Manifold Filters and Neural Networks: Geometric Graph Signal Processing in the Limit

Zhiyang Wang

Electrical and Systems Engineering
University of Pennsylvania

zhiyangw@seas.upenn.edu
http://zhiyangwang.net

April 10 2025

5-10

Large-scale Geometric Graphs in Practice

- Neural networks have been the choice - Efficient solutions are needed over massive amounts of data
- Graph neural networks process information over very large graphs – scalable and stable solutions

  $\Rightarrow$ E.g., wireless communication systems, robotic control systems, point clouds

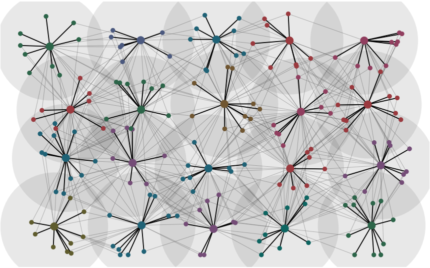

Cellular network

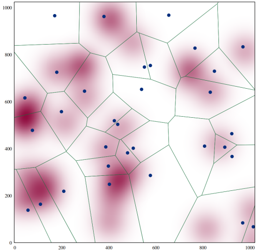

Collaborative robots

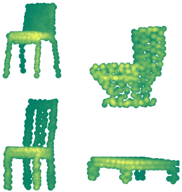

Point clouds

- We study continuous limits of graph NNs as the size of graph grows to infinity $\Rightarrow$ manifold NNs

- Continuous limit model brings insights into sampled discrete models $\Rightarrow$ graphs and images

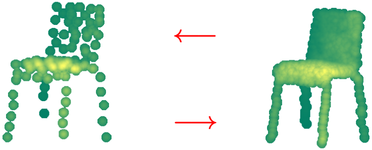

100 nodes   800 nodes

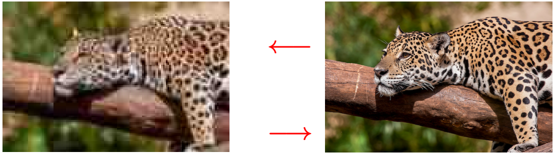

143×95 600×399

- Continuous models easier for theoretical insights $\Leftrightarrow$ Discrete models easier for practical application

Convlutional Filters and Neural Networks in Time

- Convolutional filters in time are linear combinations of time shifted inputs with the time shift operator

$\Rightarrow$

$$y_n = \sum_{k=0}^{K-1} h_k \, x_{n-k}$$

$h_0 x_n \quad + \quad h_1 x_{n-1} \quad + \quad h_2 x_{n-2} \quad + \quad h_3 x_{n-3}$

- Convolutional neural networks (CNNs) compose layers of convolutional filters and non-linearities

Graph Filters and Graph Neural Networks

- Graph convolutional filters are linear combinations of polynomials on graph matrix representations

$\Rightarrow$

$$\mathbf{y} = \sum_{k=0}^{K-1} h_k \, \mathbf{S}^k \, \mathbf{x}$$

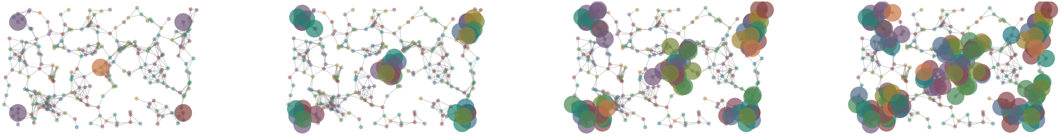

$h_0\mathbf{S}^0\mathbf{x} \quad + \quad h_1\mathbf{S}^1\mathbf{x} \quad + \quad h_2\mathbf{S}^2\mathbf{x} \quad + \quad h_3\mathbf{S}^3\mathbf{x}$

- Graph neural networks (GNNs) compose layers of graph filters and point-wise non-linearities

Manifold Filters and Manifold Neural Networks

- Manifold convolutional filters are linear combinations of Laplace-Beltrami operator exponentials $\Rightarrow$

$$g(x) = \int_0^\infty h(t)\, e^{-t\mathcal{L}} f(x) \mathrm{d}t$$

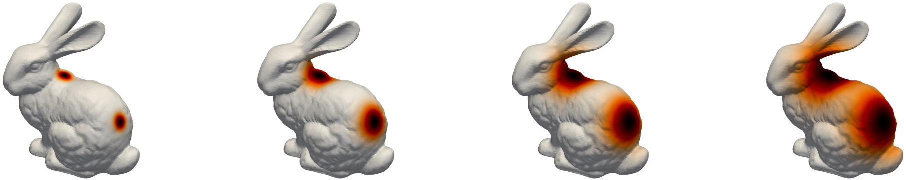

$$h(0T_s)e^{-0T_s\mathcal{L}}f \quad + \quad h(1T_s)e^{-1T_s\mathcal{L}}f \quad + \quad h(2T_s)e^{-2T_s\mathcal{L}}f \quad + \quad h(3T_s)e^{-3T_s\mathcal{L}}f$$

- Manifold neural networks (MNNs) compose layers of manifold filters and point-wise non-linearities

- Manifold convolutional filters are linear combinations of Laplace-Beltrami operator exponentials $\Rightarrow$

$$g(x) \approx \sum_{k=0}^{\infty} h(kT_s)\, e^{-kT_s\mathcal{L}} f(x)$$

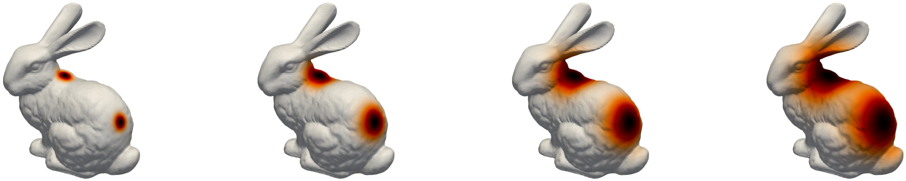

$$h(0T_s)e^{-0T_s\mathcal{L}}f \quad + \quad h(1T_s)e^{-1T_s\mathcal{L}}f \quad + \quad h(2T_s)e^{-2T_s\mathcal{L}}f \quad + \quad h(3T_s)e^{-3T_s\mathcal{L}}f$$

- Manifold neural networks (MNNs) compose layers of manifold filters and point-wise non-linearities

My research focuses on utilizing MNNs to understand fundamental properties of GNNs

Fundamental Properties of Convolutional Neural Networks

- CNNs on discrete time/image signals converge to CNNs on continuous time/image signals

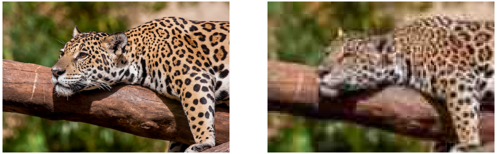

Sample from high res to low res

Deform from high res

- CNNs have two fundamental properties derived from continuous limits that explain their performances
  - $\Rightarrow$ Scalability: Training CNNs with small images is sufficient for transferring to larger images
  - $\Rightarrow$ Stability: CNNs are stable to deformations, which captures the invariance of nature

D. Owerko et al., *Transferability of Convolutional Neural Networks in Stationary Learning Tasks,* arXiv:2307.11588

S. Mallat, *Group invariant scattering,* Communications on Pure and Applied Mathematics

- Graph convolutions are algebraically equivalent to standard convolutions on images

Sample from high res to low res

Deform from high res

- Can we derive these two fundamental properties for GNNs to explain their performances?

Fundamental Properties of GNNs Derived from MNNs

- Graph convolutions are algebraically equivalent to standard convolutions on images

Sample from high res to low res

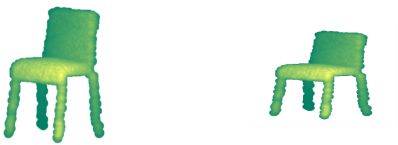

Deform from high res

- GNNs have two fundamental properties derived from MNNs to understand their performances

  $\Rightarrow$ Scalability: Convergence of GNNs to MNNs implies transferability of GNNs across scales

  $\Rightarrow$ Stability: Stability of MNNs to manifold deformations reveals stability of GNNs

- GNNs converge to the underlying MNNs provided the filters satisfy spectral continuity conditions

$$\left\| \Phi(\mathbf{H}, \mathbf{L}_n, \mathbf{P}_n f) - \mathbf{P}_n \Phi(\mathbf{H}, \mathcal{L}, f) \right\| = O\left[\left(\frac{N}{\alpha} + A_h\right)\sqrt{\xi} + \frac{\log(n)}{n}\right] \|f\|_{L^2(\mathcal{M})}$$

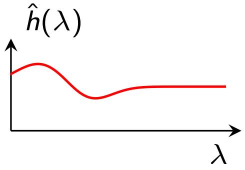

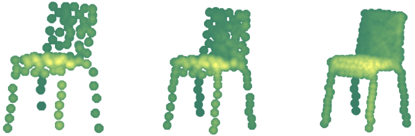

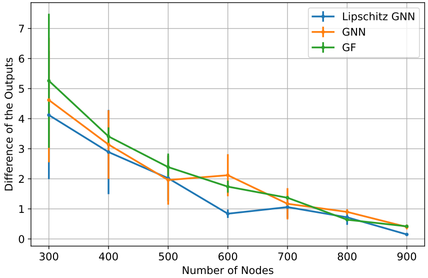

- GNNs converge to the underlying MNNs provided the filters satisfy spectral continuity conditions

$$\left\| \mathbf{\Phi}(\mathbf{H}, \mathbf{L}_n, \mathbf{P}_n f) - \mathbf{P}_n \mathbf{\Phi}(\mathbf{H}, \mathcal{L}, f) \right\| = O\left[\left(\frac{N}{\alpha} + A_h \right) \sqrt{\xi} + \frac{\log(n)}{n} \right] \|f\|_{L^2(\mathcal{M})}$$

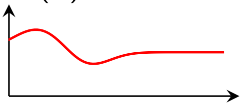

GNNs trained on small graphs with continuous filters are able to transfer to large graphs

Z. Wang et al, *Geometric Graph Filters and Neural Networks: Limit Properties and Discriminability Trade-offs*, IEEE Trans on Signal Processing

- Stability of MNNs to manifold deformations implies stability of GNNs with continuous filters

$$\left\|\boldsymbol{\Phi}(\mathbf{H},\mathcal{L},f)-\boldsymbol{\Phi}(\mathbf{H},\mathcal{L}',f)\right\| = O\left[\left(\frac{N}{\alpha}+A_h+\frac{M}{\gamma}+B_h\right)\epsilon\right]\|f\|_{L^2(\mathcal{M})}$$

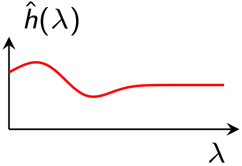

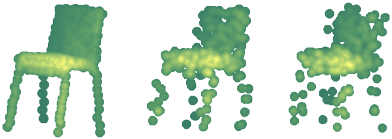

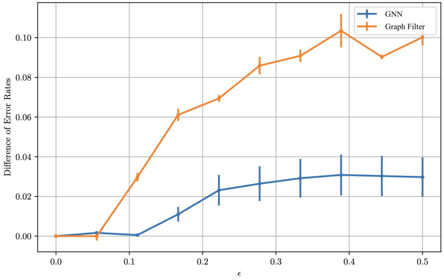

- Stability of MNNs to manifold deformations implies stability of GNNs with continuous filters

$$\left\| \mathbf{\Phi}(\mathbf{H}, \mathcal{L}, f) - \mathbf{\Phi}(\mathbf{H}, \mathcal{L}', f) \right\| = O\left[\left(\frac{N}{\alpha} + A_h + \frac{M}{\gamma} + B_h\right)\epsilon\right] \|f\|_{L^2(\mathcal{M})}$$

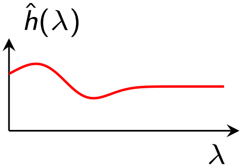

GNNs with continuous filters are stable to deformations

Z. Wang et al., *Stability to Deformations of Manifold Filters and Manifold Neural Networks,* IEEE Trans on Signal Processing

- Train GNNs for optimal resource allocation policies under system constraints in ad-hoc networks

  ⇒ GNN is trained over a family of wireless networks ⇒ Possible because of stability

  ⇒ GNN transfers to larger networks without retraining ⇒ Possible because of transferability

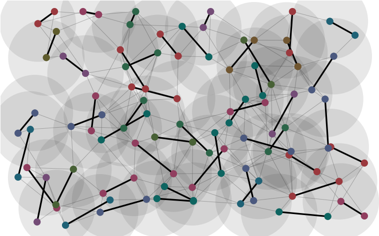

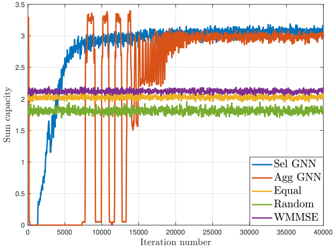

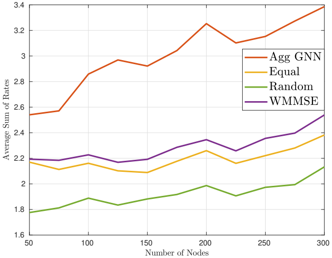

Z. Wang et al., *Learning decentralized wireless resource allocations with graph neural networks*, IEEE Trans on Signal Processing

- MNNs process scalar signals over manifolds $\Rightarrow$ vector fields arise in some applications
- We define tangent bundle convolution and further construct tangent bundle neural networks

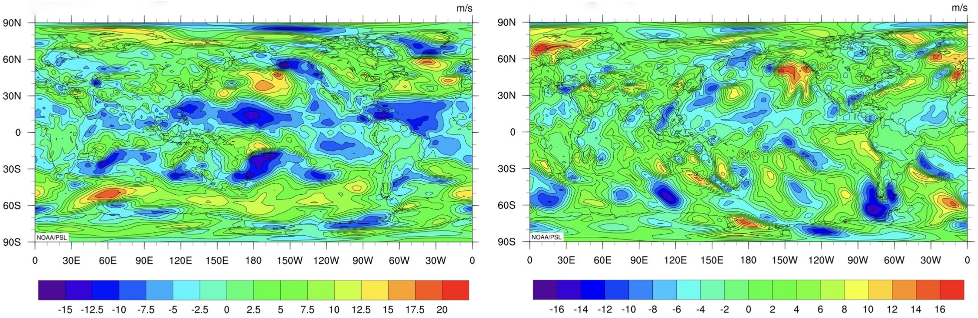

Visualization of Earth wind field

C. Battiloro, **Z. Wang**, et al., *Tangent Bundle Convolutional Learning: from Manifolds to Cellular Sheaves and Back*, IEEE Trans on Signal Processing

Graph Filters and Graph Neural Networks

15-3

- Graph $\mathbf{G}$ with matrix representation $\mathbf{S} \in \mathbb{R}^{n \times n}$ – graph shift operator – and graph signal $\mathbf{x} \in \mathbb{R}^n$

- Graph convolutional filter is defined as a summation of iterative graph data diffusions

$$\mathbf{y} = \mathbf{h}_{\mathbf{G}}(\mathbf{S})\mathbf{x} = \sum_{k=0}^{K-1} h_k \mathbf{S}^k \mathbf{x} \quad \text{– filter with coefficients } h_k$$

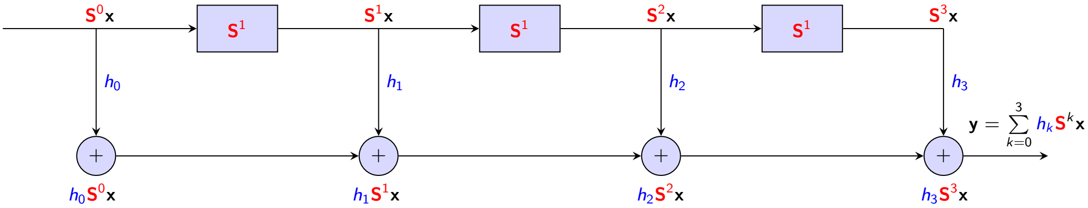

- The matrix $\mathbf{S}$ (which is symmetric) admits the eigenvector decomposition $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^H$

Spectral Representation of Graph Filters

Graph filter with coefficients h_k, graph signal $\mathbf{x}$ and the filtered signal $\mathbf{y}$

$$\mathbf{y} = \sum_{k=0}^{K-1} h_k \mathbf{S}^k \mathbf{x} = \mathbf{h}(\mathbf{S})\mathbf{x}$$

The Graph Fourier Transforms (GFTs) $\hat{\mathbf{x}} = \mathbf{V}^H \mathbf{x}$ and $\hat{\mathbf{y}} = \mathbf{V}^H \mathbf{y}$ are related by

$$\hat{\mathbf{y}} = \sum_{k=0}^{K-1} h_k \mathbf{\Lambda}^k \hat{\mathbf{x}} = \hat{h}(\mathbf{\Lambda})\hat{\mathbf{x}}$$

- The graph filter frequency response is point-wise on a scalar variable – $\hat{h}(\lambda) = \sum_{k=0}^{K-1} h_k \lambda^k$

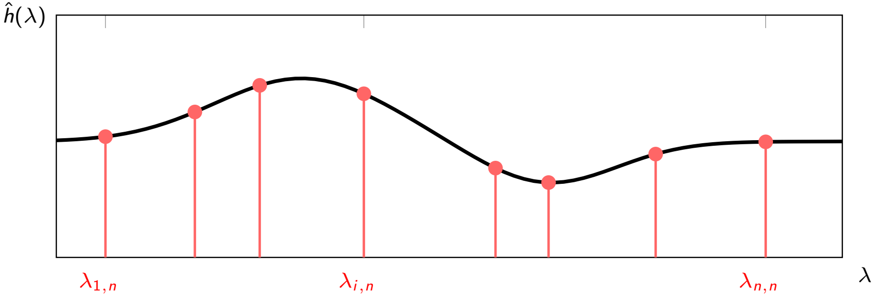

- A given graph instantiates the frequency response on its given specific eigenvalues $\lambda_{i,n}$
- Eigenvectors do not appear in the frequency response. They determine the meaning of frequencies

Graph Neural Networks (GNNs)

- Graph neural network is a cascade of L layers
- Each of the layers is composed of graph convolutions $\mathbf{h}_{\mathbf{G}}(\mathbf{S})$ and pointwise nonlinearities σ
- Define the learnable parameter set in $\mathbf{h}_{\mathbf{G}}(\mathbf{S})$ as $\mathbf{H}$
- GNN can be written as a map $\mathbf{y} = \mathbf{\Phi}_{\mathbf{G}}(\mathbf{H}, \mathbf{S}, \mathbf{x})$

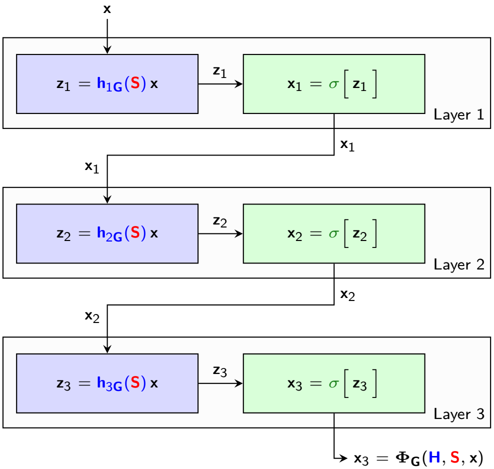

Recap:

- Graph convolutions; Spectral representation of graph filters; GNN architecture

Manifold Filters and Manifold Neural Networks

18-4

- d-dimensional manifold $\mathcal{M}$ with Laplace-Beltrami (LB) operator $\mathcal{L}$ and manifold signal f
- A Manifold filter with coefficients h is defined by the input-output relationship

$$g(x) = \int_0^\infty h(t)\, e^{-t\mathcal{L}} f(x) \mathrm{d}t = \mathbf{h}(\mathcal{L})\, f(x)$$

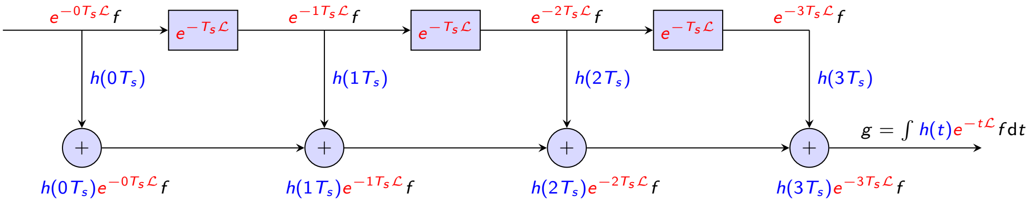

Manifold Convolutional Filters

- d-dimensional manifold $\mathcal{M}$ with Laplace-Beltrami (LB) operator $\mathcal{L}$ and manifold signal f
- A Manifold filter with coefficients h is defined by the input-output relationship

$$g(x) = \int_0^\infty h(t)\, e^{-t\mathcal{L}} f(x) \mathrm{d}t = \mathbf{h}(\mathcal{L})\, f(x)$$

- Manifold convolutions generalize graph convolutions and standard (time) convolutions

$\Rightarrow$ Discretizing a manifold filter yields a graph filter with shift operator $e^{-T_s \mathbf{L}_n}$

$$\mathbf{g} = \sum_{k=0}^{K-1} h(kT_s)\, e^{-kT_s \mathbf{L}_n}\, \mathbf{f} \approx \sum_{k=0}^{K-1} h(kT_s)\, (\mathbf{I} - T_s \mathbf{L}_n)^k\, \mathbf{f}$$

Manifold Convolutional Filters

- d-dimensional manifold $\mathcal{M}$ with Laplace-Beltrami (LB) operator $\mathcal{L}$ and manifold signal f
- A Manifold filter with coefficients h is defined by the input-output relationship

$$g(x) \;=\; \int_0^\infty h(t)\, e^{-t\mathcal{L}}\, f(x)\mathrm{d}t \;=\; \mathbf{h}(\mathcal{L})\, f(x)$$

- Manifold convolutions generalize graph convolutions and standard (time) convolutions

$\Rightarrow$ Recover standard convolutions if we make the particular choice $\mathcal{L} = \mathrm{d}/\mathrm{d}x$

$$g(x) \;=\; \int_0^\infty h(t)\, e^{-t\mathrm{d}/\mathrm{d}x} f(x)\,\mathrm{d}t \;=\; \int_0^\infty h(t)\, f(x-t)\,\mathrm{d}t$$

- $\mathcal{L}$ is self-adjoint and positive semi-definite, which leads to a discrete spectrum $\{\lambda_i, \phi_i\}_{i\in\mathbb{N}^+}$

Spectral Representation of Manifold Filters

Manifold filter with filter function $h(t)$, manifold signal $f(x)$ and the filtered signal $g(x)$

$$g(x) = \int_0^\infty h(t) e^{-t\mathcal{L}} \mathrm{d}t f(x) = \mathbf{h}(\mathcal{L}) f(x)$$

The frequency components $\hat{f}(i) = \langle f, \phi_i \rangle_{L^2(\mathcal{M})}$ and $\hat{g}(i) = \langle g, \phi_i \rangle_{L^2(\mathcal{M})}$ are related by

$$\hat{g}(i) = \int_0^\infty h(t) e^{-t\lambda_i} \mathrm{d}t \hat{f}(i) = \hat{h}(\mathbf{\Lambda}) \hat{f}(i)$$

- The manifold filter frequency response is point-wise on a scalar variable – $\hat{h}(\lambda) = \int_0^\infty h(t) e^{-t\lambda} dt$

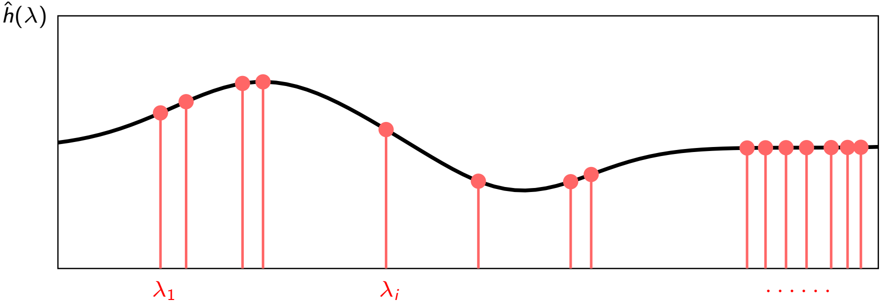

- A given manifold instantiates the frequency response on its given specific eigenvalues λ_i
- Laplace-Beltrami operator possesses infinite spectrum with $\lambda_i \propto i^{2/d}$ according to Weyl's law

- Manifold neural network is a cascade of L layers
- Each of the layers is composed of manifold convolutions $\mathbf{h}(\mathcal{L})$ and pointwise nonlinearities σ
- Define the learnable parameter set in $\mathbf{h}(\mathcal{L})$ as $\mathbf{H}$
- MNN can be written as a map $\mathbf{y} = \mathbf{\Phi}(\mathbf{H}, \mathcal{L}, f)$

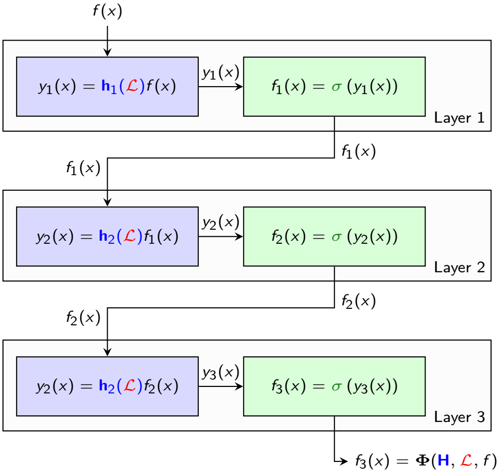

Recap:

- Graph and manifold convolutions; Spectral representation of graph and manifold filters; GNN and MNN architectures

Scalability of Graph Neural Networks

22-7

- Geometric graph filters and GNNs converge to their underlying manifold filters and MNNs

  $\Rightarrow$ Convergence enables transferability of geometric GNNs from small to large graphs

- Sample the manifold at $\{x_i\}_{i=1}^n$. Construct graph $\mathbf{G}_n$ with edge weights $w_{ij} = K_\xi \left(\frac{\|x_i - x_j\|^2}{\xi} \right)$

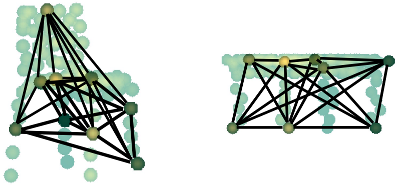

Gaussian kernel-based graphs

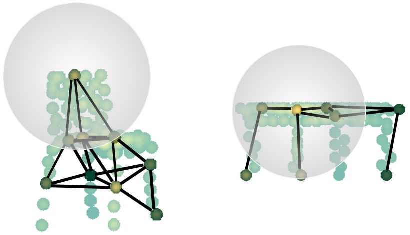

ϵ-graphs

Geometric Graph Filters and Neural Networks

- Geometric graph filter is defined by replacing Laplace-Beltrami operator with graph Laplacians $\mathbf{L}_n$

$$\mathbf{g} = \int_0^\infty h(t) e^{-t\mathbf{L}_n} \mathrm{d}t \mathbf{f} = \mathbf{h}(\mathbf{L}_n)\mathbf{f}, \qquad [\mathbf{f}]_i = f(x_i)$$

- Geometric graph neural networks on $\mathbf{G}_n$ $\Rightarrow$ cascading graph filters and non-linearities $\mathbf{\Phi}(\mathbf{H}, \mathbf{L}_n, \mathbf{f})$

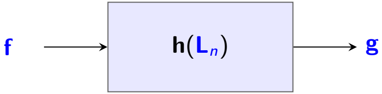

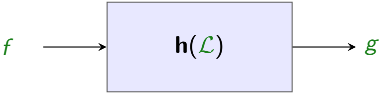

- Analyze the properties of GNNs and MNNs with the spectral structures of graphs and manifolds

- Geometric graph filter is defined by replacing Laplace-Beltrami operator with graph Laplacians $\mathbf{L}_n$

$$\mathbf{g} = \int_0^\infty h(t) e^{-t\mathbf{L}_n} \mathrm{d}t \mathbf{f} = \mathbf{h}(\mathbf{L}_n)\mathbf{f}, \qquad [\mathbf{f}]_i = f(x_i)$$

- Geometric graph neural networks on $\mathbf{G}_n$ $\Rightarrow$ cascading graph filters and non-linearities $\boldsymbol{\Phi}(\mathbf{H}, \mathbf{L}_n, \mathbf{f})$

$$\hat{\mathbf{f}} \longrightarrow \boxed{\sum_{i=0}^{n} \hat{h}(\lambda_{i,n})} \longrightarrow \hat{\mathbf{g}} \qquad\qquad \hat{f} \longrightarrow \boxed{\sum_{i=0}^{\infty} \hat{h}(\lambda_i)} \longrightarrow \hat{g}$$

- Analyze the properties of GNNs and MNNs with the spectral structures of graphs and manifolds

- A filter is A_h-Lipschitz if its frequency response function $\hat{h}(\lambda)$ is A_h-Lipschitz continuous

Definition (α-separated spectrum)

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

$$|\lambda_i - \lambda_j| > \alpha.$$

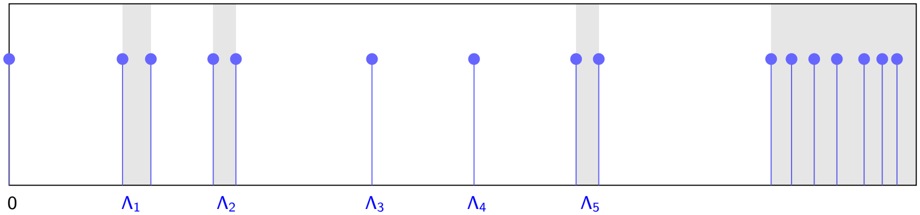

- A filter is A_h-Lipschitz if its frequency response function $\hat{h}(\lambda)$ is A_h-Lipschitz continuous

Definition (α-FDT filter)

The frequency response of α-frequency Difference threshold (α-FDT) filter $\mathbf{h}(\mathcal{L})$ satisfies

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

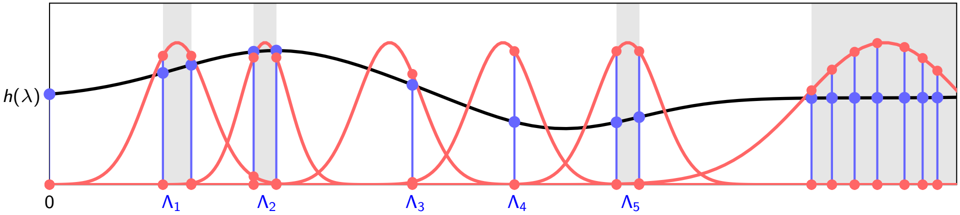

Theorem (Convergence of Geometric GNNs)

If an L-layer GNN $\mathbf{\Phi}(\mathbf{H}, \mathbf{L}_n, \cdot)$ on $\mathbf{G}_n$ and MNN $\mathbf{\Phi}(\mathbf{H}, \mathcal{L}, \cdot)$ on $\mathcal{M}$ have normalized Lipschitz nonlinearities, it holds in high probability that

$$\left\| \mathbf{\Phi}(\mathbf{H}, \mathbf{L}_n, \mathbf{P}_n f) - \mathbf{P}_n \mathbf{\Phi}(\mathbf{H}, \mathcal{L}, f) \right\|_{L^2(\mathbf{G}_n)} = O\left[\left(\frac{N}{\alpha} + A_h\right)\sqrt{\xi} + \frac{\log(n)}{n}\right] \|f\|_{L^2(\mathcal{M})}$$

with filters that are α-FDT with $\delta_D \leq O(\sqrt{\xi}/\alpha)$ and A_h-Lipschitz continuous.

- The properties of large GNNs can be analyzed via MNN $\Rightarrow$ Transferability across graph scales
- The error bound shows trade-off between convergence and discriminability $\Rightarrow$ nonlinearities lift

Z. Wang et al, *Geometric Graph Filters and Neural Networks: Limit Properties and Discriminability Trade-offs,* IEEE Trans on Signal Processing

- We evaluate the implementations of GNNs with ModelNet10 classification

Z. Wu et al, *3D shapenets: A deep representation for volumetric shapes*, IEEE CVPR 2015

- Compare the graph output differences between trained small graphs and large graphs ($n = 1000$)

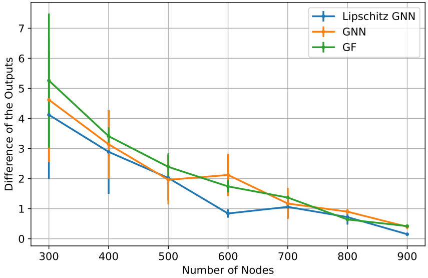

- GNNs can converge to MNNs as more points are sampled; Lipschitz GNNs have smaller differences

- We verify the transferability by testing the trained GNNs on graphs with $n = 1000$

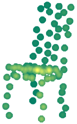

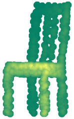

Baseline GNN		GF	GNN	Lipschitz GNN
16.95 ± 5.42	$n = 300$	21.97 ± 4.17	10.10 ± 1.40	8.60 ± 2.95
13.11 ± 4.97	$n = 500$	19.83 ± 5.94	7.74 ± 4.05	7.68 ± 3.75
10.02 ± 3.87	$n = 700$	16.62 ± 2.38	7.92 ± 3.14	8.02 ± 2.77
6.83 ± 3.96	$n = 900$	13.85 ± 3.81	7.45 ± 4.03	7.44 ± 3.30

Table: Error rates tested on $n = 1000$

- Transferability allows the GNNs trained on a small graph directly applied to a large graph

Recap:

- Graph and manifold convolutions; Spectral representation of graph and manifold filters; GNN and MNN architectures
- Transferability of GNNs across scales based on the convergence of GNNs to MNNs

Stability of GNNs Implied by MNNs

29-6

Manifold Deformations as Operator Perturbations

- We investigate the stability of MNNs to the manifold deformations

  $\Rightarrow$ Consider manifold signal f and a deformation $\tau(x) \in \mathcal{M}$ over the manifold (ϵ-small, ϵ-smooth)

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

  $\Rightarrow$ Translate manifold signal perturbations as LB operator perturbations (ϵ-small)

Theorem (Manifold deformations)

Let the deformation $\tau(x) : \mathcal{M} \to \mathcal{M}$ satisfies $\text{dist}(x, \tau(x)) \leq \epsilon$ and $J(\tau_*) = I + \Delta$ with $\|\Delta\|_F \leq \epsilon$. If the gradient field is smooth, it holds that

$$\mathcal{L} - \mathcal{L}' = \mathbf{E}\mathcal{L} + \mathcal{A},$$

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

- Manifold filters are parameterized by Laplace-Beltrami operator $\mathcal{L}$ and perturbed operator $\mathcal{L}'$
- In the spatial domain

- In the spectral domain

- Compare the difference of MNNs and perturbed MNNs with the spectral analysis of $\mathcal{L}$ and $\mathcal{L}'$

- A filter is B_h-Integral Lipschitz if its frequency response satisfies

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

Definition (γ-separated spectrum)

The γ-separated spectrum of a LB operator $\mathcal{L}$ is defined as the partition $\Lambda_1(\gamma) \cup \ldots \cup \Lambda_N(\gamma)$ such that all $\lambda_i \in \Lambda_k(\gamma)$ and $\lambda_j \in \Lambda_l(\gamma)$, $k \neq l$, satisfy

$$\left|\frac{\lambda_i}{\lambda_j} - 1\right| > \gamma.$$

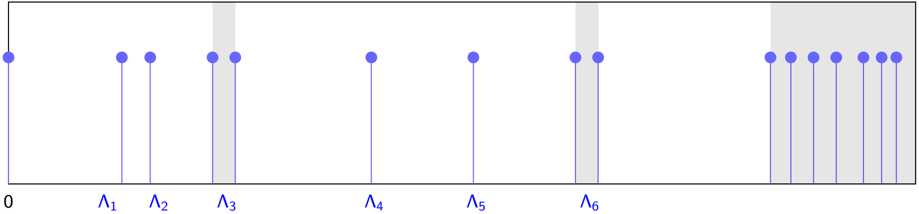

- A filter is B_h-Integral Lipschitz if its frequency response satisfies

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

Definition (γ-FRT filter)

The frequency response of γ-Frequency Ratio Threshold (γ-FRT) filter $\mathbf{h}(\mathcal{L})$ satisfies

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

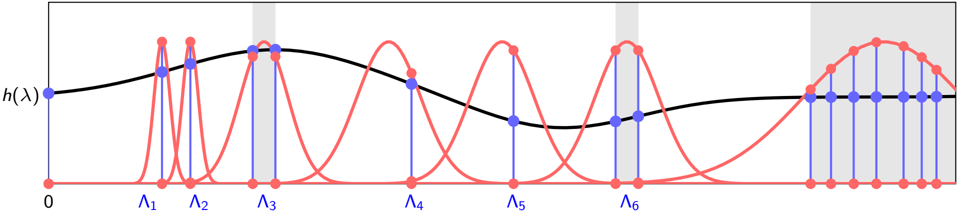

Theorem (Stability of MNNs to deformations)

An L-layer MNN $\mathbf{\Phi}(\mathbf{H}, \mathcal{L}, f)$ have normalized Lipschitz continuous nonlinearities. Let $\mathcal{L}'$ be the deformed LB operator with $\max\{\alpha, 2, |\gamma/1-\gamma|\} \gg \epsilon$, then

$$\left\| \mathbf{\Phi}(\mathbf{H}, \mathcal{L}, f) - \mathbf{\Phi}(\mathbf{H}, \mathcal{L}', f) \right\|_{L^2(\mathcal{M})} = O\left[\left(\frac{N}{\alpha} + A_h + \frac{M}{\gamma} + B_h\right)\epsilon\right] \|f\|_{L^2(\mathcal{M})}$$

if the manifold filters are α-FDT with $\delta_D \leq O(\epsilon/\alpha)$, γ-FRT with $\delta_R \leq O(\epsilon/\gamma)$, A_h-Lipschitz continuous and B_h-integral Lipschitz continuous.

- The difference bound shows a trade-off between stability and discriminability $\Rightarrow$ nonlinearities lift

Z. Wang et al., *Stability to Deformations of Manifold Filters and Manifold Neural Networks,* IEEE Trans on Signal Processing

- We verify the stability by comparing the performance on normal and deformed point clouds

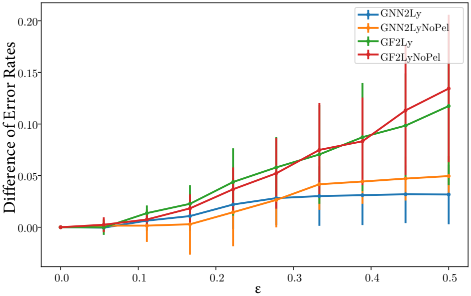

Architecture	$\epsilon = 0.2$	$\epsilon = 0.4$
GNN2Ly	$7.37\% \pm 1.43\%$	$7.71\% \pm 3.96\%$
GF2Ly	$13.76\% \pm 6.82\%$	$13.54\% \pm 7.16\%$
Architecture	$\epsilon = 0.6$	$\epsilon = 0.8$
GNN2Ly	$8.04\% \pm 2.83\%$	$11.01\% \pm 6.33\%$
GF2Ly	$14.76\% \pm 5.67\%$	$16.04\% \pm 6.34\%$

Techinical Summary

- We introduce manifold neural networks (MNNs) as the limits of graph neural networks
- And study their fundamental properties:

  $\Rightarrow$ Scalability: GNNs converge to MNNs $\Rightarrow$ the transferability of GNNs across scales

  $\Rightarrow$ Stability: MNNs are stable to deformations $\Rightarrow$ the stability of large-scale GNNs

- Manifold convolutional filters are linear combinations of Laplace-Beltrami operator exponentials $\Rightarrow$

$$g(x) = \int_0^\infty h(t)\, e^{-t\mathcal{L}}\, f(x) \mathrm{d}t$$

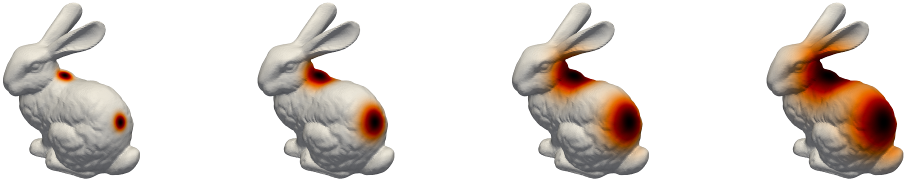

$$h(0T_s)e^{-0T_s\mathcal{L}}f \quad + \quad h(1T_s)e^{-1T_s\mathcal{L}}f \quad + \quad h(2T_s)e^{-2T_s\mathcal{L}}f \quad + \quad h(3T_s)e^{-3T_s\mathcal{L}}f$$

- Manifold neural networks (MNNs) compose layers of manifold filters and point-wise non-linearities

- GNNs converge to the underlying MNNs provided the filters satisfy spectral continuity conditions

$$\left\| \Phi(\mathbf{H}, \mathbf{L}_n, \mathbf{P}_n f) - \mathbf{P}_n \Phi(\mathbf{H}, \mathcal{L}, f) \right\| = O\left[\left(\frac{N}{\alpha} + A_h \right) \sqrt{\xi} + \frac{\log(n)}{n} \right] \|f\|_{L^2(\mathcal{M})}$$

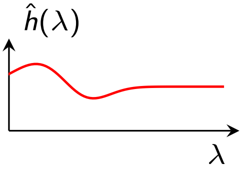

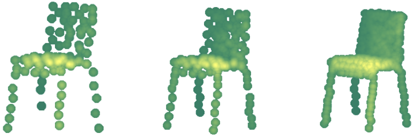

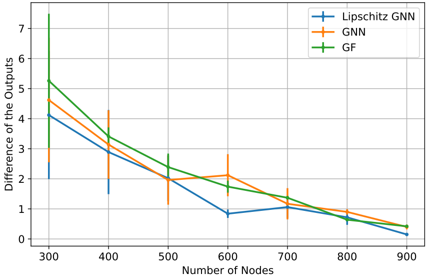

- Stability of MNNs to manifold deformations implies stability of GNNs with continuous filters

$$\left\| \Phi(\mathbf{H}, \mathcal{L}, f) - \Phi(\mathbf{H}, \mathcal{L}', f) \right\| = O\left[\left(\frac{N}{\alpha} + A_h + \frac{M}{\gamma} + B_h\right)\epsilon\right] \|f\|_{L^2(\mathcal{M})}$$

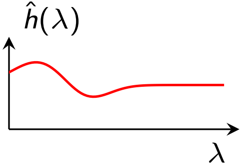

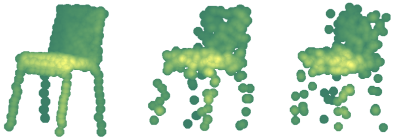

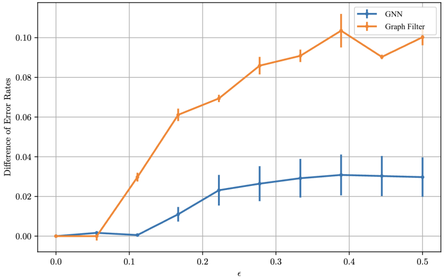

Wireless Resource Allocation

- We test the trained GNN in other ad-hoc networks of fixed size and density

  ⇒ The GNN remains optimal across permutations of ad-hoc networks

Ad-hoc network with 25 pairs

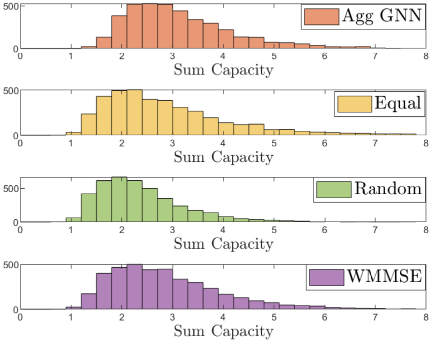

Ad-hoc network with 50 pairs

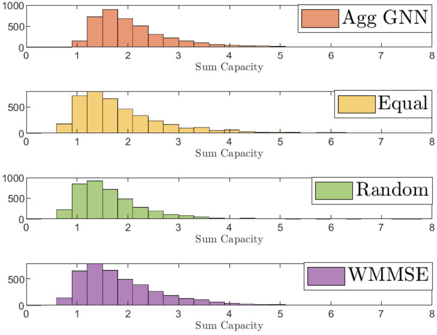

- We test in other networks of increasing size and fixed density
  - $\Rightarrow$ The GNN transfers to larger ad-hoc networks with no need of retraining

Ad-hoc network with 25 pairs

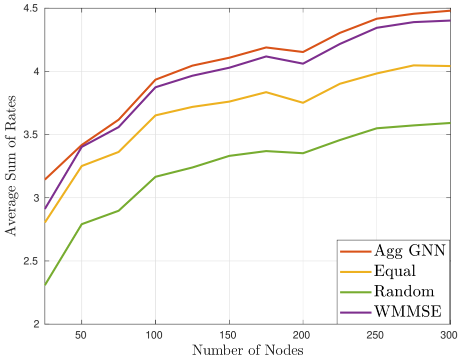

Ad-hoc network with 50 pairs

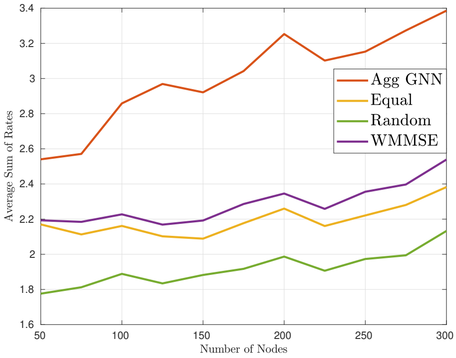

Z. Wang et al., *Learning decentralized wireless resource allocations with graph neural networks*, IEEE Trans on Signal Processing

Wind Field Reconstruction with Tangent Bundle NNs

		$E\{\tilde{n}\} = 0.5n$	$E\{\tilde{n}\} = 0.3n$	$E\{\tilde{n}\} = 0.1n$
$E\{n\} = 200$	DD-TNN	$\mathbf{1.99 \cdot 10^{-2}} \pm 2.30 \cdot 10^{-3}$	$\mathbf{1.18 \cdot 10^{-2}} \pm 1.62 \cdot 10^{-3}$	$\mathbf{3.67 \cdot 10^{-3}} \pm 1.23 \cdot 10^{-3}$
	MNN	$3.19 \cdot 10^{-2} \pm 1.31 \cdot 10^{-2}$	$2.74 \cdot 10^{-2} \pm 1.55 \cdot 10^{-2}$	$2.58 \cdot 10^{-2} \pm 1.82 \cdot 10^{-2}$
	MLP	$2.03 \cdot 10^{-2} \pm 2.28 \cdot 10^{-3}$	$1.20 \cdot 10^{-2} \pm 1.68 \cdot 10^{-3}$	$3.69 \cdot 10^{-3} \pm 1.17 \cdot 10^{-3}$
$E\{n\} = 300$	DD-TNN	$\mathbf{1.88 \cdot 10^{-2}} \pm 1.72 \cdot 10^{-3}$	$\mathbf{1.13 \cdot 10^{-2}} \pm 1.54 \cdot 10^{-3}$	$\mathbf{3.96 \cdot 10^{-3}} \pm 1.00 \cdot 10^{-3}$
	MNN	$2.68 \cdot 10^{-2} \pm 7.64 \cdot 10^{-3}$	$2.41 \cdot 10^{-2} \pm 1.05 \cdot 10^{-2}$	$2.09 \cdot 10^{-2} \pm 1.76 \cdot 10^{-2}$
	MLP	$1.95 \cdot 10^{-2} \pm 1.74 \cdot 10^{-3}$	$1.18 \cdot 10^{-2} \pm 1.56 \cdot 10^{-3}$	$4.00 \cdot 10^{-3} \pm 8.85 \cdot 10^{-4}$
$E\{n\} = 400$	DD-TNN	$\mathbf{1.95 \cdot 10^{-2}} \pm 1.66 \cdot 10^{-3}$	$\mathbf{1.14 \cdot 10^{-2}} \pm 1.38 \cdot 10^{-3}$	$\mathbf{3.70 \cdot 10^{-3}} \pm 8.55 \cdot 10^{-4}$
	MNN	$2.48 \cdot 10^{-2} \pm 6.55 \cdot 10^{-3}$	$2.52 \cdot 10^{-2} \pm 1.20 \cdot 10^{-2}$	$8.16 \cdot 10^{-2} \pm 1.87 \cdot 10^{-1}$
	MLP	$2.01 \cdot 10^{-2} \pm 1.66 \cdot 10^{-3}$	$1.19 \cdot 10^{-2} \pm 1.24 \cdot 10^{-3}$	$3.81 \cdot 10^{-3} \pm 8.46 \cdot 10^{-4}$

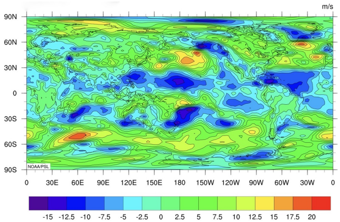

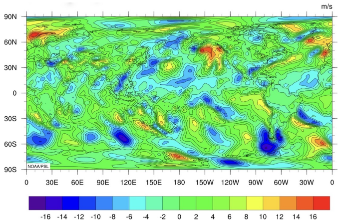

C. Battiloro, **Z. Wang**. et al., *Tangent bundle convolutional learning: from manifolds to cellular sheaves and back*, IEEE Trans on Signal Processing

- We introduce manifold neural networks (MNNs) as the limits of graph neural networks

- And study their fundamental properties:

  ⇒ Scalability: GNNs converge to MNNs ⇒ the transferability of GNNs across scales

  ⇒ Stability: MNNs are stable to deformations ⇒ the stability of large-scale GNNs

- Informs the practical design of graph neural networks for large-scale geometric graphs

  ⇒ Point-cloud analysis, Wireless communications, Wind field reconstructions etc.