to outperform filters in learning tasks in which high frequency components are important (Section V).

Related Work and Significance. We focus on the stability analysis of MNNs – the limit version of GNNs – because all of the existing GNN stability results have bounds that grow with the number of nodes in the graph [21], [24]. To overcome this limitation, many works have studied neural networks on graphons [25], [26], [27] and more general graph models with variable sparsity [28]. Results in these settings are independent of graph size – same as the results for a limit object presented here. Manifolds can provide limit models for relatively sparse graphs, i.e. ϵ -graphs and k -NN graphs [29]. Even if there exist other random graph models allowing to model moderately sparse graphs (e.g., [28]), these models do not have the physical

interpretation of a manifold, which is often a better descriptor of real-world domains. Of particular relevance to our paper is the work on GNN transferability for graphs that are sampled from a general topological space [30].

Organization. Section II introduces preliminary definitions (Section II-A), defines manifold convolutions (Section II-B), and introduces the spectral domain representation of manifold filters (Section II-C). Section III studies the stability of manifold filters to manifold deformations. It shows that a diffeomorphism results in a perturbation of the LB operator that involves additive and multiplicative terms (Theorem 1). It then goes on to study the effect on manifold filters of additive (Section III-A) and multiplicative (Section III-B) perturbations of the LB operator. Section IV extends the analysis of manifold filter stability to manifold neural networks and Section V discusses the implications of the results derived in Sections III and IV. Section VI explains how to recover graph filters and GNNs from the discretization of a manifold. Section VII illustrates the results of Sections III and IV with numerical examples. Section VIII concludes the paper. Proofs are deferred to appendices and supplementary material.

II. MANIFOLD CONVOLUTIONAL FILTERS

Consider a compact, smooth and differentiable d -dimensional submanifold $\mathcal{M} \subset \mathbb{R}^N$. For simplicity, in this paper we use the words submanifold and manifold interchangeably, assuming the manifold $\mathcal{M}$ to always be embedded in $\mathbb{R}^N$. This embedding induces a Riemannian structure [31]. In turn, the Riemannian structure allows defining a measure μ over the manifold as well as a notion of *length* for smooth curves on $\mathcal{M}$. Given two points $x, y \in \mathcal{M}$, the length of the shortest curve between x and y is denoted $\text{dist}(x, y)$ and called the geodesic distance between these points.

We consider the manifold $\mathcal{M}$ to be the support of data that we represent as smooth real scalar functions $f : \mathcal{M} \rightarrow \mathbb{R}$. We call these scalar functions *manifold signals*. We focus on manifold signals that have finite energy, such that $f \in L^2(\mathcal{M})$. Since $L^2(\mathcal{M})$ is a Hilbert space, it is equipped with an inner product given by

$$\langle f, g \rangle_{L^2(\mathcal{M})} = \int_{\mathcal{M}} f(x)g(x)d\mu(x) \quad (4)$$

where $d\mu(x)$ is the d -dimensional volume element corresponding to measure μ . Thus, the energy of the signal f is given by $\|f\|_{L^2(\mathcal{M})}^2 = \langle f, f \rangle_{L^2(\mathcal{M})}$.

A. Laplace-Beltrami Operator

On manifolds, differentiation is implemented by the *intrinsic gradient* – a local operator acting on a neighborhood of each point on the manifold that is homeomorphic to a d -dimensional Euclidean space. This neighborhood contains all the vectors tangent to $\mathcal{M}$ at x , which is called the tangent space of $x \in \mathcal{M}$ and is denoted as $T_x\mathcal{M}$. The disjoint union of all tangent spaces on $\mathcal{M}$ is defined as the tangent bundle $T\mathcal{M}$. Formally, the intrinsic gradient is the operator $\nabla : L^2(\mathcal{M}) \rightarrow L^2(T\mathcal{M})$

mapping scalar functions $f \in L^2(\mathcal{M})$ to *tangent vector functions* $\nabla f \in L^2(T\mathcal{M})$, where $L^2(T\mathcal{M})$ is the Hilbert space of vector fields over the manifold $\mathcal{M}$. The tangent vector function $\nabla f(x) \in T_x\mathcal{M}$ indicates the direction of the fastest change of manifold signal f at point x . The adjoint of the intrinsic gradient is the *intrinsic divergence*, defined as $\text{div} : L^2(T\mathcal{M}) \rightarrow L^2(\mathcal{M})$. Interpreting the tangent vector field as the velocity field of a fluid, the intrinsic divergence can be seen as a measure of the net motion of the fluid [14].

The Laplace-Beltrami (LB) operator of a manifold $\mathcal{M}$ is defined as the operator $\mathcal{L} : L^2(\mathcal{M}) \rightarrow L^2(\mathcal{M})$ given by the function composition of the intrinsic divergence and the intrinsic gradient. When considered in the local coordinates supported on $T_x\mathcal{M}$ [32], the LB operator can be written as

$$\mathcal{L}f = -\text{div} \circ \nabla f = -\nabla \cdot \nabla f. \quad (5)$$

Much like the Laplace operator in Euclidean domains (or the Laplacian matrix, in the case of graphs [33]), the LB operator measures the total variation of function f , i.e., how much the value of f at a point deviates from local average of the values of f in its surroundings [14]. Since the LB operator $\mathcal{L}$, like the gradient ∇ , is a local operator depending on the tangent space $T_x\mathcal{M}$ of each point $x \in \mathcal{M}$, in the following we omit this dependence for the ease of presentation by writing $\mathcal{L} = \mathcal{L}_x$ and $\nabla = \nabla_x$.

The LB operator plays an important role in partial differential equations (PDEs), as it governs the dynamics of the diffusion of heat over manifolds as given by the *heat equation*

$$\frac{\partial u(x, t)}{\partial t} + \mathcal{L}u(x, t) = 0. \quad (6)$$

If $u \in L^2(\mathcal{M})$ measures the temperature over the manifold with $u(x, t)$ representing the temperature of point $x \in \mathcal{M}$ at time $t \in \mathbb{R}^+$, equation (6) can be interpreted to mean that, at point x on manifold $\mathcal{M}$, the rate at which the manifold “cools down” is proportional to the difference between the temperature of x and the local average of the temperature of the points in its neighborhood. With initial condition $u(x, 0) = f(x)$ for all $x \in \mathcal{M}$, the solution to this equation is given by

$$u(x, t) = e^{-t\mathcal{L}}f(x), \quad (7)$$

which is the key support to implement the LB operator in the definitions proposed later.

The LB operator $\mathcal{L}$ is self-adjoint and positive-semidefinite. Considering that $\mathcal{M}$ is compact, the LB operator $\mathcal{L}$ has a real positive eigenvalue spectrum $\{\lambda_i\}_{i=1}^\infty$ satisfying

$$\mathcal{L}\phi_i = \lambda_i\phi_i \quad (8)$$

where ϕ_i is the eigenfunction associated with eigenvalue λ_i . The indices i are such that the eigenvalues are ordered in increasing order as $0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$ repeated according to their multiplicity. In particular, for a d -dimensional manifold, we have that $\lambda_i \propto i^{2/d}$ as a consequence of Weyl’s law [34], [35]. The eigenfunctions ϕ_i are orthonormal and form a generalized eigenbasis of $L^2(\mathcal{M})$ in the intrinsic sense. Since $\mathcal{L}$ is a total variation operator, the eigenvalues λ_i can be interpreted as the canonical frequencies and the eigenfunctions ϕ_i as the

canonical oscillation modes of $\mathcal{M}$. This further allows us to implement operator $\mathcal{L}$ in the spectral domain.

B. Manifold Filters

Time signals are processed by filters which compute the continuous time convolution of the input signal and the filter impulse response [18]; images and high-dimensional Euclidean signals are processed by filters implementing multidimensional convolutions [36]; and graph signals are filtered by computing graph convolutions [6]. In this paper, we define a manifold filter as the convolution of the filter impulse response $\tilde{h}(t)$ and the manifold signal f . Note that the definition of the convolution operation, denoted as $\star_{\mathcal{M}}$, leverages the heat diffusion dynamics described in (7).

Definition 1 (Manifold filter): Let $\tilde{h} : \mathbb{R}^+ \rightarrow \mathbb{R}$ and let $f \in L^2(\mathcal{M})$ be a manifold signal. The manifold filter with impulse response $\tilde{h}$, denoted as $\mathbf{h}$, is given by

$$g(x) = (\mathbf{h}f)(x) := (\tilde{h} \star_{\mathcal{M}} f)(x) := \int_0^\infty \tilde{h}(t)u(x, t)dt, \quad (9)$$

where $\tilde{h} \star_{\mathcal{M}} f$ is the *manifold convolution* of $\tilde{h}$ and f , and $u(x, t)$ is the solution of the heat equation (6) with the initial condition $u(x, 0) = f(x)$.

In a slight abuse of nomenclature, in the following we will use the terms manifold filter and manifold convolution interchangeably.

From Definition 1, we see that the manifold filter operates on manifold signals f by (i) scaling the diffusion process (7) starting at f by $\tilde{h}$ and (ii) aggregating the outcome of the scaled diffusion process from $t = 0$ to $t = \infty$.

Substituting (6) into the convolution definition in (9) yields

$$g(x) = (\mathbf{h}f)(x) = \int_0^\infty \tilde{h}(t)e^{-t\mathcal{L}}f(x)dt = \mathbf{h}(\mathcal{L})f(x). \quad (10)$$

This alternative form uncovers the fact that the manifold convolution is a map from functions f to functions g that is *parametric* on the Laplacian $\mathcal{L}$ – i.e., manifold convolutions are completely determined by the LB operator of the manifold. The exponential term $e^{-t\mathcal{L}}$ can be seen as a diffusion or shift operation akin to a time delay in a linear time-invariant (LTI) filter [18], or as the graph shift operator in a linear shift-invariant (LSI) graph filter [6]. Indeed, if we consider the manifold $\mathcal{M}$ to be the real line, the manifold filter defined in (10) recovers a LTI filter. If we consider it to be a set of points connected by a geometric graph, (10) recovers an LSI graph filter. We discuss these special cases in detail in Appendix F and Section VI.

C. Frequency Representation of Manifold Filters

A manifold signal $f \in L^2(\mathcal{M})$ can be represented in the frequency domain of the manifold by projecting f onto the LB operator eigenbasis (8) as

$$[\hat{f}]_i = \langle f, \phi_i \rangle_{L^2(\mathcal{M})} = \int_{\mathcal{M}} f(x)\phi_i(x)d\mu(x), \quad (11)$$

where we claim that $\hat{f}$ is the *frequency representation* of the corresponding signal with $f = \sum_{i=1}^\infty [\hat{f}]_i \phi_i$.

Frequency representations are useful because they help understand the frequency behavior of the manifold filter $\mathbf{h}(\mathcal{L})$. To see this, we consider the frequency representation of the manifold filter output g in (10), which is

$$[\hat{g}]_i = \int_{\mathcal{M}} \int_0^\infty \tilde{h}(t) e^{-t\mathcal{L}} f(x) dt \phi_i(x) d\mu(x). \quad (12)$$

Rearranging the integrals and substituting $e^{-t\mathcal{L}} \phi_i = e^{-t\lambda_i} \phi_i$, we can get

$$[\hat{g}]_i = \int_0^\infty \tilde{h}(t) e^{-t\lambda_i} dt [\hat{f}]_i. \quad (13)$$

The expression relating $\hat{g}$ and $\hat{f}$ is called the *frequency response* of the filter $\mathbf{h}(\mathcal{L})$.

Definition 2 (Frequency response): The frequency response of manifold filter $\mathbf{h}(\mathcal{L})$ is given by

$$\hat{h}(\lambda) = \int_0^\infty \tilde{h}(t) e^{-t\lambda} dt, \lambda \in (0, \infty). \quad (14)$$

An important consequence of Definition 2 is that, since $\hat{h}(\lambda)$ is parametric on λ , the manifold filter is pointwise in the frequency domain. This can be seen by plugging (14) into (13), and is stated explicitly in Proposition 1.

Proposition 1: Manifold filter $\mathbf{h}(\mathcal{L})$ is pointwise in the frequency domain, which is written as

$$[\hat{g}]_i = \hat{h}(\lambda_i) [\hat{f}]_i, i \in \mathbb{N}^+ \quad (15)$$

Definition 2 also emphasizes that the frequency response of a manifold filter is independent of the underlying manifold. Note that, in (14), $\hat{h}(\lambda)$ is a function of an arbitrary scalar variable λ . To obtain the frequency behavior of this filter on a given manifold $\mathcal{M}$, we need to evaluate $\hat{h}$ at the corresponding LB operator eigenvalues λ_i [cf. (8)]. If the manifold changes (or if we want to deploy the same filter on a different manifold $\mathcal{M}'$), it suffices to reevaluate $\hat{h}$ at λ'_i , i.e., at the eigenvalues of the new LB operator $\mathcal{L}'$.

III. STABILITY OF MANIFOLD FILTERS WITH RESPECT TO MANIFOLD DEFORMATIONS

On the manifold $\mathcal{M}$, we define a deformation as function $\tau: \mathcal{M} \rightarrow \mathcal{M}$ and the curvature distance between x and the displaced $\tau(x)$ $\text{dist}(x, \tau(x))$ is upper bounded, i.e., $\tau(x)$ is a displaced point in the neighborhood of x , which holds for all $x \in \mathcal{M}$. The deformation τ induces a corresponding tangent map $\tau_{*,x}: T_x \mathcal{M} \rightarrow T_{\tau(x)} \mathcal{M}$ which is a linear map between the tangent spaces [37]. With the coordinate description of tangent map, the tangent map $\tau_{*,x}$ can be exactly represented by the Jacobian matrix $J_x(\tau)$. When $\text{dist}(x, \tau(x))$ is bounded, the Frobenius norm of $J_x(\tau) - I$ can also be upper bounded, and these bounds are used to measure the size of the deformation τ .

Let $f: \mathcal{M} \rightarrow \mathbb{R}$ be a manifold signal. Because $\mathcal{M}$ is the codomain of τ , $g = f \circ \tau$ maps points $\tau(x) \in \mathcal{M}$ to $f(\tau(x)) \in \mathbb{R}$, so that the effect of a manifold deformation on the signal f is a signal perturbation leading to a new signal g supported on the same manifold. To understand the effect of this deformation

on the LB operator, let $p = \mathcal{L}g$. Since p is also a signal on $\mathcal{M}$, we may define an operator $\mathcal{L}'$ mapping f directly into p ,

$$p(x) = \mathcal{L}' f(x) = \mathcal{L}g(x) = \mathcal{L}f(\tau(x)). \quad (16)$$

The operator $\mathcal{L}'$ is the perturbed LB operator, which is effectively the new LB operator resulting from the deformation τ . Assuming that the gradient field is smooth, the difference between $\mathcal{L}'$ and $\mathcal{L}$ is given by the following theorem. The proof is deferred to Appendix A.

Theorem 1: Let $\mathcal{L}$ be the LB operator of manifold $\mathcal{M}$. Let $\tau: \mathcal{M} \rightarrow \mathcal{M}$ be a manifold perturbation such that $\text{dist}(x, \tau(x)) \leq \epsilon$ and $J_x(\tau) = I + \Delta_x$ with $\|\Delta_x\|_F \leq \epsilon$ for all $x \in \mathcal{M}$. If the gradient field is smooth, it holds that

$$\mathcal{L} - \mathcal{L}' = \mathbf{E}\mathcal{L} + \mathcal{A}, \quad (17)$$

where $\mathbf{E}$ and $\mathcal{A}$ satisfy $\|\mathbf{E}\| = O(\epsilon)$ and $\|\mathcal{A}\|_{op} = O(\epsilon)$.

Therefore, the perturbation of the LB operator incurred by a manifold deformation τ is a combination of an absolute perturbation $\mathcal{A}$ [cf. Definition 3] and a relative perturbation $\mathbf{E}\mathcal{L}$ [cf. Definition 7]. This largely simplifies our analysis of stability. Since manifold filters are parametric on $\mathcal{L}$ [cf. Proposition 1], it is sufficient to characterize their stability to deformations of the manifold by analyzing their behavior in the presence of absolute and relative LB perturbations. This is what we do in Sections III-A and III-B.

A. Stability of Manifold Filters to Absolute Perturbations

We start by analyzing the stability of manifold filters to absolute perturbations of the LB operator, which are introduced in Definition 3.

Definition 3 (Absolute perturbations): Let $\mathcal{L}$ be the LB operator of manifold $\mathcal{M}$. An absolute perturbation of $\mathcal{L}$ is defined as

$$\mathcal{L}' - \mathcal{L} = \mathcal{A}, \quad (18)$$

where the absolute perturbation operator $\mathcal{A}$ is self-adjoint.

Like $\mathcal{L}$, the operator $\mathcal{L}'$ resulting from the absolute perturbation of $\mathcal{L}$ is self-adjoint due to the symmetry of $\mathcal{A}$. Hence, it admits an eigendecomposition similar to (5). When the filter coefficients are fixed, the frequency response of the manifold filter (15) can be obtained by evaluating $\hat{h}(\lambda)$ at each λ_i . Thus, to understand the effect of the perturbation on the filter we need to look at how the perturbation of the LB operator changes the eigenvalues λ_i . The challenge in this case is that the spectrum of $\mathcal{L}$ is infinite-dimensional, i.e., there is an infinite (though countable) number of eigenvalue perturbations that need to be taken into account, which leads to an untractable infinite summation. As demonstrated by Proposition 2, however, large eigenvalues of $\mathcal{L}$ tend to accumulate in certain parts of the real line. This suggests a strategy to partition the spectrum into finite number of partitions, by treating the eigenvalues in the same partition similarly, we can turn to analyze the stability under the perturbations on these finite number of partitions with the discriminability of certain frequency components sacrificed.

Proposition 2: Consider a d -dimensional compact manifold $\mathcal{M} \subset \mathbb{R}^N$ and let $\mathcal{L}$ be its LB operator with eigenvalues

$\{\lambda_k\}_{k=1}^\infty$. Let C_0 be an arbitrary constant, $K_0(C_0)$ some finite constant depends on C_0 and C_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 ,

$$N_1 = \left\lceil \left(\frac{\alpha d}{C_0 4\pi^2} \right)^{d/(2-d)} (C_d \text{Vol}(\mathcal{M}))^{2/(2-d)} \right\rceil \quad (19)$$

such that, for all $k > \max\{N_1, K_0(C_0)\}$,

$$\lambda_{k+1} - \lambda_k \leq \alpha$$

Proof: This is a direct consequence of Weyl's law [38, Chapter 1], [35], See Section A in supplementary material. $\square$

Given this asymptotic behavior, we can divide the eigenvalues into a finite number of partitions by placing eigenvalues that are less than $\alpha > 0$ apart from each other in groups. This spectrum separation strategy is described in Definition 4. To achieve it, we will need a specific type of manifold filter called Frequency Difference Threshold (FDT) filter as introduced in Definition 5.

Definition 4 (α -separated spectrum): An α -separated spectrum of a LB operator $\mathcal{L}$ is defined as a partition $\Lambda_1(\alpha) \cup \dots \cup \Lambda_N(\alpha)$ such that all $\lambda_i \in \Lambda_k(\alpha)$ and $\lambda_j \in \Lambda_l(\alpha)$, $k \neq l$, satisfy

$$|\lambda_i - \lambda_j| > \alpha. \quad (20)$$

Definition 5 (α -FDT filter): An α -frequency difference threshold (α -FDT) filter is defined as a filter $\mathbf{h}(\mathcal{L})$ whose frequency response satisfies

$$|\hat{h}(\lambda_i) - \hat{h}(\lambda_j)| \leq \delta_k \text{ for all } \lambda_i, \lambda_j \in \Lambda_k(\alpha) \quad (21)$$

with $\delta_k \leq \delta$ for $k = 1, \dots, N$.

In an α -separated spectrum, eigenvalues $\lambda_i \in \Lambda_k(\alpha)$ and $\lambda_j \in \Lambda_l(\alpha)$ in different sets ($k \neq l$) are at least α away from each other. Conversely, eigenvalues $\lambda_i, \lambda_j \in \Lambda_k(\alpha)$ are no more than α apart. This partitioning creates several eigenvalue groups spaced by at least α . Note that the sets $\Lambda_k(\alpha)$ can have any size and, in particular, they can be singletons.

The partitioning of the spectrum described in Definition 4 is achieved by an α -FDT filter. This filter separates the spectrum of the manifold by assigning similar frequency responses—that deviate no more than δ_k from each other—to eigenvalues $\lambda_i \in \Lambda_k(\alpha)$, $1 \leq k \leq N$. In other words, an α -FDT filter does not discriminate between eigenvalues $\lambda_i, \lambda_j \in \Lambda_k(\alpha)$. Importantly, the δ_k in Definition 4 are finite, so that they can be bounded by some δ .

To obtain manifold filters that are stable to absolute perturbations of $\mathcal{L}$, we also need these filters to be Lipschitz continuous as shown in Definition 6.

Definition 6 (Lipschitz filter): A filter is A_h -Lipschitz if its frequency response is Lipschitz continuous with Lipschitz constant A_h , i.e.,

$$|\hat{h}(a) - \hat{h}(b)| \leq A_h |a - b| \text{ for all } a, b \in (0, \infty). \quad (22)$$

Between the eigenvalue groups, the filters that we consider are assumed to be A_h -Lipschitz continuous. This means that, in regions of the spectrum where the $\Lambda_k(\alpha)$ are singletons, the filter can vary with slope at most A_h as shown in Fig. 1. Note




Fig. 1. Illustration of an α -FDT filter. The x -axis stands for the spectrum with each sample representing an eigenvalue. The gray shaded areas show the grouping of the eigenvalues according to Definition 4. The red lines show a set of α -FDT filters that can discriminate each eigenvalue group.

that we can always construct convolutional filters (14) that are both Lipschitz continuous and α -FDT.

Under mild assumptions on the amplitude of the frequency response $\hat{h}$ (Assumption 1), it can be shown that Lipschitz continuous α -FDT filters are stable to absolute perturbations of the LB operator. This result is stated in Theorem 2.

Assumption 1 (Non-amplifying filters): The filter frequency response $\hat{h} : \mathbb{R}^+ \rightarrow \mathbb{R}$ is non-amplifying. I.e., for all $\lambda \in (0, \infty)$, $\hat{h}$ satisfies $|\hat{h}(\lambda)| \leq 1$.

Note that this assumption is rather reasonable, because the filter frequency response $\hat{h}(\lambda)$ can always be normalized.

Theorem 2 (Manifold filter stability to absolute perturbations): Consider a manifold $\mathcal{M}$ with LB operator $\mathcal{L}$. Let $\mathbf{h}(\mathcal{L})$ be an α -FDT manifold filter [cf. Definition 5] and A_h -Lipschitz [cf. Definition 6]. Consider an absolute perturbation $\mathcal{L}' = \mathcal{L} + \mathcal{A}$ of the LB operator $\mathcal{L}$ [cf. Definition 3] where $\|\mathcal{A}\| \leq \epsilon < \alpha$. Then, under Assumption 1 it holds that

$$\|\mathbf{h}(\mathcal{L})f - \mathbf{h}(\mathcal{L}')f\|_{L^2(\mathcal{M})} \leq \left(\frac{\pi N_s \epsilon}{\alpha - \epsilon} + A_h \epsilon + 2(N - N_s)\delta \right) \|f\|_{L^2(\mathcal{M})}, \quad (23)$$

where N is the size of the α -separated spectrum partition [cf. Definition 4] and N_s is the number of singletons.

Proof: See Appendix C. $\square$

Provided that $\epsilon \ll \alpha$, FDT filters are thus stable to absolute perturbations of the LB operator $\mathcal{L}$. The stability bound depends on (i) the continuity of the FDT filter as measured by the Lipschitz constant A_h and (ii) its frequency difference threshold α , which affects the bound directly as well as indirectly through the number of partitions N . Note that this bound consists of three terms. The first corresponds to the difference between the eigenfunctions of $\mathcal{L}$ and $\mathcal{L}'$, which affects the stability bound by changing projection directions. The second stems from the distance between the original and perturbed eigenvalues. Finally, the third reflects the bounded fluctuation of the filter frequency response within the same eigenvalue group.

The bound in Theorem 2 can be simplified by setting $\delta = \pi\epsilon/(2\alpha - 2\epsilon)$ as in Corollary 1.

Corollary 1: Setting $\delta = \pi\epsilon/(2\alpha - 2\epsilon)$, under the same assumptions of Theorem 2 it holds that

$$\|\mathbf{h}(\mathcal{L})f - \mathbf{h}(\mathcal{L}')f\|_{L^2(\mathcal{M})} \leq \left(\frac{\pi N}{\alpha - \epsilon} + A_h \right) \epsilon \|f\|_{L^2(\mathcal{M})}. \quad (24)$$




Fig. 2. Illustration of a γ -FRT filter. The x -axis stands for the spectrum with each sample representing an eigenvalue. The gray shaded area shows the grouping of the eigenvalues according to Definition 8. The red lines show a set of α -FDT filters that can discriminate each eigenvalue group.

A particular case of Theorem 2, the simplified stability bound in Corollary 1 is helpful to understand the effect of the filter spectrum on stability as well as of the size of the perturbation. In particular, from Corollary 1 we can tell that the filter is more stable if the Lipschitz constant A_h is small and the frequency difference threshold α is large. On the other hand, small A_h and large α mean that the filter is less discriminative. This reveals a stability-discriminability trade-off where discriminability should be understood as the ability to tell frequencies apart. In other words, we propose α -FDT filter to maintain the stability by trying to discriminate only eigenvalue groups instead of every single eigenvalue. Importantly, this trade-off is not related to the magnitude of the frequencies that the filters amplify (as is the case in, e.g., [20], [39]). We will keep observing this trade-off throughout the stability analysis of both manifold filters and MNNs. More details about this characteristic will be discussed in Section V.

B. Stability of Manifold Filters to Relative Perturbations

Relative perturbations of the LB operator are defined similarly as follows.

Definition 7 (Relative perturbations): Let $\mathcal{L}$ be the LB operator of manifold $\mathcal{M}$. A relative perturbation of $\mathcal{L}$ is defined as

$$\mathcal{L}' - \mathcal{L} = \mathbf{E}\mathcal{L}, \quad (25)$$

where the relative perturbation term $\mathbf{E}\mathcal{L}$ is self-adjoint.

Like absolute perturbations, relative perturbations also perturb the eigenvalues and eigenfunctions of $\mathcal{L}$. However, in the case of relative perturbations, the perturbations to the eigenvalues are proportional to their absolute values [cf. Lemma 3 in Appendix B]. Relative perturbations thus require a different spectrum separation strategy to guarantee stability. With relative perturbations, larger eigenvalues are impacted with larger perturbation values, which can send eigenvalues originally α -close to each other to different groups as Fig. 2 shows. Therefore, we will need a different type of filter implementing a different type of spectrum separation. Our strategy is inspired by Proposition 3, which is another variation of Weyl's law.

Proposition 3: Let $\mathcal{M}$ be a d -dimensional compact embedded manifold in $\mathbb{R}^N$ with LB operator $\mathcal{L}$, and let $\{\lambda_k\}_{k=1}^\infty$ denote the eigenvalues of $\mathcal{L}$. Let C_0 denote an arbitrary constant and

$K_0(C_0)$ some finite constant depends on C_0 . For any $\gamma > 0$, there exists N_2 given by

$$N_2 = \lceil (((\gamma + 1)/C_0)^{d/2} - 1)^{-1} \rceil \quad (26)$$

such that, for all $k > \max\{N_2, K_0(C_0)\}$, it holds that

$$\lambda_{k+1} - \lambda_k \leq \gamma \lambda_k.$$

Proof: This is a direct consequence of Weyl's law [38, Chapter 1], [35]. See Section A in supplementary material. $\square$

Hence, to enforce stability we need to separate the spectrum relatively to the ratio between neighboring eigenvalues. This partitioning is called γ -separated spectrum and formalized in Definition 8. A γ -separated spectrum is achieved by a so-called Frequency Ratio Threshold (FRT) filters. We introduce them in Definition 9.

Definition 8 (γ -separated spectrum): A γ -separated spectrum of a LB operator $\mathcal{L}$ is defined as a partition $\Lambda_1(\gamma) \cup \dots \cup \Lambda_M(\gamma)$ such that all $\lambda_i \in \Lambda_k(\gamma)$ and $\lambda_j \in \Lambda_l(\gamma)$, $k \neq l$, satisfy

$$\left| \frac{\lambda_i}{\lambda_j} - 1 \right| > \gamma. \quad (27)$$

Definition 9 (γ -FRT filter): A γ -frequency ratio threshold (γ -FRT) filter is defined as a filter $\hat{h}(\mathcal{L})$ whose frequency response satisfies

$$|\hat{h}(\lambda_i) - \hat{h}(\lambda_j)| \leq \delta_k, \text{ for all } \lambda_i, \lambda_j \in \Lambda_k(\gamma) \quad (28)$$

with $\delta_k \leq \delta$ for $k = 1, 2, \dots, M$.

In a γ -separated spectrum, the sets $\Lambda_k(\gamma)$ are built based on eigenvalue distances *relative* to the eigenvalues' magnitudes and weighted by the parameter γ . Eigenvalues $\lambda_j \in \Lambda_k(\gamma)$ and $\lambda_i \in \Lambda_l(\gamma)$ in different groups (i.e., $k \neq l$) are at least $\gamma \min(\lambda_i, \lambda_j)$ apart from each other. This means that, for $\lambda_i, \lambda_{i+1} \in \Lambda_k(\gamma)$, $\lambda_{i+1} - \lambda_i \leq \gamma \lambda_i$.

A γ -FRT filter achieves a spectrum separation in Definition 8 by giving eigenvalues $\lambda_i, \lambda_j \in \Lambda_k(\gamma)$ very similar frequency responses differing by at most plus or minus $\delta_k \leq \delta$. Meanwhile, eigenvalues belonging to different sets $\Lambda_k(\gamma)$ and $\Lambda_l(\gamma)$, $k \neq l$, are treated independently, and their frequency responses can vary a lot.

To make a manifold filter stable to relative perturbations of the LB operator, we need a further restriction on their continuity. Lipschitz continuity [cf. Definition 6] is not enough because for a Lipschitz filter the difference in frequency response for a perturbed eigenvalue grows with the eigenvalue magnitude, since the eigenvalue perturbation is relative. Therefore, we need our filters to be *integral Lipschitz* continuous as is described in Definition 10.

Definition 10 (Integral Lipschitz filter): A filter is integral Lipschitz continuous with constant B_h if its frequency response satisfies

$$|\hat{h}(a) - \hat{h}(b)| \leq \frac{B_h |a - b|}{(a + b)/2} \text{ for all } a, b \in (0, \infty). \quad (29)$$

Integral Lipschitz filters can be seen as Lipschitz filters with variable Lipschitz constant, which decreases with λ . E.g., on the interval (a, b) , the filter in Definition 10 behaves as a Lipschitz filter with Lipschitz constant $2B_h/(a + b)$. When a and b are

close, this condition can be approximated by $|ah'(a)| \leq B_h$ for all $a \in (0, \infty)$. This implies that the filter function flattens for high-frequency eigenvalues as shown in Fig. 2.

Under Assumption 1, integral Lipschitz γ -FRT filters are stable to relative perturbations as stated in Theorem 3.

Theorem 3 (Manifold filter stability to relative perturbations): Consider a manifold $\mathcal{M}$ with LB operator $\mathcal{L}$. Let $\mathbf{h}(\mathcal{L})$ be a γ -FRT filter with $\delta = \pi\epsilon/(2\gamma - 2\epsilon + 2\gamma\epsilon)$ [cf. Definition 9] and B_h -integral Lipschitz [cf. Definition 10]. Consider a relative perturbation $\mathcal{L}' = \mathcal{L} + \epsilon\mathcal{L}$ of the LB operator $\mathcal{L}$ [cf. Definition 7] where $\|\mathbf{E}\| \leq \epsilon < \gamma$. Then, under Assumption 1 it holds that

$$\|\mathbf{h}(\mathcal{L})f - \mathbf{h}(\mathcal{L}')f\|_{L^2(\mathcal{M})} \leq \left(\frac{\pi M \epsilon}{\gamma - \epsilon + \gamma\epsilon} + \frac{2B_h \epsilon}{2 - \epsilon} \right) \|f\|_{L^2(\mathcal{M})} \quad (30)$$

where M is the size of the γ -separated spectrum partition [cf. Definition 8].

Proof: See Appendix E. $\square$

When ϵ is sufficiently small ($\epsilon \ll \min(\gamma, 2)$, which is typically the case with deformations such as the one in Theorem 1), the denominators on the right hand side of (30) are approximately equal to γ and 2 respectively. Hence, γ -FRT integral Lipschitz filters are stable to relative perturbations of the LB operator. Besides appearing in the bound in Theorem 3), the frequency ratio threshold γ also affects stability indirectly through the partition size M . With a larger γ , fewer eigenvalues will be in singleton sets, thus decreasing M and improving stability. A smaller integral Lipschitz constant B_h also increases stability. However, small B_h and large γ make for smoother filters which in turn lead to a less discriminative manifold filter. Therefore, integral Lipschitz γ -FRT filters also exhibit a trade-off between discriminability and stability.

Remark 1: By comparing the illustrations of α -FDT filter (Definition 5) and γ -FRT filter (Definition 9) in Fig. 1 and Fig. 2, we see that in practice these filters have a similar frequency behavior because, due to Weyl's law [cf. Proposition 2 and 3], high frequency components will eventually be grouped in the same group and thus share similar frequency responses. Therefore, the main difference between these filters is their effects on the low-frequency components. In the low frequency spectrum, the eigengaps $\lambda_{i+1} - \lambda_i$ may be smaller than the difference threshold α , but larger than the relative ratio threshold $\gamma\lambda_i$ due to λ_i being small. However, for appropriate values of γ a filter may be both FDT and FRT. This will be shown in Section IV.

IV. STABILITY OF MANIFOLD NEURAL NETWORKS

Manifold neural networks (MNNs) are deep convolutional architectures comprised of L layers, where each layer consists of two components: a convolutional filter bank and a pointwise nonlinearity. At each layer $l = 1, 2, \dots, L$, the convolutional filters map the incoming F_{l-1} features from layer $l-1$ into F_l intermediate linear features given by

$$y_l^p(x) = \sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q(x), \quad (31)$$

where $\mathbf{h}_l^{pq}(\mathcal{L})$ is the filter mapping the q -th feature from layer $l-1$ to the p -th feature of layer l as in (15), for $1 \leq q \leq F_{l-1}$ and $1 \leq p \leq F_l$. The intermediate features are then processed by a pointwise nonlinearity $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ as

$$f_l^p(x) = \sigma(y_l^p(x)). \quad (32)$$

The nonlinearity σ processes each feature individually and we further make an assumption on its continuity as follows.

Assumption 2 (Normalized Lipschitz activation functions): The activation function σ is normalized Lipschitz continuous, i.e., $|\sigma(a) - \sigma(b)| \leq |a - b|$, with $\sigma(0) = 0$.

Note that this assumption is rather reasonable, since most common activation functions (e.g., the ReLU, the modulus and the sigmoid) are normalized Lipschitz by design.

At the first layer of the MNN, the input features are the input data f^q for $1 \leq q \leq F_0$. At the output of the MNN, the output features are given by the outputs of the L -th layer, i.e., f_L^p for $1 \leq p \leq F_L$. To represent the MNN more succinctly, we may gather the impulse responses of the manifold convolutional filters $\mathbf{h}_l^{pq}$ across all layers in a function set $\mathbf{H}$, and define the MNN map $\Phi(\mathbf{H}, \mathcal{L}, f)$. This map emphasizes that the MNN is parameterized by both the filter functions and the LB operator $\mathcal{L}$. We next will analyze the stability of $\Phi(\mathbf{H}, \mathcal{L}, f)$ with respect to perturbations on the underlying manifold.

A. Stability of MNNs to LB Operator Perturbations

MNNs inherit stability to perturbations of the LB operator from the manifold filters that compose the filterbanks in each one of their layers. This result is stated in general form—encompassing both absolute and relative perturbations—in the following theorem.

Theorem 4 (MNN stability): Consider a compact embedded manifold $\mathcal{M}$ with LB operator $\mathcal{L}$. Let $\Phi(\mathbf{H}, \mathcal{L}, f)$ be an L -layer MNN on $\mathcal{M}$ (32) with $F_0 = F_L = 1$ input and output features and $F_l = F, l = 1, 2, \dots, L-1$ features per layer. The filters $\mathbf{h}(\mathcal{L})$ and nonlinearity functions satisfy Assumptions 1 and 2 respectively. Let $\mathcal{L}'$ be the perturbed LB operator [cf. Definition 3 or Definition 7] with $\max\{\alpha, 2, |\gamma/1 - \gamma|\} \gg \epsilon$. If the manifold filters satisfy $\|\mathbf{h}(\mathcal{L})f - \mathbf{h}(\mathcal{L}')f\|_{L^2(\mathcal{M})} \leq C_{per}\epsilon\|f\|_{L^2(\mathcal{M})}$, it holds that

$$\|\Phi(\mathbf{H}, \mathcal{L}, f) - \Phi(\mathbf{H}, \mathcal{L}', f)\|_{L^2(\mathcal{M})} \leq LF^{L-1}C_{per}\epsilon\|f\|_{L^2(\mathcal{M})}.$$

Proof: See Appendix D. $\square$

Theorem 4 reflects that the stability of the MNN is affected by the hyperparameters of the MNN architecture and the stability constant of the manifold filters C_{per} . More explicitly, the stability 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 stability result also shows that there is a linear dependence on the stability constant C_{per} of manifold filters $\mathbf{h}(\mathcal{L})$ and the perturbation size ϵ . As we have shown in Section III-A and III-B, the stability constant is determined by the form of the perturbations (Definition 3 or Definition 7) as well as the spectrum separation achieved by the specific manifold filters (Definition 5 or Definition 9) with corresponding

Lipschitz conditions (Definition 6 or Definition 10). We address the specific cases as follows.

Proposition 4: With the same conditions as Theorem 4, consider the following perturbation models.

- 1) If the perturbed LB operator $\mathcal{L}'$ is an absolute perturbation, i.e., $\mathcal{L}' = \mathcal{L} + \mathbf{A}$ [cf. Definition 3] with $\|\mathbf{A}\| \leq \epsilon$ and the manifold filters $\mathbf{h}(\mathcal{L})$ are α -FDT [cf. Definition 5] with $\alpha \gg \epsilon$ and A_h -Lipschitz continuous [Definition 6] with $\delta = \pi\epsilon/(2\alpha)$, we have

$$C_{per} = \frac{\pi N}{\alpha} + A_h, \quad (33)$$

where N is the size of the α -separated spectrum partition [cf. Definition 4].

- 2) If the perturbed LB operator $\mathcal{L}'$ is a relative perturbation, i.e. $\mathcal{L}' = \mathcal{L} + \mathbf{E}\mathcal{L}$ [cf. Definition 7] with $\|\mathbf{E}\| \leq \epsilon$, and the manifold filters $\mathbf{h}(\mathcal{L})$ are γ -FRT [cf. Definition 9] with $\gamma/(1-\gamma) \gg \epsilon$ and B_h -integral Lipschitz continuous [Definition 10] with $\delta = \pi\epsilon/(2\gamma)$, we have

$$C_{per} = \frac{\pi M}{\gamma} + B_h, \quad (34)$$

where M is the size of the γ -separated spectrum partition [cf. Definition 8].

Proof: The conclusions follow directly from Theorem 4 combined with Theorem 2 or Theorem 3 under the corresponding assumptions. $\square$

Combining Theorem 4 with Proposition 4, we observe that α -FDT manifold filters with Lipschitz continuity can be composed to construct MNNs which are stable to absolute perturbations; while γ -FRT manifold filters with integral Lipschitz continuity can be composed to construct MNNs which are stable relative perturbations of the LB operator. Explicitly, by inserting the stability constant C_{per} in (23), we see that other than the perturbation size ϵ , there are three terms that determine the stability of MNNs. The first term is LF^{L-1} , which, as we have already discussed, is decided by the number of layers and filters in the MNN architecture. This term arises due to the propagation of the underlying operator perturbations across all the manifold filters in all layers of the MNN. The second term is $\pi N/\alpha$ or $\pi M/\gamma$, which results from the deviations of the eigenfunctions as well as from the frequency response variations within the same eigenvalue partition. Finally, the third term, A_h or B_h , is given by the Lipschitz or integral Lipschitz constants which are decided during the filter design or the training process. It is important to note that the stability constant C_{per} brings along the trade-off between stability and discriminability. However, unlike manifold filters, MNNs can be both stable and discriminative. This arises from the effects of nonlinear activation functions, as we discuss in further detail in Section V.

B. Stability of MNNs to Manifold Deformations

In Theorem 4 and Proposition 4, we established the conditions under which MNNs are stable to either absolute or relative perturbations of the LB operator as defined in Definitions 3 and 7. Since a manifold deformation $\tau: \mathcal{M} \rightarrow \mathcal{M}$,

with $\text{dist}(x, \tau(x)) \leq \epsilon$ and $\|J_x(\tau) - I\|_F \leq \epsilon$ for all $x \in \mathcal{M}$, translates into both an absolute and a relative perturbation of the Laplace-Beltrami operator, MNNs composed of manifold filters meeting all of these conditions in items 1 and 2 of Proposition 4, i.e., the manifold filters are α -FDT and γ -FRT, and both Lipschitz continuous and integral Lipschitz continuous, can be proved to be stable under the manifold deformation. The spectrum can be made to be both α -separated and γ -separated by making sure the eigenvalues in different partitions satisfy both (20) and (27). Assuming that all of these conditions are met, we can state our main result—that MNNs are stable to deformations of the manifold—as follows.

Theorem 5: Let $\mathcal{M}$ be a compact embedded manifold with LB operator $\mathcal{L}$ and f be a manifold signal. We construct $\Phi(\mathbf{H}, \mathcal{L}, f)$ as a MNN on $\mathcal{M}$ (32) where the filters $\mathbf{h}(\mathcal{L})$ are α -FDT [cf. Definition 5], α/λ_1 -FRT [cf. Definition 9], A_h -Lipschitz [cf. Definition 6] and B_h -integral Lipschitz [cf. Definition 10]. Consider a deformation on $\mathcal{M}$ as $\tau: \mathcal{M} \rightarrow \mathcal{M}$ where $\text{dist}(x, \tau(x)) \leq \epsilon$ and $J_x(\tau) = I + \Delta_x$ with $\|\Delta_x\|_F \leq \epsilon$ for all $x \in \mathcal{M}$ and $\epsilon \ll \min(\alpha/\lambda_1, \alpha, 2)$. Under Assumptions 1 and 2 it holds that

$$\|\Phi(\mathbf{H}, \mathcal{L}, f) - \Phi(\mathbf{H}, \mathcal{L}', f)\|_{L^2(\mathcal{M})} = O(\epsilon)\|f\|_{L^2(\mathcal{M})}. \quad (35)$$

Together, Theorem 1 and Theorem 4 imply that MNNs are stable to the manifold deformations v introduced in the beginning of this section. This is because these deformations spawn a perturbation of the LB operator that consists of both an absolute and a relative perturbation. For stability to hold, the filters that make up the layers of the MNN need to be α -FDT [cf. Definition 5], γ -FRT [cf. Definition 9], Lipschitz [cf. Definition 6] and integral Lipschitz [cf. Definition 10]. We can propose an easier special case to relate α and γ by utilizing the spectrum property of LB operator. By setting the α -FDT filter with $\alpha = \gamma\lambda_1$, eigenvalues $\lambda_i, \lambda_{i+1} \in \Lambda_k(\alpha)$ would lead to $\lambda_i, \lambda_{i+1} \in \Lambda_l(\gamma)$ due to the fact that

$$\lambda_{i+1} - \lambda_i \leq \alpha = \gamma\lambda_1 \leq \gamma\lambda_i, \quad (36)$$

with λ_1 indexed as the smallest eigenvalue in the spectrum. The requirement that the filter be α -FDT can be removed as long as $\lambda_1 > 0$ and $\alpha = \gamma\lambda_1$, since a γ -FRT filter is always $\gamma\lambda_1$ -FDT, i.e. α -FDT.

V. DISCUSSION

Stability vs. discriminability tradeoff. In both stability theorems for manifold filters (Theorems 2, 3) and in the stability theorem for MNNs (Theorem 4), the stability bounds depend on the frequency partition threshold (α or γ), the number of total partitions (N or M) and the Lipschitz continuity constant (A_h or B_h). The frequency partition threshold and the number of partitions have a combined effect on stability. As indicated by Definitions 4 and 8, a larger frequency threshold leads to a smaller number of singletons, as eigenvalues that would otherwise be separated for small thresholds end up being grouped when the threshold is large. While a large frequency threshold results in a larger number of partitions that contain more than one eigenvalue, the total number of partitions (N

or M) either stays the same or decreases because the number of eigenvalues does not exceed the number of partitions [cf. Proposition 2 or 3]. Thus, a larger frequency threshold and a smaller number of partitions both lead to a smaller stability bound. Simultaneously, a large frequency threshold makes the spectrum separated more sparsely. Therefore, a large number of eigenvalues are amplified in a similar manner, which makes the filter function less discriminative. Considering in the limit, if the frequency threshold goes to infinity, all the eigenvalues tend to be grouped and the filter would not discriminate the whole spectrum. This would lead to a very stable filter but there is no discriminability at all. The Lipschitz constant (A_h or B_h) affects stability and discriminability in similar ways. Smaller Lipschitz constants decrease the stability bound, but lead to smoother filter functions giving similar frequency responses to different eigenvalues. Hence, in both manifold filters and MNNs we observe a trade-off between stability and discriminability. Nevertheless, in MNNs this trade-off is alleviated due to the presence of nonlinearities as discussed below.

Pointwise nonlinearity. As demonstrated by Propositions 2 and 3, large eigenvalues of LB operator tend to be grouped together in one large group and share similar frequency responses. This is part of the reason why manifold filters have a stability-discriminability tradeoff, which implies that they cannot be stable and discriminative at the same time. However, in MNNs this problem is circumvented with the addition of nonlinearities. Nonlinearities have the effect of scattering the spectral components all over the eigenvalue spectrum. In the MNN, they mix the frequency components by spilling spectral components associated with the large eigenvalues that tend to be very close onto the smaller eigenvalues that could be more separated, where they can then be discriminated by the manifold filters in the following layer. This is consistent with the role of nonlinear activation functions in graph neural networks (GNNs) [20], which can be seen as instantiations of MNNs on discrete samples of the manifold as further discussed in Section VI.

Comparison with graphons. The graphon is another infinite-dimensional model that can represent the limit of convergent sequences of graphs, and a series of works have proved stability of graphon neural networks and the transferability of GNNs sampled from them [24], [26], [27], [28], [39]. Manifolds can represent the limits of relatively sparse graphs including ϵ -graphs and k -NN graphs [29]. While graphons can also be seen as the limit model of relatively sparse graphs [28], embedded manifolds in high-dimensional spaces are more realistic geometric models with physical interpretations in a number of application scenarios, such as point clouds, 3D shape segmentation and classification. Other important differences are that (i) the stability analysis on graphon models in [25], [39] focuses on deformations to the adjacency matrix of the graph, which can be translated directly as perturbations of the graphon operator, and that (ii) in the case of graphons, only an absolute perturbation model makes sense since given that the graphon spectrum is bounded a relative perturbation can always be bounded by an absolute perturbation. Meanwhile, deformations to the manifold domain translate into a combination of absolute

and relative perturbations of the LB operator, and the fact that the LB operator spectrum is unbounded makes the effects of absolute and relative perturbations distinct, especially in the high-frequency domain.

VI. FROM MANIFOLD NEURAL NETWORKS TO GRAPH NEURAL NETWORKS

MNNs are built from manifold convolutional filters (Definition 1) operating on a continuous manifold and over an infinite time horizon. This makes it impractical to implement directly the architecture described by (32) in applications. In this section, we discuss how MNNs are implemented in practice over a set of discrete samples from the manifold in a finite and discrete time frame.

A. Discretization in the Space Domain

In practice, the explicit form of the manifold and of its LB operator are unknown. What we typically have access to is a point cloud representation of the manifold, i.e., a discrete set of sampling points. From these points' coordinates, the structure of the manifold is approximated by a geometric or a nearest neighbor graph [29], [40], [41]. The LB operator is then approximated by the graph Laplacian, which can be shown to converge to the LB operator as the number of sampling points grows [29], [40].

Explicitly, suppose that $X = \{x_1, x_2, \dots, x_n\}$ is a set of n points sampled i.i.d. from measure μ of manifold $\mathcal{M}$, which is embedded in $\mathbb{R}^N$. We can construct a complete weighted symmetric graph $\mathbf{G}_n$ by taking the sampled points to be the vertices of the graph and setting the edge weights based on the Euclidean distance between pairs of points. Specifically, the weight w_{ij} associated with edge (i, j) is given by

$$w_{ij} = \frac{1}{n} \frac{1}{t_n (4\pi t_n)^{k/2}} \exp\left(-\frac{\|x_i - x_j\|^2}{4t_n}\right), \quad (37)$$

where $\|x_i - x_j\|$ is the Euclidean distance between points x_i and x_j while t_n is a parameter associated with the chosen Gaussian kernel [41]. The adjacency matrix $\mathbf{A}_n \in \mathbb{R}^{n \times n}$ is thus defined as $[\mathbf{A}_n]_{ij} = w_{ij}$ for $1 \leq i, j \leq n$ and the corresponding graph Laplacian matrix $\mathbf{L}_n$ [42] is given by

$$\mathbf{L}_n = \text{diag}(\mathbf{A}_n \mathbf{1}) - \mathbf{A}_n. \quad (38)$$

We interpret $\mathbf{L}_n$ the Laplacian operator of the constructed graph $\mathbf{G}_n$. Similarly, we define a uniform sampling operator $\mathbf{P}_n : L^2(\mathcal{M}) \rightarrow L^2(\mathbf{G}_n)$ to sample manifold signals. Given a manifold signal f , we can use operator $\mathbf{P}_n$ to sample graph signals $\mathbf{x}_n \in \mathbb{R}^n$ as

$$\mathbf{x}_n = \mathbf{P}_n f \text{ with } [\mathbf{x}_n]_i = f(x_i), \quad x_i \in X, \quad (39)$$

where the i -th entry of the graph signal $\mathbf{x}_n$ is the manifold signal f evaluated at the sample point x_i .

In Section II-B, we have shown that the manifold filter $\mathbf{h}$ is parametric with the LB operator $\mathcal{L}$. Therefore, we can also

parameterize $\mathbf{h}$ with the discrete graph Laplacian operator $\mathbf{L}_n$, which is written as,

$$\mathbf{z}_n = \int_0^\infty \tilde{h}(t) e^{-t\mathbf{L}_n} dt \mathbf{x}_n = \mathbf{h}(\mathbf{L}_n) \mathbf{x}_n, \quad \mathbf{x}_n, \mathbf{z}_n \in \mathbb{R}^n, \quad (40)$$

where $\mathbf{z}_n$, the output of the filter, is now a discrete graph signal. By cascading these discrete filters operated on graph $\mathbf{G}_n$ and pointwise nonlinearities layer after layer, we can then approximate the MNN on $\mathbf{G}_n$ as

$$\mathbf{x}_l^p = \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathbf{L}_n) \mathbf{x}_{l-1}^q \right), \quad (41)$$

where $\mathbf{h}_l^{pq}(\mathbf{L}_n)$ maps the q -th feature in the $l-1$ -th layer to the p -th feature in the l -th layer, $1 \leq q \leq F_{l-1}$ and $1 \leq p \leq F_l$, and F_l denotes the number of features in the l -th layer (we have dropped the subscript n in $\mathbf{x}_l^p$ and $\mathbf{x}_{l-1}^q$ for simplicity). After gathering the filter functions in the set $\mathbf{H}$, this neural network can be represented more succinctly as $\Phi(\mathbf{H}, \mathbf{L}_n, \mathbf{x})$.

Equation (41) is a consistent approximation of the MNN because, as n goes to infinity, the discrete graph Laplacian operator $\mathbf{L}_n$ of the graph $\mathbf{G}_n$ converges to the LB operator $\mathcal{L}$ of the manifold $\mathcal{M}$, and the sampled graph signal $\mathbf{x}_n$ converges to the manifold signal f [41]. These facts combinly imply that the output of the neural network on the graph $\mathbf{G}_n$ converges to the output of the neural network on the continuous manifold as stated in the following.

Proposition 5: Let $X = \{x_1, x_2, \dots, x_n\}$ be n points sampled i.i.d. from measure μ of d -dimensional manifold $\mathcal{M} \subset \mathbb{R}^N$, with corresponding sampling operator $\mathbf{P}_n$ (39). Let $\mathbf{G}_n$ be a discrete graph approximation of $\mathcal{M}$ constructed from X as in (37) with $t_n = n^{-1/(d+2+a)}$ and $a > 0$. Let $\Phi(\mathbf{H}, \cdot, \cdot)$ be a neural network parameterized either by the LB operator $\mathcal{L}$ of the manifold $\mathcal{M}$ or the graph Laplacian operator $\mathbf{L}_n$ of $\mathbf{G}_n$ with the filters in $\mathbf{H}$ satisfying Definition 5 and 6. It holds that

$$\lim_{n \rightarrow \infty} \|\Phi(\mathbf{H}, \mathbf{L}_n, \mathbf{P}_n f) - \mathbf{P}_n \Phi(\mathbf{H}, \mathcal{L}, f)\|_{L^2(\mathbf{G}_n)} = 0, \quad (42)$$

with the limit taken in probability.

Proof: See Section B in supplementary material. $\square$

This proposition provides theoretical support to state that neural networks constructed from the discrete Laplacian $\mathbf{L}_n$ converge to MNN and thus can inherit the stability properties of the MNN.

B. Discretization in the Time Domain

In order to learn an MNN (32), we need to learn the manifold convolutional filters $\mathbf{h}_l^{pq}$. This means that we need to learn the impulse responses $\tilde{h}(t)$ in Definition 1. However, learning continuous functions $\tilde{h}$ is computationally infeasible, so we sample $\tilde{h}$ over fixed intervals of duration T_s and parameterize the filter with coefficients $h_k = \tilde{h}(kT_s)$, $k = 0, 1, 2, \dots$. Setting the sampling interval to $T_s = 1$ for simplicity, the discrete-time manifold convolution can be written as

$$\mathbf{h}(\mathcal{L})f(x) = \sum_{k=0}^{\infty} h_k e^{-k\mathcal{L}} f(x) \quad (43)$$

where $\{h_k\}_{k=0}^{\infty}$ are called the filter coefficients or taps.

Yet, learning (43) is still impractical because there is an uncountable number of parameters h_k . To address this, we fix a time horizon of K time steps and rewrite (43) as

$$\mathbf{h}(\mathcal{L})f(x) = \sum_{k=0}^{K-1} h_k e^{-k\mathcal{L}} f(x) \quad (44)$$

which can be seen as a finite impulse response (FIR) filter with shift operator $e^{-\mathcal{L}}$. Indeed, the frequency response of this filter [cf. Proposition 1] is given by

$$\hat{h}(\lambda) = \sum_{k=0}^{K-1} h_k e^{-k\lambda}. \quad (45)$$

Combining (40) and (44), we can bring the discretization over the spatial and time domains together to rewrite the convolution operation on the discretized manifold and in the discrete-time domain, explicitly,

$$\mathbf{z}_n = \mathbf{h}(\mathbf{L}_n) \mathbf{x}_n = \sum_{k=0}^{K-1} h_k e^{-k\mathbf{L}_n} \mathbf{x}_n. \quad (46)$$

Equation (46) recovers the definition of the graph convolution [6] with graph shift operator $e^{-\mathbf{L}_n}$. This means that in practice we implement MNNs as graph neural networks (GNNs). Therefore, the stability behavior of the GNN can be seen as a proxy for the stability behavior of the MNN. We will leverage this idea in the numerical experiments of Section VII.

Remark 2: We analyze the convergence of GNNs to MNNs when graphs are constructed based on uniformly sampled points from the manifold in Proposition 5. The convergence holds when GNNs and MNNs share the same filter parameters while the filter functions are continuous in the time domain. We discuss the discretization of manifold filters in the time domain from a practical aspect that neural networks are usually trained and operated in digital systems with digital filters.

VII. NUMERICAL EXPERIMENTS

Dataset. We evaluate our MNN stability results on the ModelNet10 [43] classification problem. The dataset contains 3991 meshed CAD models from 10 categories for training and 908 models for testing. For each model, 300 points are uniformly randomly sampled from all points of the model to form the point cloud. Each point is characterized by the 3D coordinates as features. We formulate the problem by modeling a dense graph neural network model to approximate MNN. Each node in the graph can be modeled as the sampling point and each edge weight is constructed based on the distance between each pair of nodes. In this work our goal is to identify the CAD model for chairs as is illustrated in Fig. 3 with the models for chair labeled as 1 and the others as 0. We deform the underlying manifold structure by adding random perturbations to the coordinates of the sampling points. By comparing the differences of the classification error rates, we aim to show that MNNs with Lipschitz continuous and integral Lipschitz continuous manifold filters are stable via looking into the performance of the approximated GNNs.




Fig. 3. Point cloud models with 300 sampling points in each model. Our goal is to identify chair models from other models such as toilet and table.

TABLE I
CLASSIFICATION ERROR RATES FOR MODEL ‘CHAIR’ IN THE
TEST DATASET. AVERAGE OVER 5 DATA REALIZATIONS.
THE NUMBER OF NODES IS $n = 300$

Architecture	Error Rates
GNN1Ly	$8.04\% \pm 0.88\%$
GNN2Ly	$4.30\% \pm 2.64\%$
GF1Ly	$13.77\% \pm 6.87\%$
GF2Ly	$12.22\% \pm 7.89\%$

Neural network architectures. We build dense graphs to approximate the point cloud models. We use the coordinates of each point as node features. By connecting a point with all the other points in the point cloud, the edge weight is defined based the distance between every two points and a Gaussian kernel. The Laplacian matrix is calculated for each input point cloud model. We implement different architectures, including Graph Filters (GF) and Graph Neural Networks (GNN) with 1 and 2 layers, to solve the classification problem. The architectures with a single layer contain $F_0 = 3$ input features which are the 3d coordinates of each point, $F_1 = 64$ output features and $K = 5$ filter taps. While the architectures with 2 layers has another layer with $F_2 = 32$ features and 5 filter taps. We use the ReLU as nonlinearity. The learned graph filters are not regularized in architectures with ‘NoPel’ while graph filters in the other architectures are both Lipschitz and integral Lipschitz. We approximate the spectrum partitions with the continuous assumptions of the filter functions. All architectures also include a linear readout layer mapping the final output features to a binary scalar that estimates the classification.

Discriminability experiment. We train all the architectures with an ADAM optimizer [44] with learning rate set as 0.005 and decaying factors as 0.9, 0.999 by minimizing the entropy loss. The training point cloud models are divided in batches of 10 over 40 epochs. We run 5 random point samplings for all the architectures and we show the average classification error rates across these realizations as well as the standard deviation in Table I. We can observe that with the use of non-linearity, Graph Neural Networks perform better compared with Graph Filters. Architectures with more layers learn more accurate models which also leads to better performances.

Stability experiment. We test the same trained Graph Neural Networks and Graph Filters with 2 layers on perturbed test point cloud models with different perturbation levels. We perturb the test point clouds by adding a Gaussian random variable with




Fig. 4. Difference between error rates on the original test dataset and the deformed one.




Fig. 5. Difference between error rates on the original test dataset and the deformed one.

mean ϵ and variance 2ϵ to each coordinate of every sampling point, which can be seen as a deformation of the underlying manifold. We measure the stability by computing the difference between the error rates achieved based on the original test point cloud models and the perturbed ones. In Fig. 5, we see that this difference increases when the perturbations become larger, but overall the differences are small. We also observe that Graph Neural Network is more stable compared with Graph Filters. Furthermore, the Graph Neural Networks and Graph Filters with Lipschitz continuous and integral Lipschitz continuous filters are more stable. Both of these observations validate our stability results.

To further verify the discriminability under perturbations, we trained and tested the architectures with perturbed dataset. We can see from Table II that both GNN and GF can identify the chair model with small error rates while the error rates grow

TABLE II
CLASSIFICATION ERROR RATES FOR MODEL ‘CHAIR’ WITH PERTURBED
TRAINING AND TEST DATASET. AVERAGE OVER 5 DATA REALIZATIONS.
THE NUMBER OF NODES IS $n = 300$

Architecture	$\epsilon = 0.2$	0.4
GNN2Ly	$7.37\% \pm 1.43\%$	$7.71\% \pm 3.96\%$
GF2Ly	$13.76\% \pm 6.82\%$	$13.54\% \pm 7.16\%$
Architecture	$\epsilon = 0.6$	0.8
GNN2Ly	$8.04\% \pm 2.83\%$	$11.01\% \pm 6.33\%$
GF2Ly	$14.76\% \pm 5.67\%$	$16.04\% \pm 6.34\%$




Fig. 6. Difference between error rates on the original test dataset and the deformed one on growing size of graphs.

slightly with the increase of perturbation levels. GNNs still outperform GFs in discriminability with the help of nonlinearity.

We further look into the stability of GNNs with respect to the size of the graph with different perturbation levels. We study the stability of a 2-layer GNN with Lipschitz and integral Lipschitz continuous filters on graphs with a maximum of $n = 896$ points. Then we randomly sub-sample to generate graphs with the number of nodes $n = 128, 256, \dots, 896$. We plot the difference of the error rates achieved by the GNN between the original point cloud models and the perturbed point cloud models. The position of each point in the perturbed cloud is displaced by $\epsilon/\sqrt[3]{n}$. Since this is a point cloud in three dimensions, the normalization by $1/\sqrt[3]{n}$ is such that the gradient of the deformation stays constant across different sub-samplings. In Fig. 6, we show the stability of a 2-layer GNN with Lipschitz and integral Lipschitz continuous filters as we vary the number of points in the point cloud. We observe from the figure that the stability bound does not grow with the number of nodes. This experiment shows that MNN stability bounds are a more accurate model of the behavior of GNNs when graphs are sampled from a manifold while generic stability bounds predict a stability error that can grow with $\sqrt{n}$ [20]. We emphasize that this growing stability bounds of [20] are due to the fact that the graph and perturbation model are generic. Our bounds are tighter with a specific assumption on the graph – it is sampled from a manifold – and the perturbation model – it is a bounded manifold deformation with bounded gradient norm.

VIII. CONCLUSION

In this paper, we have defined manifold convolutions and manifold neural networks. We prove that the deformations on

the embedded submanifolds can be represented as a form of perturbations to the Laplace-Beltrami operator. Considering the infinite dimensionality of LB operators, we import the definition of frequency difference threshold filters and frequency ratio threshold filters to help separate the spectrum. By assigning similar frequency responses to the eigenvalues that are close enough, these filters can be proved to be stable under absolute and relative perturbations to the LB operator respectively with Lipschitz continuous assumptions. While the manifold filters need to trade-off between the stability and discriminability. MNNs composed with layers of manifold filters and pointwise nonlinearities can be proved to be stable to absolute and relative perturbations to the LB operators. While the frequency mixing brought by pointwise nonlinearity can help with the discriminability. We conclude that the MNNs are thus both stable to deformations and discriminative. We also show the discretizations of MNNs in both spatial and time domains to make our proposed MNNs implementable. We finally verified our results numerically with a point cloud classification problem with ModelNet10 dataset.

APPENDIX

A. Proof of Theorem 1

We first introduce the concepts of the tangent map and the Jacobian matrix [45].

Definition 11: Let $\mathcal{M}$ be the manifold and $\tau : \mathcal{M} \rightarrow \mathcal{M}$. For any $x \in \mathcal{M}$, we define the tangent map to be the linear map

$$\tau_{*,x} : T_x \mathcal{M} \rightarrow T_{\tau(x)} \mathcal{M}. \quad (47)$$

When τ is a smooth map and $\mathcal{M}$ is an embedded manifold in $\mathbb{R}^N$, for $x \in \mathcal{M}$, the tangent space $T_x \mathcal{M}$ can be canonically identified with $\mathbb{R}^N$, using the basis

$$\left. \frac{\partial}{\partial x^1} \right|_x, \left. \frac{\partial}{\partial x^2} \right|_x \cdots \left. \frac{\partial}{\partial x^N} \right|_x \in T_x \mathcal{M}, \quad (48)$$

similarly the tangent space $T_{\tau(x)} \mathcal{M}$ is equipped with the basis

$$\left. \frac{\partial}{\partial y^1} \right|_{\tau(x)}, \left. \frac{\partial}{\partial y^2} \right|_{\tau(x)} \cdots \left. \frac{\partial}{\partial y^N} \right|_{\tau(x)} \in T_{\tau(x)} \mathcal{M}, \quad (49)$$

and is just the Jacobian matrix $J_x(\tau)$ of τ at x ,

$$\tau_* \left(\left. \frac{\partial}{\partial x^i} \right|_x \right) = \sum_{j=1}^N \frac{\partial \tau^j}{\partial x^i} \Big|_x \cdot \left. \frac{\partial}{\partial y^j} \right|_{\tau(x)}, \quad (50)$$

where $\left. \frac{\partial \tau^j}{\partial x^i} \right|_x$ is the entry of the Jacobian matrix $[J_x(\tau)]_{i,j}$ and τ^j is the j -th component of τ .

With the tangent map $\tau_{*,x}$ as the linear map between the tangent spaces $T_x \mathcal{M}$ to $T_{\tau(x)} \mathcal{M}$ and based on the definition of $\mathcal{L}'$ in equation (16) together with the definition of Laplace-Beltrami operator in (5), the operation carried out on the deformed manifold data f can be written as

$$-\mathcal{L}' f(x) = (\nabla \cdot \nabla) f(\tau(x)) \quad (51)$$

$$= (J_x(\tau)^T \nabla_\tau \cdot J_x(\tau)^T \nabla_\tau) f(\tau(x)). \quad (52)$$

The equality in (52) results from the chain rule of gradient operator where ∇_τ is denoted as the intrinsic gradient around $\tau(x)$

in the tangent space $T_{\tau(x)}\mathcal{M}$. By replacing $J_x(\tau) = I + \Delta_x$ the inner product term, (52) can be rewritten as

$$J_x(\tau)^T \nabla_\tau \cdot J_x(\tau)^T \nabla_\tau = \nabla_\tau \cdot \nabla_\tau + 2(\Delta_x^T \nabla_\tau \cdot \nabla_\tau) + \Delta_x^T \nabla_\tau \cdot \Delta_x^T \nabla_\tau. \quad (53)$$

With $\mathcal{L} = -\nabla_\tau \cdot \nabla_\tau$, the perturbed operator is

$$\mathcal{L} - \mathcal{L}' = 2(\Delta_x^T \nabla_\tau \cdot \nabla_\tau) + \Delta_x^T \nabla_\tau \cdot \Delta_x^T \nabla_\tau \quad (54)$$

$$= 2\|\Delta_x\|_F^2 (\nabla_\tau \cdot \nabla_\tau) + \|\Delta_x\|_F^2 (\nabla_\tau \cdot \nabla_\tau) + \mathcal{A}. \quad (55)$$

From (54) to (55), we extract the relative term and use $\mathcal{A}$ to represent the compliment terms. This leads to $\mathbf{E} = \|\Delta_x\|_F^2 + 2\|\Delta_x\|_F$, as the relative perturbation term, the norm of which is bounded by the leading term as $O(\epsilon)$. The norm of the compliment term therefore can be written as

$$\|\mathcal{A}\| = \|\mathbf{E}(\nabla_\tau \cdot \nabla_\tau) - 2(\Delta_x^T \nabla_\tau \cdot \nabla_\tau) - \Delta_x^T \nabla_\tau \cdot \Delta_x^T \nabla_\tau\| \quad (56)$$

$$\leq \|2\|\Delta_x\|_F(\nabla_\tau \cdot \nabla_\tau) - 2(\Delta_x^T \nabla_\tau \cdot \nabla_\tau)\| + \|\|\Delta_x\|_F^2(\nabla_\tau \cdot \nabla_\tau) - \Delta_x^T \nabla_\tau \cdot \Delta_x^T \nabla_\tau\|, \quad (57)$$

which can be also bounded by the leading terms as $O(\epsilon)$ combining with the boundedness of the gradient field.

B. Lemmas

Now we need to include two important lemmas to analyze the influence on eigenvalues and eigenfunctions caused by the perturbation.

Lemma 1: [Weyl's Theorem] The eigenvalues of LB operators $\mathcal{L}$ and perturbed $\mathcal{L}' = \mathcal{L} + \mathcal{A}$ satisfy

$$|\lambda_i - \lambda'_i| \leq \|\mathcal{A}\|, \text{ for all } i = 1, 2, \dots \quad (58)$$

Proof of Lemma 1: The minimax principle asserts that

$$\lambda_i(\mathcal{L}) = \max_{\text{codim} T = i-1} \lambda[\mathcal{L}, T] = \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} \langle \mathcal{L}u, u \rangle. \quad (59)$$

Then for any $1 \leq i$, we have

$$\lambda_i(\mathcal{L}') = \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} \langle (\mathcal{L} + \mathcal{A})u, u \rangle \quad (60)$$

$$= \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} (\langle \mathcal{L}u, u \rangle + \langle \mathcal{A}u, u \rangle) \quad (61)$$

$$\geq \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} \langle \mathcal{L}u, u \rangle + \lambda_1(\mathcal{A}) \quad (62)$$

$$= \lambda_1(\mathcal{A}) + \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} \langle \mathcal{L}u, u \rangle \quad (63)$$

$$= \lambda_i(\mathcal{L}) + \lambda_1(\mathcal{A}). \quad (64)$$

Similarly, we can have $\lambda_i(\mathcal{L}') \leq \lambda_i(\mathcal{L}) + \max_k \lambda_k(\mathcal{A})$. This leads to $\lambda_1(\mathcal{A}) \leq \lambda_i(\mathcal{L}') - \lambda_i(\mathcal{L}) \leq \max_k \lambda_k(\mathcal{A})$. This leads to the conclusion that:

$$|\lambda'_i - \lambda_i| \leq \|\mathcal{A}\|. \quad (65)$$

□

To measure the difference of eigenfunctions, we introduce the Davis-Kahan sin θ theorem as follows.

Lemma 2: [Davis-Kahan sin θ Theorem] Suppose the spectra of operators $\mathcal{L}$ and $\mathcal{L}'$ are partitioned as $\sigma \cup \Sigma$ and $\omega \cup \Omega$ respectively, with $\sigma \cap \Sigma = \emptyset$ and $\omega \cap \Omega = \emptyset$. Then we have

$$\|E_{\mathcal{L}}(\sigma) - E_{\mathcal{L}'}(\omega)\| \leq \frac{\pi \|(\mathcal{L}' - \mathcal{L})E_{\mathcal{L}}(\sigma)\|}{2d} \leq \frac{\pi \|\mathcal{L}' - \mathcal{L}\|}{2d}, \quad (66)$$

where d satisfies $\min_{x \in \sigma, y \in \Omega} |x - y| \geq d$ and $\min_{x \in \Sigma, y \in \omega} |x - y| \geq d$.

Proof of Lemma 2: See [46]. □

Lemma 3: The eigenvalues of LB operators $\mathcal{L}$ and perturbed $\mathcal{L}' = \mathcal{L} + \mathbf{E}\mathcal{L}$ with $\|\mathbf{E}\| \leq \epsilon$ satisfy

$$|\lambda_i - \lambda'_i| \leq \epsilon |\lambda_i|, \text{ for all } i = 1, 2, \dots \quad (67)$$

Proof of Lemma 3: With the assumption that $\mathcal{L}' = \mathcal{L} + \mathbf{E}\mathcal{L}$, we have

$$\lambda_i(\mathcal{L} + \mathbf{E}\mathcal{L}) = \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} \langle (\mathcal{L} + \mathbf{E}\mathcal{L})u, u \rangle \quad (68)$$

$$= \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} (\langle \mathcal{L}u, u \rangle + \langle \mathbf{E}\mathcal{L}u, u \rangle) \quad (69)$$

$$= \lambda_i(\mathcal{L}) + \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} \langle \mathbf{E}\mathcal{L}u, u \rangle. \quad (70)$$

For the second term, we have

$$|\langle \mathbf{E}\mathcal{L}u, u \rangle| \leq \|\mathbf{E}\mathcal{L}\| |u| \leq \epsilon \sum_i |\lambda_i(\mathcal{L})| |\xi_i|^2 = \epsilon \langle \mathcal{L}u, u \rangle \quad (71)$$

Therefore, we have

$$\lambda_i(\mathcal{L} + \mathbf{E}\mathcal{L}) \leq \lambda_i(\mathcal{L}) + \epsilon \max_{\text{codim} T \leq i-1} \min_{u \in T, \|u\|=1} \langle \mathcal{L}u, u \rangle \quad (72)$$

$$= \lambda_i(\mathcal{L}) + \epsilon |\lambda_i(\mathcal{L})|, \quad (72)$$

$$\lambda_i(\mathcal{L} + \mathbf{E}\mathcal{L}) \geq \lambda_i(\mathcal{L}) - \epsilon |\lambda_i(\mathcal{L})|, \quad (73)$$

$$\lambda_i(\mathcal{L}) - \epsilon |\lambda_i(\mathcal{L})| \leq \lambda_i(\mathcal{L} + \mathbf{E}\mathcal{L}) \leq \lambda_i(\mathcal{L}) + \epsilon |\lambda_i(\mathcal{L})|, \quad (74)$$

which concludes the proof. □

C. Proof of Theorem 2

In the following, we denote $\langle \cdot, \cdot \rangle_{L^2(\mathcal{M})}$ as $\langle \cdot, \cdot \rangle$ and $\|\cdot\|_{L^2(\mathcal{M})}$ as $\|\cdot\|$ for simplicity. We start by bounding the norm difference between two outputs of filter functions on operators $\mathcal{L}$ and $\mathcal{L}'$ defined in (15) as

$$\|\mathbf{h}(\mathcal{L})f - \mathbf{h}(\mathcal{L}')f\| = \left\| \sum_{i=1}^{\infty} \hat{h}(\lambda_i) \langle f, \phi_i \rangle \phi_i - \sum_{i=1}^{\infty} \hat{h}(\lambda'_i) \langle f, \phi'_i \rangle \phi'_i \right\|. \quad (75)$$

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

$$h^{(0)}(\lambda) = \begin{cases} \hat{h}(\lambda) - \sum_{l \in \mathcal{K}_m} \hat{h}(C_l) & \lambda \in [\Lambda_k(\alpha)]_{k \in \mathcal{K}_s} \\ 0 & \text{otherwise} \end{cases} \quad (76)$$

$$h^{(l)}(\lambda) = \begin{cases} \hat{h}(C_l) & \lambda \in [\Lambda_k(\alpha)]_{k \in \mathcal{K}_s} \\ \hat{h}(\lambda) & \lambda \in \Lambda_l(\alpha) \\ 0 & \text{otherwise} \end{cases} \quad (77)$$

where C_l is some constant in $\Lambda_l(\alpha)$. We can start by analyzing the output difference of $h^{(0)}(\lambda)$. With the triangle inequality, the norm difference can then be written as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle \phi_i - h^{(0)}(\lambda'_i) \langle f, \phi'_i \rangle \phi'_i \right\| \\ &= \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle \phi_i - h^{(0)}(\lambda_i) \langle f, \phi'_i \rangle \phi'_i \right. \\ & \quad \left. + h^{(0)}(\lambda_i) \langle f, \phi'_i \rangle \phi'_i - h^{(0)}(\lambda'_i) \langle f, \phi'_i \rangle \phi'_i \right\| \end{aligned} \quad (78)$$

$$\begin{aligned} & \leq \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle \phi_i - h^{(0)}(\lambda_i) \langle f, \phi'_i \rangle \phi'_i \right\| \\ & \quad + \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi'_i \rangle \phi'_i - h^{(0)}(\lambda'_i) \langle f, \phi'_i \rangle \phi'_i \right\| \end{aligned} \quad (79)$$

$$\begin{aligned} & \leq \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \left(\langle f, \phi_i \rangle \phi_i - \langle f, \phi_i \rangle \phi'_i + \langle f, \phi_i \rangle \phi'_i \right. \right. \\ & \quad \left. \left. - \langle f, \phi'_i \rangle \phi'_i \right) \right\| + \left\| \sum_{i=1}^{\infty} (h^{(0)}(\lambda_i) - h^{(0)}(\lambda'_i)) \langle f, \phi'_i \rangle \phi'_i \right\| \end{aligned} \quad (80)$$

$$\begin{aligned} & \leq \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle (\phi_i - \phi'_i) \right\| \\ & \quad + \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i - \phi'_i \rangle \phi'_i \right\| \\ & \quad + \left\| \sum_{i=1}^{\infty} (h^{(0)}(\lambda_i) - h^{(0)}(\lambda'_i)) \langle f, \phi'_i \rangle \phi'_i \right\| \end{aligned} \quad (81)$$

For the first term in (81), we employ Lemma 2 and therefore we have $\sigma = \lambda_i$ and $\omega = \lambda'_i$, for $\lambda_i \in [\Lambda_k(\alpha)]_{k \in \mathcal{K}_s}$ we can have

$$\|\phi_i - \phi'_i\| \leq \frac{\pi}{2} \frac{\|\mathcal{A}\|}{\alpha - \epsilon} = \frac{\pi}{2} \frac{\epsilon}{\alpha - \epsilon}. \quad (82)$$

Here d can be seen as $d = \min_{\lambda_i \in \Lambda_k(\alpha), \lambda_j \in \Lambda_l(\alpha), k \neq l} |\lambda_i - \lambda'_j|$. Combined with the fact that $|\lambda_i - \lambda_j| > \alpha$ and $|\lambda_i - \lambda'_i| \leq \epsilon$ for all $\lambda_i \in \Lambda_k(\alpha), \lambda_j \in \Lambda_l(\alpha), k \neq l$, we have $d \geq \alpha - \epsilon$. With Cauchy-Schwartz inequality, we have the first term in (81) bounded as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle (\phi_i - \phi'_i) \right\| \\ & \leq \sum_{i=1}^{\infty} |h^{(0)}(\lambda_i)| |\langle f, \phi_i \rangle| \|\phi_i - \phi'_i\| \leq \frac{N_s \pi \epsilon}{2(\alpha - \epsilon)} \|f\|. \end{aligned} \quad (83)$$

The second term in (81) is bounded as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i - \phi'_i \rangle \phi'_i \right\| \\ & \leq \sum_{i=1}^{\infty} |h^{(0)}(\lambda_i)| \|\phi_i - \phi'_i\| \|f\| \leq \frac{N_s \pi \epsilon}{2(\alpha - \epsilon)} \|f\|. \end{aligned} \quad (84)$$

These two bounds are obtained by noting that $|h^{(0)}(\lambda)| < 1$ and $h^{(0)}(\lambda) = 0$ for $\lambda \in [\Lambda_k(\alpha)]_{k \in \mathcal{K}_m}$. The number of eigenvalues

within $[\Lambda_k(\alpha)]_{k \in \mathcal{K}_s}$ is denoted as N_s . The third term in (81) can be bounded by the Lipschitz continuity of h combined with Lemma 1.

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} (h^{(0)}(\lambda_i) - h^{(0)}(\lambda'_i)) \langle f, \phi'_i \rangle \phi'_i \right\|^2 \\ & \leq \sum_{i=1}^{\infty} |h^{(0)}(\lambda_i) - h^{(0)}(\lambda'_i)|^2 |\langle f, \phi'_i \rangle|^2 \\ & \leq \sum_{i=1}^{\infty} A_h^2 |\lambda_i - \lambda'_i|^2 |\langle f, \phi'_i \rangle|^2 \leq A_h^2 \epsilon^2 \|f\|^2. \end{aligned} \quad (85)$$

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

$$\begin{aligned} & \left\| \mathbf{h}^{(l)}(\mathcal{L})f - \mathbf{h}^{(l)}(\mathcal{L}')f \right\| \\ & \leq \left\| (\hat{h}(C_l) + \delta)f - (\hat{h}(C_l) - \delta)f \right\| \leq 2\delta \|f\|, \end{aligned} \quad (86)$$

where $\mathbf{h}^{(l)}(\mathcal{L})$ and $\mathbf{h}^{(l)}(\mathcal{L}')$ are manifold filters with filter function $h^{(l)}(\lambda)$ on the LB operators $\mathcal{L}$ and $\mathcal{L}'$ respectively. Combining the filter functions, we can write

$$\begin{aligned} & \left\| \mathbf{h}(\mathcal{L})f - \mathbf{h}(\mathcal{L}')f \right\| \\ &= \left\| \mathbf{h}^{(0)}(\mathcal{L})f + \sum_{l \in \mathcal{K}_m} \mathbf{h}^{(l)}(\mathcal{L})f - \mathbf{h}^{(0)}(\mathcal{L}')f - \sum_{l \in \mathcal{K}_m} \mathbf{h}^{(l)}(\mathcal{L}')f \right\| \end{aligned} \quad (87)$$

$$\begin{aligned} & \leq \left\| \mathbf{h}^{(0)}(\mathcal{L})f - \mathbf{h}^{(0)}(\mathcal{L}')f \right\| + \sum_{l \in \mathcal{K}_m} \left\| \mathbf{h}^{(l)}(\mathcal{L})f - \mathbf{h}^{(l)}(\mathcal{L}')f \right\| \end{aligned} \quad (88)$$

$$\leq \frac{N_s \pi \epsilon}{\alpha - \epsilon} \|f\| + A_h \epsilon \|f\| + 2(N - N_s) \delta \|f\|, \quad (89)$$

which concludes the proof.

D. Proof of Theorem 4

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

$$\left\| \Phi(\mathbf{H}, \mathcal{L}, f) - \Phi(\mathbf{H}, \mathcal{L}', f) \right\| = \left\| \sum_{q=1}^{F_L} f_L^q - \sum_{q=1}^{F_L} f_L'^q \right\|. \quad (90)$$

The output signal of layer l of MNN $\Phi(\mathbf{H}, \mathcal{L}, f)$ can be written as

$$f_l^p = \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q \right). \quad (91)$$

Similarly, for the perturbed $\mathcal{L}'$ the corresponding MNN is $\Phi(\mathbf{H}, \mathcal{L}', f)$ the output signal can be written as

$$f_l'^p = \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}'^q \right). \quad (92)$$

The difference therefore becomes

$$\begin{aligned} & \|f_l^p - f_l'^p\| \\ &= \left\| \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q \right) - \sigma \left(\sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}'^q \right) \right\|. \end{aligned} \quad (93)$$

With the assumption that σ is normalized Lipschitz, we have

$$\|f_l^p - f_l'^p\| \leq \left\| \sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q - \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}'^q \right\| \quad (94)$$

$$\leq \sum_{q=1}^{F_{l-1}} \left\| \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q - \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}'^q \right\|. \quad (95)$$

By adding and subtracting $\mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}^q$ from each term, combined with the triangle inequality we can get

$$\begin{aligned} & \left\| \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q - \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}'^q \right\| \\ & \leq \left\| \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q - \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}^q \right\| \\ & \quad + \left\| \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}^q - \mathbf{h}_l^{pq}(\mathcal{L}') f_{l-1}'^q \right\| \end{aligned} \quad (96)$$

The first term can be bounded with (89) for absolute perturbations. The second term can be decomposed by Cauchy-Schwartz inequality and non-amplifying of the filter functions as

$$\left\| f_l^p - f_l'^p \right\| \leq \sum_{q=1}^{F_{l-1}} C_{per} \epsilon \|f_{l-1}^q\| + \sum_{q=1}^{F_{l-1}} \|f_{l-1}^q - f_{l-1}'^q\|, \quad (97)$$

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

$$\begin{aligned} \|f_l^p\| & \leq \left\| \sum_{q=1}^{F_{l-1}} \mathbf{h}_l^{pq}(\mathcal{L}) f_{l-1}^q \right\| \leq \sum_{q=1}^{F_{l-1}} \|\mathbf{h}_l^{pq}(\mathcal{L})\| \|f_{l-1}^q\| \\ & \leq \sum_{q=1}^{F_{l-1}} \|f_{l-1}^q\| \leq \prod_{l'=1}^{l-1} F_{l'} \sum_{q=1}^{F_0} \|f^q\|. \end{aligned} \quad (98)$$

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

$$\left\| f_l^p - f_l'^p \right\| \leq l C_{per} \epsilon \left(\prod_{l'=1}^{l-1} F_{l'} \right) \sum_{q=1}^{F_0} \|f^q\|. \quad (99)$$

Replace l with L we can obtain

$$\begin{aligned} & \|\Phi(\mathbf{H}, \mathcal{L}, f) - \Phi(\mathbf{H}, \mathcal{L}', f)\| \\ & \leq \sum_{q=1}^{F_L} \left(L C_{per} \epsilon \left(\prod_{l'=1}^{L-1} F_{l'} \right) \sum_{q=1}^{F_0} \|f^q\| \right). \end{aligned} \quad (100)$$

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

$$\|\Phi(\mathbf{H}, \mathcal{L}, f) - \Phi(\mathbf{H}, \mathcal{L}', f)\| \leq L F^{L-1} C_{per} \epsilon \|f\|, \quad (101)$$

which concludes the proof.

E. Proof of Theorem 3

The decomposition follows the same routine as (75) shows. By decomposing the filter function as (102) and (103), the norm difference can also be bounded separately.

$$h^{(0)}(\lambda) = \begin{cases} \hat{h}(\lambda) - \sum_{l \in \mathcal{K}_m} \hat{h}(C_l) & \lambda \in [\Lambda_k(\gamma)]_{k \in \mathcal{K}_s} \\ 0 & \text{otherwise} \end{cases} \quad (102)$$

$$h^{(l)}(\lambda) = \begin{cases} \hat{h}(C_l) & \lambda \in [\Lambda_k(\gamma)]_{k \in \mathcal{K}_s} \\ \hat{h}(\lambda) & \lambda \in \Lambda_l(\gamma) \\ 0 & \text{otherwise} \end{cases} \quad (103)$$

where now $\hat{h}(\lambda) = h^{(0)}(\lambda) + \sum_{l \in \mathcal{K}_m} h^{(l)}(\lambda)$ with $\mathcal{K}_s$ defined as the group index set of singletons and $\mathcal{K}_m$ the set of partitions that contain multiple eigenvalues. For manifold filter $\mathbf{h}^{(0)}(\mathcal{L})$ with filter function $h^{(0)}(\lambda)$, the norm difference can also be written as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle \phi_i - h^{(0)}(\lambda'_i) \langle f, \phi'_i \rangle \phi'_i \right\| \\ & \leq \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle (\phi_i - \phi'_i) \right\| \\ & \quad + \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i - \phi'_i \rangle \phi'_i \right\| \\ & \quad + \left\| \sum_{i=1}^{\infty} (h^{(0)}(\lambda_i) - h^{(0)}(\lambda'_i)) \langle f, \phi'_i \rangle \phi'_i \right\|. \end{aligned} \quad (104)$$

The difference of the eigenvalues due to relative perturbations can be similarly addressed by Lemma 3.

The first two terms of (104) rely on the differences of eigenfunctions, which can be derived with Davis-Kahan Theorem in Lemma 2, the difference of eigenfunctions can be written as

$$\|\mathbf{E} \mathcal{L} \phi_i\| = \|\mathbf{E} \lambda_i \phi_i\| = \lambda_i \|\mathbf{E} \phi_i\| \leq \lambda_i \|\mathbf{E}\| \|\phi_i\| \leq \lambda_i \epsilon. \quad (105)$$

The first term in (104) then can be bounded as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle (\phi_i - \phi'_i) \right\| \\ & \leq \sum_{i=1}^{\infty} |h^{(0)}(\lambda_i)| |\langle f, \phi_i \rangle| \|\phi_i - \phi'_i\| \leq \sum_{i \in \mathcal{K}_s} \frac{\pi \lambda_i \epsilon}{2 d_i} \|f\|. \end{aligned} \quad (106)$$

Because $d_i = \min\{|\lambda_i - \lambda'_{i-1}|, |\lambda'_i - \lambda_{i-1}|, |\lambda'_{i+1} - \lambda_i|, |\lambda_{i+1} - \lambda'_i|\}$, with Lemma 3 implied, we have

$$|\lambda_i - \lambda'_{i-1}| \geq |\lambda_i - (1 + \epsilon) \lambda_{i-1}|, \quad (107)$$

$$|\lambda'_i - \lambda_{i-1}| \geq |(1 - \epsilon) \lambda_i - \lambda_{i-1}|, \quad (108)$$

$$|\lambda'_{i+1} - \lambda_i| \geq |(1 - \epsilon) \lambda_{i+1} - \lambda_i|, \quad (109)$$

$$|\lambda_{i+1} - \lambda'_i| \geq |\lambda_{i+1} - (1 + \epsilon) \lambda_i|. \quad (110)$$

Combine with Lemma 3 and Definition 8, $d_i \geq \epsilon \gamma + \gamma - \epsilon$:

$$|(1 - \epsilon) \lambda_{i+1} - \lambda_i| \geq |\gamma \lambda_i - \epsilon \lambda_{i+1}| \quad (111)$$

$$= \epsilon \lambda_i \left| 1 - \frac{\lambda_{i+1}}{\lambda_i} + \frac{\gamma}{\epsilon} - 1 \right| \quad (112)$$

$$\geq \lambda_i (\gamma - \epsilon + \gamma \epsilon) \quad (113)$$

This leads to the bound as

$$\left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i \rangle (\phi_i - \phi'_i) \right\| \leq \frac{M_s \pi \epsilon}{2(\gamma - \epsilon + \gamma \epsilon)} \|f\|. \quad (114)$$

The second term in (104) can also be bounded as

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} h^{(0)}(\lambda_i) \langle f, \phi_i - \phi'_i \rangle \phi'_i \right\| \\ & \leq \sum_{i=1}^{\infty} |h^{(0)}(\lambda_i)| \|\phi_i - \phi'_i\| \|f\| \leq \frac{M_s \pi \epsilon}{2(\gamma - \epsilon + \gamma \epsilon)} \|f\|, \end{aligned} \quad (115)$$

which similarly results from the fact that $|h^{(0)}(\lambda)| < 1$ and $h^{(0)}(\lambda) = 0$ for $\lambda \in [\Lambda_k(\gamma)]_{k \in \mathcal{K}_m}$. The number of eigenvalues within $[\Lambda_k(\gamma)]_{k \in \mathcal{K}_s}$ is denoted as M_s .

The third term in (104) is:

$$\begin{aligned} & \left\| \sum_{i=1}^{\infty} (h^{(0)}(\lambda_i) - h^{(0)}(\lambda'_i)) \langle f, \phi'_i \rangle \phi'_i \right\|^2 \\ & \leq \sum_{i=1}^{\infty} \left(\frac{B_h \epsilon |\lambda_i|}{(\lambda_i + \lambda'_i)/2} \right)^2 \langle f, \phi'_i \rangle^2 \leq \left(\frac{2B_h \epsilon}{2 - \epsilon} \right)^2 \|f\|^2, \end{aligned} \quad (116)$$

with the use of Lemma 3 and Definition 10.

Then we need to analyze the output difference of $h^{(l)}(\lambda)$.

$$\begin{aligned} & \left\| \mathbf{h}^{(l)}(\mathcal{L})f - \mathbf{h}^{(l)}(\mathcal{L}')f \right\| \\ & \leq \left\| (\hat{h}(C_l) + \delta)f - (\hat{h}(C_l) - \delta)f \right\| \leq 2\delta \|f\|. \end{aligned} \quad (117)$$

Combine the filter function, we could get

$$\begin{aligned} & \left\| \mathbf{h}(\mathcal{L})f - \mathbf{h}(\mathcal{L}')f \right\| \\ & = \left\| \mathbf{h}^{(0)}(\mathcal{L})f + \sum_{l \in \mathcal{K}_m} \mathbf{h}^{(l)}(\mathcal{L})f - \mathbf{h}^{(0)}(\mathcal{L}')f - \sum_{l \in \mathcal{K}_m} \mathbf{h}^{(l)}(\mathcal{L}')f \right\| \end{aligned} \quad (118)$$

$$\leq \left\| \mathbf{h}^{(0)}(\mathcal{L})f - \mathbf{h}^{(0)}(\mathcal{L}')f \right\| + \sum_{l \in \mathcal{K}_m} \left\| \mathbf{h}^{(l)}(\mathcal{L})f - \mathbf{h}^{(l)}(\mathcal{L}')f \right\| \quad (119)$$

$$\leq \frac{M_s \pi \epsilon}{\gamma - \epsilon + \gamma \epsilon} \|f\| + \frac{2B_h \epsilon}{2 - \epsilon} \|f\| + 2(M - M_s)\delta \|f\|, \quad (120)$$

which concludes the proof.

F. Definition 1 and Convolutional Filters in Continuous Time

The manifold convolution in Definition 1 can also be motivated with a connection to linear time invariant filters. This requires that we consider the differential equation

$$\frac{\partial u(x, t)}{\partial t} = \frac{\partial}{\partial x} u(x, t). \quad (121)$$

This is a one-sided wave equation and it is therefore not an exact analogous of the diffusion equation in (6) – this would require that the second derivative be used in the right of (121). The important observation to make here is that the exponential of the derivative operator is a time shift so that we can write $u(x, t) = e^{t\partial/\partial x} f(x) = f(x - t)$. This is true because $e^{t\partial/\partial x} f(x)$ and $f(x - t)$ are both solutions of (121). It then

follows that Definition 1 particularized to (121) yields the convolution definition

$$g(x) = \int_0^\infty \tilde{h}(t) e^{t\partial/\partial x} f(x) dt = \int_0^\infty \tilde{h}(t) f(x - t) dt. \quad (122)$$

This is the standard definition of time convolutions.

The frequency representation result in Proposition 1 holds for (122) and it implies that standard convolutional filters in continuous time are completely characterized by the frequency response in Definition 2. The more standard definition of a filter's frequency response as the Fourier transform of the impulse response $\tilde{h}(t)$ – as opposed to the Laplace transform we use in Definition 2 – suffices because complex exponentials $e^{j\omega}$ are an orthonormal basis of eigenfunctions of the derivative operator with associated eigenvalues $j\omega$.

REFERENCES

- [1] Z. Wang, L. Ruiz, and A. Ribeiro, "Stability of neural networks on Riemannian manifolds," in *Proc. 29th Eur. Signal Process. Conf. (EUSIPCO)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 1845–1849.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] E. Isufi, A. Loukas, A. Simonetto, and G. Leus, "Autoregressive moving average graph filtering," *IEEE Trans. Signal Process.*, vol. 65, no. 2, pp. 274–288, Jan. 2017.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] J. Zhou et al., "Graph neural networks: A review of methods and applications," *AI Open*, vol. 1, pp. 57–81, 2020.
- [10] W. Fan et al., "Graph neural networks for social recommendation," in *Proc. World Wide Web Conf.*, 2019, pp. 417–426.
- [11] S. Wu, Y. Tang, Y. Zhu, L. Wang, X. Xie, and T. Tan, "Session-based recommendation with graph neural networks," in *Proc. AAAI Conf. Artif. Intell.*, 2019, vol. 33, no. 1, pp. 346–353.
- [12] A. Chowdhury, G. Verma, C. Rao, A. Swami, and S. Segarra, "Unfolding WMMSE using graph neural networks for efficient power allocation," *IEEE Trans. Wireless Commun.*, vol. 20, no. 9, pp. 6004–6017, Sep. 2021.
- [13] Z. Wang, M. Eisen, and A. Ribeiro, "Unsupervised learning for asynchronous resource allocation in ad-hoc wireless networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 8143–8147.
- [14] 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.
- [15] M. M. Bronstein, J. Bruna, T. Cohen, and P. Velicković, "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges," 2021, *arXiv:2104.13478*.
- [16] 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.

- [17] A. Sandryhaila and J. M. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.
- [18] A. V. Oppenheim et al., *Signals & Systems*. London, UK: Pearson, 1997.
- [19] S. Mallat, "Group invariant scattering," *Commun. Pure Appl. Math.*, vol. 65, no. 10, pp. 1331–1398, 2012.
- [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] F. Gama, A. Ribeiro, and J. Bruna, "Stability of graph scattering transforms," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 8038–8048.
- [22] S. Verma and Z.-L. Zhang, "Stability and generalization of graph convolutional neural networks," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery & Data Mining*, 2019, pp. 1539–1548.
- [23] D. Zou and G. Lerman, "Graph convolutional neural networks via scattering," *Appl. Comput. Harmon. Anal.*, vol. 49, no. 3, pp. 1046–1074, 2020.
- [24] L. Ruiz, F. Gama, and A. Ribeiro, "Graph neural networks: Architectures, stability, and transferability," *Proc. IEEE*, vol. 109, no. 5, pp. 660–682, May 2021.
- [25] L. Ruiz, Z. Wang, and A. Ribeiro, "Graphon and graph neural network stability," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 5255–5259.
- [26] 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, 2023.
- [27] L. Ruiz, L. F. Chamon, and A. Ribeiro, "Transferability properties of graph neural networks," *IEEE Trans. Signal Process.*, vol. 71, pp. 3474–3489, 2023.
- [28] 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.
- [29] 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, 2022.
- [30] 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.
- [31] J. Gallier and J. Quaintance, *Differential Geometry and Lie Groups: A Computational Perspective*, vol. 12. New York, NY, USA: Springer-Verlag, 2020.
- [32] Y. Canzani, *Analysis on Manifolds via the Laplacian*. Lecture Notes, ch. 4, 2013, pp. 41–44. [Online]. Available: <http://www.math.harvard.edu/canzani/docs/Laplacian.pdf>
- [33] P. Moon and D. E. Spencer, *Field Theory Handbook: Including Coordinate Systems, Differential Equations and Their Solutions*. New York, NY, USA: Springer-Verlag, 2012.
- [34] W. Arendt, R. Nittka, W. Peter, and F. Steiner, "Weyl's law: Spectral properties of the Laplacian in mathematics and physics," in *Mathematical Analysis of Evolution, Information, and Complexity*, New Jersey, USA: Wiley Online Library, 2009, pp. 1–71.
- [35] S. Musser, "Weyl's law on Riemannian manifolds." Dept. Math., Univ. Chicago, Chicago, IL, USA. [Online]. Available: <http://math.uchicago.edu/may/REU2016/REUPapers/Musser.pdf>
- [36] C. M. Bishop and N. M. Nasrabadi, *Pattern Recognition and Machine Learning*, vol. 4, no. 4. New York, NY, USA: Springer-Verlag, 2006.
- [37] L. W. Tu, *An Introduction to Manifolds*. New York, NY, USA: Springer-Verlag, 2011.
- [38] W. Arendt and W. P. Schleich, *Mathematical Analysis of Evolution, Information, and Complexity*. Hoboken, NJ, USA: Wiley, 2009.
- [39] 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.
- [40] 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, 2021.
- [41] 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.
- [42] R. Merris, "A survey of graph Laplacians," *Linear Multilinear Algebra*, vol. 39, nos. 1–2, pp. 19–31, 1995.
- [43] Z. Wu et al., "3D ShapeNets: A deep representation for volumetric shapes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1912–1920.
- [44] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [45] G. Gross and E. Meinrenken, *Manifolds, Vector Fields, and Differential Forms. An Introduction to Differential Geometry*. New York, NY, USA: Springer-Verlag, 2023.
- [46] A. Seelmann, "Notes on the $\sin 2\theta$ theorem," *Integral Equ. Operator Theory*, vol. 79, no. 4, pp. 579–597, 2014.