Graph neural networks as dynamical systems

Francesco Di Giovanni

Twitter

First Italian School in GDL: July 25–28, Pescara

Presentation outline

- Graph preliminaries
- Spectral analysis and Dirichlet energy on graphs
- Dynamical systems on graphs
- MPNNs as multi-particle systems and the gradient flow framework (GRAFF)
- Presentation of *Graph Neural Networks as Gradient Flows*

Introduction

Preliminaries on graph operators

- $G = (V, E)$ is an *undirected* graph with $|V| = n$ and $i \quad j$ if $(i, j) \quad E$
- A, D are $n \times n$ adjacency and (diagonal) degree matrices
- The *normalized* adjacency is $\bar{A} = D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$
- The Laplacian $\Delta = I - \bar{A}$ is an operator acting on signals $f : V \to R$ as

$$(\Delta f)_i = f_i - \sum_{j \quad i} \frac{f_j}{\sqrt{d_i d_j}}$$

Preliminaries on graph operators

- $G = (V, E)$ is an *undirected* graph with $|V| = n$ and $i \quad j$ if $(i, j) \quad E$
- A, D are $n \times n$ adjacency and (diagonal) degree matrices
- The *normalized* adjacency is $\bar{A} = D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$
- The Laplacian $\Delta = I - \bar{A}$ is an operator acting on signals $f : V \to R$ as

$$(\Delta f)_i = f_i - \sum_{j \quad i} \frac{f_j}{\sqrt{d_i d_j}}$$

The Laplacian $\Delta \quad 0 \to$ eigenvalues satisfy $0 = \lambda_0^{\Delta} \leq \ldots \leq \lambda_{n-2}^{\Delta} \leq \rho_{\Delta}$, with $\rho_{\Delta} \leq 2$, and are called (graph) *frequencies*, eigenvectors are denoted by $\{\varphi^{\Delta}\}_{=0}^{n-1}$

Signal on graphs: Dirichlet energy and smoothness

Consider a signal (feature) $\mathbf{f} : V \to \mathbb{R}$ e.g. temperature of each node

We write $\mathbf{f} = (f_1, \ldots, f_n) \quad \to \mathbf{f} = \sum_{i=1}^{n} c_i \varphi_i$

Δ can be used to measure smoothness of $\mathbf{f}$: the Dirichlet energy[1] E^{Dir} is defined by

$$E^{\text{Dir}}(\mathbf{f}) := \frac{1}{4} \sum_{i \sim j} \left\| \frac{f_i}{\sqrt{d_i}} - \frac{f_j}{\sqrt{d_j}} \right\|^2 = \frac{1}{2} \langle \mathbf{f}, \Delta \mathbf{f} \rangle = \frac{1}{2} \sum_{i=1}^{n} \lambda_i c_i^2.$$

[1] Zhou and Schölkopf (2005)

Signal on graphs: Dirichlet energy and smoothness

Consider a signal (feature) $\mathbf{f} : V \to \mathbb{R}$ e.g. temperature of each node

We write $\mathbf{f} = (f_1, \ldots, f_n) \quad \to \mathbf{f} = \sum_i c_i \varphi_i$

Δ can be used to measure smoothness of $\mathbf{f}$: the Dirichlet energy[1] E^{Dir} is defined by

$$E^{\text{Dir}}(\mathbf{f}) := \frac{1}{4} \sum_{i \sim j} \left\| \frac{f_i}{\sqrt{d_i}} - \frac{f_j}{\sqrt{d_j}} \right\|^2 = \frac{1}{2} \langle \mathbf{f}, \Delta \mathbf{f} \rangle = \frac{1}{2} \sum_i \lambda_i c_i^2.$$

$\to$ the frequency components of $\mathbf{f}$ determine the variations of the signal along edges

The quantity $f_i / \sqrt{d_i} - f_j / \sqrt{d_j} := \nabla \mathbf{f}(i, j)$ is the **gradient** of $\mathbf{f}$ along $(i, j) \in E$

[1] Zhou and Schölkopf (2005)

A rough picture: low-pass vs high-pass filtering

Consider a dynamical process $t \to \mathsf{f}(t) \quad \mathsf{R}^n$ starting at $\mathsf{f}_0 \to \mathsf{f}(t) = \square \; c\,(t)\varphi^{\Delta}$

A rough picture: low-pass vs high-pass filtering

Consider a dynamical process $t \to f(t) \quad R^n$ starting at $f_0 \to f(t) = \square \; c\,(t)\varphi^{\Delta}$

If the high-frequency components $|c\,(t)|$, with $>> 0$, decrease with time, then the process acts as '**low-pass** filtering' → smooths the signal out

A rough picture: low-pass vs high-pass filtering

Consider a dynamical process $t \to f(t) \quad \mathbb{R}^n$ starting at $f_0 \to f(t) = \square \; c\,(t)\varphi^{\Delta}$

If the high-frequency components $|c\,(t)|$, with $>> 0$, decrease with time, then the process acts as '**low-pass** filtering' → smooths the signal out

If the low-frequency components $|c\,(t)|$, with 0, decrease with time, then the process acts as '**high-pass** filtering' → sharpens the signal

Figure 1: First four Laplacian eigenvectors of Minnesota Road graph. Figure taken from Bronstein et al. (2017)

A prototypical low-pass filtering: the graph heat equation

Consider an input signal $f_0 : V \to \mathbb{R}$ and recall that $f \to E^{\mathrm{Dir}}(f) = \frac{1}{2}\langle f, \Delta f \rangle$

If we want to *minimize* E^{Dir} $\to$ take infinitesimal steps in the direction of steepest descent

$$\text{Heat equation}: \quad \dot{f}(t) = -\nabla_f E^{\mathrm{Dir}}(f(t)) = -\Delta f(t), \quad f(0) = f_0.$$

This is a **gradient flow**; $\dot{E}^{\mathrm{Dir}}(f(t)) \leq 0$ and $f(t) \to f_\infty$ s.t. $\Delta f_\infty = 0$ i.e. $f_\infty \in \mathrm{span}(\sqrt{d_1}, \ldots, \sqrt{d_n})$

Low-pass dynamics $\to$ 'features become indistinguishable' when $t >> 1$

Multiple channels

Consider $\mathsf{F} : \mathsf{V} \to \mathsf{R}^d$ with matrix representation F $\mathsf{R}^{n\times d} \to E^{\mathrm{Dir}}$ can be extended as

$$E^{\mathrm{Dir}}(\mathsf{F}) = \frac{1}{4} \sum_{(i,j)\ \mathsf{E}} \left\| \frac{\mathsf{f}_i}{\sqrt{d_i}} - \frac{\mathsf{f}_j}{\sqrt{d_j}} \right\|^2 = \frac{1}{2}\mathrm{trace}(\mathsf{F}\ \ \Delta\mathsf{F})$$

The gradient flow of E^{Dir} yields heat equation in each feature channel[2]:

$$\dot{\mathsf{f}}^r(t) = -\Delta \mathsf{f}^r(t), \quad 1 \le r \le d$$

[2] 'Channels' = 'feature components' = 'feature coordinates'

The formalism

We can vectorize a matrix signal $F \in \mathbb{R}^{n \times d}$! $\mathrm{vec}(F) \in \mathbb{R}^{nd}$

We use the Kronecker product $\otimes_d : \mathbb{R}^{nd} \times \mathbb{R}^{nd} \to \mathbb{R}^{nd}$ to rewrite E^{Dir} as

$$E^{\mathrm{Dir}}(F) = \frac{1}{2}\langle \mathrm{vec}(F), (L \otimes_d I_d)\mathrm{vec}(F)\rangle$$

The heat equation can also be rewritten by `stacking the columns as'

$$\mathrm{vec}(\dot{F}(t)) = (L \otimes_d I_d)\mathrm{vec}(F(t))$$

Upshot: formalism reduces a matrix ODE to a vector ODE ! vectorized ODEs are much easier to deal with

A motivating example

How to determine if a dynamical process on a graph is dominated by the low or high frequencies?

A motivating example

How to determine if a dynamical process on a graph is dominated by the low or high frequencies? Use E^{Dir} after normalization

A motivating example

How to determine if a dynamical process on a graph is dominated by the low or high frequencies? Use E^{Dir} after normalization

Consider $\dot{F}(t) = AF(t)$ () $\quad \text{vec}(\dot{F}(t)) = (I_d \quad A)\text{vec}(F(t))$, with $F(0) = F_0$

Recall that $A = I$ so we can solve as

$$f^r(t) = e^{At} f^r(0) = e^{(I \quad)t} f^r(0); \quad 1 \quad r \quad d$$

A motivating example

How to determine if a dynamical process on a graph is dominated by the low or high frequencies? Use E^{Dir} after normalization

Consider $\dot{F}(t) = AF(t)$ () $\text{vec}(\dot{F}(t)) = (I_d \quad A)\text{vec}(F(t))$, with $F(0) = F_0$

Recall that $A = I$ so we can solve as

$$f^r(t) = e^{At} f^r(0) = e^{(I \quad)t} f^r(0); \quad 1 \quad r \quad d$$

Expand each channel in the basis g satisfying $A\ 、 = (1 \quad 、)\ 、$:

$$f^r(t) = \sum_{、} e^{(1 \quad 、)t} h f^r(0); \quad 、 \ i \ 、$$

A motivating example

Recall that $_0$ is the smoothest eigenvector i.e. $_0 = 0$

The projection along $_0$ is the one growing the fastest[8] since

$$\langle f^r(t); \ _0 \rangle = e^{(1 \ \ 0)t} \langle f^r(0); \ _0 \rangle$$

The dynamics are `dominated' by the low-frequencies: does $E_{Dir}(F(t)) \ ! \ \ 0$?

[3] Unless $\langle f^r(0); \ _0 \rangle_{ij} = 0$ which is only true in a smaller subspace of $\mathbb{R}^{pq}$

[4] Unless $\langle f^r(0); \ _i \rangle = 0$ for all $` > 0$

A motivating example

Recall that $_0$ is the smoothest eigenvector i.e. $_0 = 0$

The projection along $_0$ is the one growing the fastest[8] since

$$\langle f^r(t); \ _0 \rangle = e^{(1 \ 0)t} \langle f^r(0); \ _0 \rangle$$

The dynamics are `dominated' by the low-frequencies: does $E^{Dir}(F(t)) \ ! \ 0$? No:[4]

$$E^{Dir}(f^r(t)) = \frac{1}{2}\langle f^r(t); f^r(t) \rangle = \sum_{` > 0} e^{(1 \ `)t} (\langle f^r(0); \ ` \rangle)^2 \ ! \ 1$$

[3] Unless $\langle f^r(0); \ _0 \rangle_{ij} = 0$ which is only true in a smaller subspace of $\mathbb{R}^n$

[4] Unless $\langle f^r(0); \ ` \rangle = 0$ for all $` > 0$

A motivating example

Looking at E^{Dir} is not enough! we should normalize rst: in fact we have

$$E^{Dir}(F(t)=jjF(t)jj) ! 0; t ! 1$$

A motivating example

Looking at E^{Dir} is not enough! we should normalize rst: in fact we have

$$E^{Dir}(F(t)=jjF(t)jj) \; ! \quad 0; \quad t \; ! \; 1$$

and for each channel r d 9 f_1^r s.t.

$$f^r(t)=jjf^r(t)jj \; ! \quad f_1^r \; ; \quad f \; {}_1^r \; = 0$$

A motivating example

Looking at E^{Dir} is not enough! we should normalize rst: in fact we have

$$E^{Dir}(F(t)=\|F(t)\|) \to 0, \quad t \to 1$$

and for each channel r d 9 f_1^r s.t.

$$f^r(t)=\|f^r(t)\| \to f_1^r, \quad f_1^r = 0$$

Upshot: Analyse $F(t)$ via $E^{Dir}(F(t)=\|F(t)\|) \to$ Rayleigh quotient of $_d$

Low-frequency-dominant: LFD

De nition

A dynamical system $\dot{F}(t)$ initialized at $F(0)$ is Low-Frequency-Dominant (LFD) if $E^{Dir}(F(t) = \|F(t)\|) \to 0$ for $t \to \infty$.

Low-frequency-dominant: LFD

De nition

A dynamical system $F(t)$ initialized at $F(0)$ is Low-Frequency-Dominant LFD if $E^{Dir}(F(t)=\|F(t)\|) \to 0$ for $t \to 1$.

Does it make sense?

Lemma

A dynamical system is LFD iff for each sequence $t_j \to 1$ there exist a subsequence $t_{j_k} \to 1$ and F_1 s.t. $F(t_{j_k})=\|F(t_{j_k})\| \to F_1$ and $f_1^r = 0$.

LFD dynamics: numerical example

A numerical example of LFD dynamics: $T = 4.0, \quad = 0.1$

$$\dot{F}(t) = AF(t) \quad ; \qquad = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

LFD dynamics: numerical example

A numerical example of LFD dynamics: $T = 4.0, \quad = 0.1$

$$\dot{F}(t) = AF(t) \quad ; \qquad = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

In both cases the eigenvector f_0 dominates the dynamics

- Top: solution becomes unbounded
- Bottom: evolution of $F(t) = \|F(t)\|$
  ! convergence to $\ker(\ )$ where
  we only distinguish nodes based on their degree

High-frequency-dominant: HFD

Note that $E^{Dir}(F) \quad \frac{1}{2} \quad jjFjj^2 ! E \quad ^{Dir}(F = jjFjj) \quad \frac{1}{2}$

High-frequency-dominant: HFD

Note that $E^{Dir}(F) \quad \frac{1}{2} \quad jjFjj^2 ! E^{Dir}(F=jjFjj) \quad \frac{1}{2}$

De nition

A dynamical system $F(t)$ initialized at $F(0)$ is High-Frequency-Dominant (HFD) if $E^{Dir}(F(t)=jjF(t)jj) ! \quad =2$ for $t ! 1$.

High-frequency-dominant: HFD

Note that $E^{Dir}(F) \geq \frac{1}{2} \|F\|^2 \geq E^{Dir}(F) - \|F\| \geq \frac{1}{2}$

Definition

A dynamical system $F(t)$ initialized at $F(0)$ is High-Frequency-Dominant (HFD) if $E^{Dir}(F(t) - \|F(t)\|) / E^{Dir}(F(t)) \to 0$ for $t \to \infty$.

Does it make sense?

Lemma

A dynamical system is HFD iff for each sequence $t_j \to \infty$ there exist a subsequence $t_{j_k} \to \infty$ and F_∞ s.t. $F(t_{j_k}) - \|F(t_{j_k})\| \to F_\infty$ and $f_1^\infty = f_1^\infty$.

Why do we need HFD?

Consider $\dot{F}(t) = AF(t)$! eigenvector dominates the dynamics

- Evolution of $F(t) = \frac{F(t)}{||F(t)||}$! convergence to $\ker(L)$ where we distinguish nodes based on the largest frequency eigenvector (right figure)

Homophily vs heterophily aka short vs long range interactions

Semi-supervised setting V_{tr} V labelled! predict labels on V_{test}

Homophily: Neighbours often share labels labels are smooth i.e. low-pass is `good'

Heterophily: 1 homophily! labels are not smooth i.e. low-pass is `bad'

Homophily vs heterophily aka short vs long range interactions

Semi-supervised setting V_{tr} V labelled! predict labels on V_{test}

Homophily: Neighbours often share labels labels are smooth i.e. low-pass is `good'

Heterophily: 1 homophily! labels are not smooth i.e. low-pass is `bad'

Dual perspective: short-range relations vs long-range relations relevant for graph classi cation and regression tasks on molecules

Graph Convolutional Networks

A layer of Graph Convolutional Network (GCN)[5] is de ned by:

$$F(t+1) = \mathrm{ReLU}\left(AF(t)W(t) \right)$$

A is the message-passing matrix and W(t) is the `channel-mixing'

[5] Kipf and Welling (2017)

[6] Nt and Maehara (2019); Oono and Suzuki (2020); Cai and Wang (2020)

Graph Convolutional Networks

A layer of Graph Convolutional Network (GCN)[5] is de ned by:

$$F(t+1) = \mathrm{ReLU}\quad AF(t)W(t)$$

A is the message-passing matrix and W(t) is the `channel-mixing'

- Poor performance on heterophilic graphs
- Degradation when increasing depth (over-smoothing)[6]

[5] Kipf and Welling (2017)

[6] Nt and Maehara (2019); Oono and Suzuki (2020); Cai and Wang (2020)

Low-pass lters and over-smoothing: review

Theorem (Cai and Wang)

Let $(1 \quad)^2 := \max \ (1 \quad)^2$ and $s_T = \max_{t \quad T} \text{sing}(W(t))$. Then the solution $F(T)$ of GCN satis es

$$E^{Dir}(F(T)) \quad (s_T(1 \quad))^{2T} E^{Dir}(F(0)):$$

Low-pass lters and over-smoothing: review

Theorem (Cai and Wang)

Let $(1 \quad)^2 := \max \ (1 \quad)^2$ and $s_T = \max_{t \ T} \text{sing}(W(t))$. Then the solution $F(T)$ of GCN satis es

$$E^{Dir}(F(T)) \quad (s_T(1 \quad))^{2T} E^{Dir}(F(0)):$$

- If singular values of $W(t)$ are controlled in terms of the spectrum of ! solution of GCN becomes increasingly smoother

Low-pass lters and over-smoothing: review

Theorem (Cai and Wang)

Let $(1 \quad)^2 := \max_{\ \ \ } (1 \quad \ \ )^2$ and $s_T = \max_{t \ \ T} \text{sing}(W(t))$. Then the solution $F(T)$ of GCN satis es

$$E^{Dir}(F(T)) \quad (s_T(1 \quad))^{2T} E^{Dir}(F(0)):$$

- If singular values of $W(t)$ are controlled in terms of the spectrum of ! solution of GCN becomes increasingly smoother
- GCN should succeed with homophily but fail with heterophily

Low-pass lters and over-smoothing: review

Theorem (Cai and Wang)

Let $(1 \quad)^2 := \max_{\ } (1 \quad \ )^2$ and $s_T = \max_{t \ T} \text{sing}(W(t))$. Then the solution $F(T)$ of GCN satis es

$$E^{Dir}(F(T)) \quad (s_T(1 \quad))^{2T} E^{Dir}(F(0)):$$

- If singular values of $W(t)$ are controlled in terms of the spectrum of! solution of GCN becomes increasingly smoother
- GCN should succeed with homophily but fail with heterophily
- If $T >> 1$, we converge to ker( ) i.e. only information to separate nodes is degree

Are graph convolutional models doomed?

- What if the singular values of W are not bounded by $(1)^2$?

Are graph convolutional models doomed?

- What if the singular values of W are not bounded by $(1)^2$?
- Can we require more structure on W?

Are graph convolutional models doomed?

- What if the singular values of W are not bounded by $(1)^2$?
- Can we require more structure on W?
- What is the interpretation of W?

Are graph convolutional models doomed?

- What if the singular values of W are not bounded by $(1+\epsilon)^2$?
- Can we require more structure on W?
- What is the interpretation of W?
- What is the `minimal requirement' for a graph convolutional framework to be UFED?

A physics-inspired approach

Figure 2: Actual GRAFF dynamics: attractive and repulsive forces lead to a non-smoothing process able to separate labels

Joint w/ J. Rowbottom, B. Chamberlain, T. Markovich, M. Bronstein (2022)

Outline of the contributions

- We propose a gradient flow framework (GRAFF) for MPNNs where the equations follow the direction of steepest descent of a learnable energy

Outline of the contributions

- We propose a gradient flow framework (GRAFF) for MPNNs where the equations follow the direction of steepest descent of a learnable energy
- We show how the channel-mixing W can learn to induce either LFD or HFD dynamics via its spectrum

Outline of the contributions

- We propose a gradient flow framework (GRAFF) for MPNNs where the equations follow the direction of steepest descent of a learnable energy
- We show how the channel-mixing W can learn to induce either LFD or HFD dynamics via its spectrum
- This allows us to interpret MPNNs as multi-particle dynamics with attractive and repulsive forces generated by positive and negative eigenvalues of W

Outline of the contributions

- We propose a gradient flow framework (GRAFF) for MPNNs where the equations follow the direction of steepest descent of a learnable energy
- We show how the channel-mixing W can learn to induce either LFD or HFD dynamics via its spectrum
- This allows us to interpret MPNNs as multi-particle dynamics with attractive and repulsive forces generated by positive and negative eigenvalues of W
- Show that LFD = HFD dynamics induced by this framework adapt to the underlying homophily/heterophily

Residual networks as discrete ODEs

A ResNet $F(t+\ ) = F(t) + \ \ \text{ResNet}(F(t))$ is the Euler discretization of an ODE[7] (as the step-size $! \ \ 0$)

$$E(t) = \text{ResNet}(F(t))$$

ODE theory $!$ analysing and improving ResNets

Figure 3: Dynamics of ResNet vs ODE. Figure taken from Chen et al. (2018)

[7] Haber and Ruthotto (2018); Chen et al. (2018)

Residual networks as discrete ODEs

A ResNet $F(t+\quad) = F(t) + \quad \text{ResNet}(F(t))$ is the Euler discretization of an ODE[7] (as the step-size $! \quad 0$)

$$E(t) = \text{ResNet}(F(t))$$

ODE theory $!$ analysing and improving ResNets

Figure 3: Dynamics of ResNet vs ODE. Figure taken from Chen et al. (2018)

What about residual MPNNs?

$$F(t+\quad) = F(t) + \quad \text{MPNN}(G; F(t)) \; ! \quad E(t) = \text{MPNN}(\; G; F(t))$$

[7] Haber and Ruthotto (2018); Chen et al. (2018)

Instances of `continuous' MPNNs

The linear GCN[8] system

$$F(t+1) = AF(t)W(t) \quad E(t) = AF(t)W(t) \quad F(t)$$

[8] Wu et al. (2019)

Instances of `continuous' MPNNs

The linear GCN[8] system

$$F(t+1) = AF(t)W(t) \quad ! \quad \dot{F}(t) = AF(t)W(t) \quad F(t)$$

If we use the -formalism: GCN is the unit step-size discretization of

$$\mathrm{vec}(\dot{F}(t)) = (W(t)^{>} \quad A \quad I)\mathrm{vec}(F(t))$$

! we'll see that the dampening term I is responsible for LFD dynamics

[8] Wu et al. (2019)

Instances of `continuous' MPNNs

Continuous Graph Neural Network (CGNN)[9] set $W = W^{\top}$!

$$E(t) = F(t) + F(t)W + F(0)$$

[9] Xhonneux et al. (2020)

Instances of `continuous' MPNNs

Continuous Graph Neural Network (CGNN)[9] set $W = W^{>}$!

$$\dot{F}(t) = F(t) + F(t)W + F(0)$$

- CGNN is a gradient flow
- We'll prove that this is never HFD
- Source term $F(0)$ increases expressive power

[9] Xhonneux et al. (2020)

Instances of `continuous' MPNNs

Graph Neural Diffusion (GRAND)[10] is the `continuous' version of GAT[11]

$$\dot{F}(t) = (I - A(F(t)))F(t)$$

- $A(F(t))$ is an attention matrix over the edge set
- (Linear) GRAND is a diffusion process with maximum principle, low-pass filter and over-smoothing

[10] Chamberlain et al. (2021)
[11] Veličković et al. (2018)

Instances of `continuous' MPNNs

PDE-GCN$_D$ [12] is a diffusion process given by

$$\dot{F}(t) = F(t)W(t)^{\top} W(t)$$

! We'll prove that this is a smoothing process and hence not suitable for heterophilic graphs

[12] Eliasof et al. (2021)

Instances of `continuous' MPNNs

Second-order variants[13] ! by design they prevent over-smoothing

$$\dot{F}(t) = \mathrm{MPNN}(G; F(t)) \quad F(t) \quad E(t)$$

However, why oscillatory behaviour? Do we need them?

[13] Eliasof et al. (2021); Rusch et al. (2022)

Instances of `continuous' MPNNs

Second-order variants[13] ! by design they prevent over-smoothing

$$\dot{F}(t) = \mathrm{MPNN}(G; F(t)) \quad F(t) \quad E(t)$$

However, why oscillatory behaviour? Do we need them?

The actual equations are parametric how to choose them ?

[13] Eliasof et al. (2021); Rusch et al. (2022)

Instances of `continuous' MPNNs

Second-order variants[13] ! by design they prevent over-smoothing

$$\dot{F}(t) = \mathrm{MPNN}(G; F(t)) \quad F(t) \quad E(t)$$

However, why oscillatory behaviour? Do we need them?

The actual equations are parametric how to choose them ?

Upshot: Learn an energy rather than the equations!

[13] Eliasof et al. (2021); Rusch et al. (2022)

Dynamical systems as gradient flows

Dynamical systems are gradient flows where $E : \mathbb{R}^N \to \mathbb{R}$:

$$\dot{F}(t) = \text{ODE}(F(t)) = -\nabla_F E(F(t)) = -\frac{\partial E}{\partial F}(F(t))$$

Gradient flows are easier to analyze and interpret since the solution $F(t)$ is minimizing E

What if we parametrize an energy rather than the MPNN equations?

Dynamical systems as gradient flows

Dynamical systems are gradient flows when $\mathrm{ODE} : \mathbb{R}^N \to \mathbb{R}$:

$$\dot{F}(t) = \mathrm{ODE}(F(t)) = -\nabla_F E(F(t)) = -\frac{dE}{dF}(F(t)) \leq 0.$$

Gradient flows are easier to analyze and interpret since the solution $F(t)$ is minimizing E

What if we parametrize an energy rather than the MPNN equations?

Goal: Learn E generalizing E^{Dir} → find right notion of smoothness for the problem

$$\dot{F}(t) = \mathrm{MPNN}(G; F(t)) = -\nabla_F E(G; F(t))$$

GNNs as Gradient Flows part 1: taking inspiration from harmonic maps

Harmonic map ow in continuous space

f : R^n ! (R^d; h) smooth with h a constant metric The Dirichlet energy of f is

$$E(f; h) = \frac{1}{2} \int_{R^n} \| r f \|_h^2 \, dx = \frac{1}{2} \sum_{q,r=1}^{d} \sum_{j=1}^{n} \int_{R^n} h_{qr} \, @f^q @f^r (x) dx$$

! measures the smoothness of f wrt h

[14] Kimmel et al. (1997); Perona and Malik (1990)

Harmonic map ow in continuous space

f : R^n ! (R^d; h) smooth with h a constant metric The Dirichlet energy of f is

$$E(f; h) = \frac{1}{2}\int_{R^n} \mathrm{kr}\ f\ k_h^2\, dx = \frac{1}{2}\sum_{q,r=1}^{d}\sum_{j=1}^{n}\int_{R^n} h_{qr}\, @_j f^{\,q}\, @_j f^{\,r}\,(x)\, dx$$

! measures the smoothness of f wrt h

Eells and Sampson (1964) studied the gradient ow of E given by $f_t(t) = r_f E(f(t))$ to nd minimizers of E ! extended to manifolds harmonic map ow

For PDE-based image processing gradient ows of E recover the Perona-Malik diffusion [14]

[14] Kimmel et al. (1997); Perona and Malik (1990)

Extending the formalism to graphs

$f : R^n \to (R^d, h)$ smooth with h a constant metric. The Dirichlet energy of f is

$$E(f; h) = \frac{1}{2} \int_{R^n} \|df\|_h^2 \, dx = \frac{1}{2} \sum_{q,r=1}^{d} \sum_{j=1}^{n} \int_{R^n} h_{qr} \partial_j f^q \partial_j f^r (x) dx$$

Extending the formalism to graphs

$f : R^n \to (R^d; h)$ smooth with h a constant metric. The Dirichlet energy of f is

$$E(f; h) = \frac{1}{2} \int_{R^n} \|r f\|_h^2 \, dx = \frac{1}{2} \sum_{q,r=1}^{d} \sum_{j=1}^{n} \int_{R^n} h_{qr} \partial_j f^q \partial_j f^r (x) dx$$

Replace $\int_{R^n}$ with $\sum_{i \in V}$ and ∂_{j_i} with $r_{(i,j) \in E}$:

$$E_W^{\mathrm{Dir}}(F) := \frac{1}{4} \sum_{q,r=1}^{d} \sum_{i \in V} \sum_{j:(i,j) \in E} h_{qr} (r f^q)_{ij} (r f^r)_{ij} = \frac{1}{4} \sum_{(i,j) \in E} \|W (r F)_{ij}\|^2 :$$

with $H = W^{>} W$ with $W \in R^{d \times d}$

Extending the formalism to graphs

$f : R^n \to (R^d; h)$ smooth with h a constant metric. The Dirichlet energy of f is

$$E(f; h) = \frac{1}{2} \int_{R^n} \|r f\|_h^2 \, dx = \frac{1}{2} \sum_{q,r=1}^{d} \sum_{j=1}^{n} \int_{R^n} h_{qr} \partial_j f^q \partial_j f^r (x) dx$$

! Replace $\int_{R^n}$ with $\sum_{i \in V}$ and $\partial_{j i}$ with $\nabla_{(i,j) \in E}$:

$$E_W^{Dir}(F) := \frac{1}{4} \sum_{q,r=1}^{d} \sum_{i \in V} \sum_{j:(i,j) \in E} h_{qr} (\nabla f^q)_{ij} (\nabla f^r)_{ij} = \frac{1}{4} \sum_{(i,j) \in E} \|W (\nabla F)_{ij}\|^2 :$$

with $H = W^\top W$ with $W \in R^{d \times d}$

If we minimize E_W^{Dir} we expect $\|(\nabla F)_{ij}\|$ to shrink `except' when inside $\ker(H)$

Generalized harmonic flow on graphs is smoothing

We treat W as learnable weights and study the gradient flow of E_W^{Dir}:

$$\dot{F}(t) = -\nabla_F E_W^{\mathrm{Dir}}(F(t)) = -F(t)W^\top W$$

Generalized harmonic flow on graphs is smoothing

We treat W as learnable weights and study the gradient flow of E_W^{Dir}:

$$\dot{F}(t) = -\nabla_F E_W^{\mathrm{Dir}}(F(t)) = -F(t)W^\top W$$

Proposition (Di Giovanni, Rowbottom, et al.)

The dynamics is smoothing. Let $P_W^{\ker}$ be the projection onto $\ker(W^\top W)$, then

$$E^{\mathrm{Dir}}(F(t)) \leq e^{-2t\,\mathrm{gap}(W^\top W)\,\mathrm{gap}(\Delta)}\|F(0)\|^2 + E^{\mathrm{Dir}}((P_W^{\ker} \otimes I_n)\mathrm{vec}(F(0)));\quad t \geq 0.$$

$\forall\, 1 \leq i \leq d$: for each $i \in V$ we have $f_i(t) \to \frac{\sqrt{d_i}}{\sum_{i=1}^{n} d_i} + P_W^{\ker} f_i(0)$.

Generalized harmonic flow on graphs is smoothing

We treat W as learnable weights and study the gradient flow of E_W^{Dir}:

$$\dot{F}(t) = -\nabla_F E_W^{\mathrm{Dir}}(F(t)) = -F(t)W^\top W$$

Proposition (Di Giovanni, Rowbottom, et al.)

The dynamics is smoothing. Let $P_W^{\ker}$ be the projection onto $\ker(W^\top W)$, then

$$E^{\mathrm{Dir}}(F(t)) \leq e^{-2t\,\mathrm{gap}(W^\top W)\,\mathrm{gap}(\cdot)}\|F(0)\|^2 + E^{\mathrm{Dir}}((P_W^{\ker} \otimes I_n)\mathrm{vec}(F(0)));\quad t \geq 0.$$

$\exists\, x_1 \in \mathbb{R}^d$: for each $i \in V$ we have $f_i(t) \to \sqrt{d_i}\, x_1 + P_W^{\ker} f_i(0)$.

Generalized harmonic flow on graphs is smoothing

We treat W as learnable weights and study the gradient flow of E_W^{Dir}:

$$\dot{F}(t) = -\nabla_F E_W^{\mathrm{Dir}}(F(t)) = -F(t)W^\top W$$

Proposition (Di Giovanni, Rowbottom, et al.)

The dynamics is smoothing. Let $P_W^{\ker}$ be the projection onto $\ker(W^\top W)$, then

$$E^{\mathrm{Dir}}(F(t)) \leq e^{-2t\,\mathrm{gap}(W^\top W)\,\mathrm{gap}(\Delta)}\|F(0)\|^2 + E^{\mathrm{Dir}}((P_W^{\ker} \otimes I_n)\mathrm{vec}(F(0)));\quad t \geq 0.$$

$\forall \in \mathbb{R}^d$: for each $i \in V$ we have $f_i(t) \to \frac{\sqrt{d_i}}{\sqrt{1+d_i}} + P_W^{\ker} f_i(0)$.

Generalized harmonic flow on graphs is smoothing

We treat W as learnable weights and study the gradient flow of E_W^{Dir}:

$$\dot{F}(t) = -\nabla_F E_W^{\mathrm{Dir}}(F(t)) = -\Delta_F(t) W^\top W$$

Proposition (Di Giovanni, Rowbottom, et al.)

The dynamics is smoothing. Let $P_W^{\ker}$ be the projection onto $\ker(W^\top W)$, then

$$E^{\mathrm{Dir}}(F(t)) \leq e^{-2t\,\mathrm{gap}(W^\top W)\,\mathrm{gap}(\Delta)} \|F(0)\|^2 + E^{\mathrm{Dir}}((P_W^{\ker} \otimes I_n)\mathrm{vec}(F(0)));\quad t \geq 0.$$

$\forall \alpha_1 \in \mathbb{R}^d$: for each $i \in V$ we have $f_i(t) \to \frac{\sqrt{d_i}}{\alpha_1} + P_W^{\ker} f_i(0)$.

Generalized harmonic flow on graphs is smoothing

We treat W as learnable weights and study the gradient flow of E_W^{Dir}:

$$\dot{F}(t) = -\nabla_F E_W^{\mathrm{Dir}}(F(t)) = -F(t)W^\top W$$

Proposition (Di Giovanni, Rowbottom, et al.)

The dynamics is smoothing. Let $P_W^{\ker}$ be the projection onto $\ker(W^\top W)$, then

$$E^{\mathrm{Dir}}(F(t)) \leq e^{-2t\,\mathrm{gap}(W^\top W)\,\mathrm{gap}(\Delta)}\|F(0)\|^2 + E^{\mathrm{Dir}}((P_W^{\ker} \otimes I_n)\mathrm{vec}(F(0)));\quad t \geq 0.$$

$\forall \alpha_1, \ldots, \alpha_d \in \mathbb{R}^d$: for each $i \in V$ we have $f_i(t) \to \prod_{j=1}^{d} \alpha_j^{\sqrt{d_i}} + P_W^{\ker} f_i(0)$.

A few comments on the graph harmonic ow

- No W separates the limit embeddings of nodes with same degree and input features

[15] Similar to Nt and Maehara (2019); Oono and Suzuki (2020)

[16] This is different from Nt and Maehara (2019); Oono and Suzuki (2020); Cai and Wang (2020)

A few comments on the graph harmonic ow

- No W separates the limit embeddings of nodes with same degree and input features
- If W has zero kernel, nodes with same degrees converge to the same representation a over-smoothing occurs[15]

[15] Similar to Nt and Maehara (2019); Oono and Suzuki (2020)
[16] This is different from Nt and Maehara (2019); Oono and Suzuki (2020); Cai and Wang (2020)

A few comments on the graph harmonic ow

- No W separates the limit embeddings of nodes with same degree and input features
- If W has zero kernel, nodes with same degrees converge to the same representation a over-smoothing occurs[15]
- Over-smoothing occurs independently of the spectral radius of W or if its eigenvalues are positive– even for equations which lead to residual MPNNs when discretized[16]

[15] Similar to Nt and Maehara (2019); Oono and Suzuki (2020)
[16] This is different from Nt and Maehara (2019); Oono and Suzuki (2020); Cai and Wang (2020)

GNNs as Gradient Flows part 2: multi-particle energy approach

A more general energy

We can rewrite $E_W^{\text{Dir}}(F) = \frac{1}{2}\sum_i \langle f_i; W^\top W f_i\rangle \quad \frac{1}{2}\sum_{i,j} a_{ij}\langle f_i; W^\top W f_j\rangle$

Replace $W^\top W$ with symmetric matrices $; W \in \mathbb{R}^{d\ d}$!

$$E^{\text{tot}}(F) := \frac{1}{2}\sum_i \langle f_i;\ f_i\rangle \quad \frac{1}{2}\sum_{i,j} a_{ij}\langle f_i; W f_j\rangle \quad E^{\text{ext}}(F) + E_W^{\text{pair}}(F)$$

A more general energy

We can rewrite $E_W^{\text{Dir}}(F) = \frac{1}{2}\sum_i \langle f_i; W^\top W f_i\rangle \quad \frac{1}{2}\sum_{i,j} a_{ij}\langle f_i; W^\top W f_j\rangle$

Replace $W^\top W$ with symmetric matrices $; W \in \mathbb{R}^{d\ d}$!

$$E^{\text{tot}}(F) := \frac{1}{2}\sum_i \langle f_i;\ f_i\rangle \quad \frac{1}{2}\sum_{i,j} a_{ij}\langle f_i; W f_j\rangle \quad E^{\text{ext}}(F) + E_W^{\text{pair}}(F)$$

The gradient flow of E^{tot} is

$$\dot{F}(t) = -\nabla_F E^{\text{tot}}(F(t)) = -F(t) + AF(t)W:$$

Attraction vs repulsion

Node-features as particles in $\mathbb{R}^d$ with energy E^{tot}

- E^{ext} is independent of the graph topology external field
- E^{pair}_W potential energy, with W defining pairwise interactions of adjacent nodes

Attraction vs repulsion

Node-features as particles in $\mathbb{R}^d$ with energy E^{tot}

- E^{ext} is independent of the graph topology — external field
- E^{pair}_W potential energy, with W defining pairwise interactions of adjacent nodes

Decompose W = $\quad$ into positive and negative eigenvalues

Attraction vs repulsion

$$W = \quad + \quad + \quad >$$

$$E^{tot}(F) = \frac{1}{2}\sum_{i} hf_i;(\quad W)f_i i + \frac{1}{4}\sum_{i;j} jj \quad + (r\ F)_{ij}\, jj^2 \quad \frac{1}{4}\sum_{i;j} jj \quad (r\ F)_{ij}\, jj^2:$$

Attraction vs repulsion

W = > + + >

$$E^{tot}(F) = \frac{1}{2}\sum_i hf_i;(\quad W)f_i i + \frac{1}{4}\sum_{i;j} jj \quad + (r\ F)_{ij}\, jj^2 \quad \frac{1}{4}\sum_{i;j} jj \quad (r\ F)_{ij}\, jj^2:$$

The gradient ow minimizes E^{tot} ! W encodes..

- attraction via its positive eigenvalues since $+ (r\ F)_{ij}\, jj^2$ decreases edge-wise
- repulsion via its negative eigenvalues since $(r\ F)_{ij}\, jj^2$ increases edge-wise

Spectrum of W induces LFD or HFD

Consider $\dot{F}(t) = AF(t)W$ ($\star$) $\quad \mathrm{vec}(\dot{F}(t)) = (W \otimes A)\mathrm{vec}(F(t))$

Spectrum of W induces LFD or HFD

Consider $\dot{F}(t) = AF(t)W$ ( ) $\quad \mathrm{vec}(\dot{F}(t)) = (W \quad A)\mathrm{vec}(F(t))$

Write the spectrum of W as $\{\lambda_r^W\}$ with $\lambda_+^W = (\max \lambda_r^W)_+$ and $\lambda^W = (\min \lambda_r^W)$

Spectrum of W induces LFD or HFD

Consider $\dot{F}(t) = AF(t) + F(t)W$ ($\star$) $\quad \mathrm{vec}(\dot{F}(t)) = (I \otimes A + W^{T} \otimes I)\mathrm{vec}(F(t))$

Write the spectrum of W as $\{\lambda_r^W\}$ with $\lambda_+^W = (\max_r \lambda_r^W)_+$ and $\lambda_-^W = (\min_r \lambda_r^W)_-$

Any eigenvalue of $W \otimes I + I \otimes A$ can be written as $\lambda_r^W + \lambda_i^A = \lambda_r^W(1 + \lambda_i)$

Spectrum of W induces LFD or HFD

Consider $\dot{F}(t) = AF(t)W$ $\Rightarrow$ $\mathrm{vec}(\dot{F}(t)) = (W \otimes A)\mathrm{vec}(F(t))$

Write the spectrum of W as $\{\lambda_r^W\}$ with $\lambda_+^W = (\max_r \lambda_r^W)_+$ and $\lambda_-^W = (\min_r \lambda_r^W)_-$

Any eigenvalue of $W \otimes A$ can be written as $\lambda_r^W \lambda_i^A = \lambda_r^W(1 - \epsilon_i)$

Let P_W be the projection onto the eigenspace of W associated with λ_+^W $:= \lim_{j \to \infty} (\lambda^W)^j/(\lambda_+^W - 1)!$ Recall that λ_+^W is the largest eigenvalue of $\mathcal{L} = I \otimes A$

Spectrum of W induces LFD or HFD

Proposition (Di G. , Rowbottom , et al.)

If $> \ _+^W$, then $E(t) = AF(t)W$ is HFD for a.e. $F(0)$: there exists $_{HFD}$ such that[17]

$$E^{Dir}(F(t)) = e^{2t} \quad \frac{}{2} \, jj P_W F(0) jj^2 + O(e^{\ 2t \ _{HFD}}) \quad ; \quad t \quad 0;$$

and $F(t) = jj F(t) jj$ converges to $F_1 \ 2 \ R^n \ ^d$ such that $f \ _1^r = \ f_1^r$, for $1 \quad r \quad d$.

[17] We have an explicit formula depending on `spectral gaps' of and W

Spectrum of W induces LFD or HFD

Proposition (Di G. , Rowbottom , et al.)

If $> \ ^{W}_{+}$, then $E(t) = AF(t)W$ is HFD for a.e. $F(0)$: there exists $_{HFD}$ such that[18]

$$E^{Dir}(F(t)) = e^{2t} \quad \frac{}{2} \, jj P_W F(0) jj^2 + O(e^{2t \ _{HFD}}) \quad ; \quad t \quad 0;$$

and $F(t) = jj F(t) jj$ converges to $F_1 \ 2 \ R^n \ ^d$ such that $f \ ^r_1 = f \ ^r_1$, for $1 \quad r \quad d$.

[18] We have an explicit formula depending on `spectral gaps' of W

Spectrum of W induces LFD or HFD

Proposition (Di G. , Rowbottom , et al.)

If $> \ _+^W$, then $E_-(t) = AF(t)W$ is HFD for a.e. $F(0)$: there exists $_{HFD}$ such that[19]

$$E^{Dir}(F(t)) = e^{2t} \quad \frac{}{2} \, jj P_W F(0) jj^2 + O(e^{2t \ _{HFD}}) \quad ; \quad t \quad 0;$$

and $F(t) = jj F(t) jj$ converges to $F_1 \ 2 \ R^n \ ^d$ such that $f_1^r = f_1^r$, for $1 \quad r \quad d$.

[19] We have an explicit formula depending on `spectral gaps' of and W

Spectrum of W induces LFD or HFD

Proposition (Di G. , Rowbottom , et al.)

If $> \ _+^W$, then $E(t) = AF(t)W$ is HFD for a.e. $F(0)$: there exists $_{HFD}$ such that[20]

$$E^{Dir}(F(t)) = e^{2t} \quad \frac{}{2} jj P_W F(0) jj^2 + O(e^{2t\ HFD}) \ ; \quad t \quad 0;$$

and $F(t) = jj F(t) jj$ converges to $F_1 \ 2 \ R^n \ ^d$ such that $f \ _1^r = \ f_1^r$, for $1 \quad r \quad d$.

[20] We have an explicit formula depending on `spectral gaps' of and W

Spectrum of W induces LFD or HFD

Proposition (Di G. , Rowbottom , et al.)

If $> \ _+^W$, then $E(t) = AF(t)W$ is HFD for a.e. $F(0)$: there exists $_{HFD}$ such that[21]

$$E^{Dir}(F(t)) = e^{2t} \quad \frac{}{2} \, jj P_W F(0) jj^2 + O(e^{2t \ _{HFD}}) \ ; \quad t \quad 0;$$

and $F(t) = jj F(t) jj$ converges to $F_1 \ 2 \ R^n \ ^d$ such that $f \ _1^r = f_1^r$, for $1 \quad r \quad d$.

[21] We have an explicit formula depending on `spectral gaps' of W

Spectrum of W induces LFD or HFD

Proposition (Di G. , Rowbottom , et al.)

If $> \|W\|_+$, then $E(t) = AF(t)W$ is HFD for a.e. $F(0)$: there exists $_{HFD}$ such that[22]

$$E^{Dir}(F(t)) = e^{2t} \quad \frac{}{2} \|P_W F(0)\|^2 + O(e^{2t\ HFD}) \quad ; \quad t \quad 0;$$

and $F(t) = \|F(t)\|$ converges to $F_1 \ 2 \ R^n \ ^d$ such that $f_1^{\ r} = f_1^{\ r}$, for $1 \quad r \quad d$.

[22] We have an explicit formula depending on `spectral gaps' of and W

Proposition (Di G. , Rowbottom , et al.)

If $> \lambda_+^W$, then $E(t) = AF(t)W$ is HFD for a.e. $F(0)$: there exists $_{HFD}$ such that[22]

$$E^{Dir}(F(t)) = e^{2t} \quad \frac{}{2} \, jj P_W F(0) jj^2 + O(e^{2t\ _{HFD}}) \quad ; \quad t \quad 0;$$

and $F(t) = jj F(t) jj$ converges to $F_1 \ 2 \ R^n \ ^d$ such that $f \ _1^r = f_1^r$, for $1 \quad r \quad d$.

If enough mass is distributed over the negative eigenvalues of the `channel-mixing', graph
high frequencies dominate what matters is how the spectra of and W interact

[22] We have an explicit formula depending on `spectral gaps' of and W

Source term and a more general family of energies

Equations with a source term may have better expressive power [23]

In our framework: add an extra energy term $E_W^{\text{source}}(F) := \langle F, F(0) \rangle_W$!

$$E(t) = F(t) + AF(t) - F(0)$$

[23] Xhonneux et al. (2020); Chen et al. (2020); Thorpe et al. (2021)

Source term and a more general family of energies

Equations with a source term may have better expressive power [23]

In our framework: add an extra energy term $E_W^{\text{source}}(F) := \langle F, F(0) W \rangle$

$$E(t) = F(t) + A F(t) W - F(0) W :$$

We can also replace A with A satisfying $A_{ij} = 0$ if $(i,j) \in E$!

$$E_{A,W}^{\text{pair}}(F) := \sum_{(i,j)} A_{ij} \langle f_i, W f_j \rangle :$$

[23] Xhonneux et al. (2020); Chen et al. (2020); Thorpe et al. (2021)

Comments on non-linear activations

Non-linear function can `activate' the inner products in the energy:

$$E^{\text{ext}}(F) + E_W^{\text{pair}}(F) = \frac{1}{2}\sum_i (\mathbf{f}_i; f_i) \quad \frac{1}{2}\sum_{i;j} a_{ij} (\mathbf{f}_i; W\mathbf{f}_j):$$

[24] Wu et al. (2019); Oono and Suzuki (2020); Chen et al. (2020)

Comments on non-linear activations

Non-linear function can `activate' the inner products in the energy:

$$E^{\text{ext}}(F) + E_W^{\text{pair}}(F) = \frac{1}{2}\sum_i (\langle f_i, f_i \rangle) \quad \frac{1}{2}\sum_{i,j} a_{ij} \, (\langle f_i, W f_j \rangle):$$

A few reasons why we keep the gradient linear

- Perform spectral analysis in closed form[24]
- We have seen no gain in performance when including non-linear activations
- We can `push the non-linear maps' in either the encoding block or the decoding one

[24] Wu et al. (2019); Oono and Suzuki (2020); Chen et al. (2020)

A comparison with (some) continuous GNN models

Recall the continuous models:

- Linear PDE GCN_D: $\dot{F}_{\mathrm{PDE\_GCN}_D}(t) = F(t)K(t)^{\top}K(t)$
- CGNN: $\dot{F}_{\mathrm{CGNN}}(t) = F(t) + F(t)\tilde{} + F(0)$ with symmetric $\tilde{}$
- Linear GRAND: $\dot{F}_{\mathrm{GRAND}}(t) = {}_{\mathrm{RW}}F(t) = (I - A(F(0)))F(t)$

A comparison with (some) continuous GNN models

Recall the continuous models:

- Linear PDE GCN$_D$: $E_{\text{PDE GCN}_D}(t) = F(t)K(t)^{\top}K(t)$
- CGNN: $E_{\text{CGNN}}(t) = F(t) + F(t)^{\sim} + F(0)$ with symmetric $^{\sim}$
- Linear GRAND: $E_{\text{GRAND}}(t) = {}_{\text{RW}}F(t) = (I - A(F(0)))F(t)$

Proposition (Di G. , Rowbottom , et al.)

(i) PDE GCN$_D$ is a smoothing model $E^{\text{Dir}}(F_{\text{PDE GCN}_D}(t)) \to 0$.

A comparison with (some) continuous GNN models

Recall the continuous models:

- Linear PDE GCN_D: $E_{PDE\ GCN_D}(t) = F(t)K(t)^{>}K(t)$
- CGNN: $E_{CGNN}(t) = F(t) + F(t)^{\sim} + F(0)$ with symmetric $^{\sim}$
- Linear GRAND: $E_{GRAND}(t) = {}_{RW}F(t) = (I\ A(F(0)))F(t)$

Proposition (Di G., Rowbottom, et al.)

(i) PDE GCN_D is a smoothing model: $E^{Dir}(F_{PDE\ GCN_D}(t)) \to 0$.

(ii) For a.e. $F(0)$ it holds: CGNN is never HFD and if we remove the source term, then $E^{Dir}(F_{CGNN}(t) = jjF_{CGNN}(t)jj) \to e^{gap(\ )t}$.

A comparison with (some) continuous GNN models

Recall the continuous models:

- Linear PDE GCN_D: $E_{PDE\ GCN_D}(t) = F\ (t) K(t)^{>} K(t)$
- CGNN: $E_{CGNN}(t) = F\ (t) + F(t)^{\sim} + F(0)$ with symmetric $\sim$
- Linear GRAND: $E_{GRAND}(t) = {}_{RW} F(t) = (I\ A(F(0))) F(t)$

Proposition (Di G. , Rowbottom , et al.)

(i) PDE GCN_D is a smoothing model $E^{Dir}(F_{PDE\ GCN_D}(t))\ 0$.

(ii) For a.e. $F(0)$ it holds: CGNN is never HFD and if we remove the source term, then $E^{Dir}(F_{CGNN}(t) = jj F_{CGNN}(t) jj)\ e^{\ gap(\ )t}$.

(iii) If G is connected, $F_{GRAND}(t)\ !\ $ as $t\ !\ 1$, with $^{r} = mean(f^{r}(0)), 1\ r\ d$.

GNNs as Gradient Flows part 3: discrete setting

The requirement for symmetry

When classical MPNNs are discretized gradient ows?

The requirement for symmetry

When classical MPNNs are discretized gradient flows?

Let $A \in \mathbb{R}^{n \times n}$ be a symmetric graph vector field $(A)_{ij} = 0, (i,j) \notin E$

Consider a family of linear GNNs with shared weights of the form

$$F(t+1) = F(t) + AF(t)W + F(0)V, \quad 0 \leq t \leq T.$$

They are gradient flow of a `multi-particle' energy iff V and W are symmetric.

Can graph convolutional models be high-frequency dominated?

Introduce step-size 1 and consider gradient ow system

$$F(t+\ ) = F(t) + \quad AF(t)W; \quad W = W^{>};$$

Let P_W be the projection into the eigenspace of W of $A = W \quad (I \quad)$ associated with the eigenvalue $:= j \ ^{W} j(\quad 1)$ and set

$$\overset{W}{+}(\quad 1))\ ^{1} < j \ ^{W} j < 2(\ (2 \quad))\ ^{1} \qquad (1)$$

Can graph convolutional models be high-frequency dominated?

Introduce step-size 1 and consider gradient ow system

$$F(t +) = F(t) + AF(t)W ; \quad W = W^{>} ;$$

Let P_W be the projection into the eigenspace of W of $A = W (I)$ associated with the eigenvalue $:= j {}^{W} j(1)$ and set

$$\overset{W}{+} (1)) {}^{1} < j {}^{W} j < 2((2)) {}^{1} \qquad (2)$$

Can graph convolutional models be high-frequency dominated?

Introduce step-size 1 and consider gradient ow system

$$F(t +) = F(t) + AF(t)W ; \quad W = W^{>} ;$$

Let P_W be the projection into the eigenspace of W of $A = W (I)$ associated with the eigenvalue $:= j {}^{W} j(1)$ and set

$$ {}^{W}_{+} (1)) {}^{1} < j {}^{W} j < 2((2)) {}^{1} \qquad (3)$$

Can graph convolutional models be high-frequency dominated?

Theorem (Di G. , Rowbottom , et al.)

If equation 3 holds then there exists $_{HFD} <$ s.t.

$$E^{Dir}(F(m)) = (1 +)^{2m} \frac{}{2} jj P_W F(0) jj^2 + O\left(\frac{1 + _{HFD}}{1 + }\right)^{2m !!} :$$

The dynamics is HFD for a.e. $F(0)$ and $F(m) = jj F(m) jj ! F_1$ s.t. $f_1^{r} = f_1^{r}$.

Can graph convolutional models be high-frequency dominated?

Theorem (Di G. , Rowbottom , et al.)

If equation 3 holds then there exists $_{HFD} <$ s.t.

$$E^{Dir}(F(m)) = (1 + \quad)^{2m} \quad \frac{}{2} jj P_W \; F(0) jj^2 + O\left(\frac{1 + \quad}{1 + \quad_{HFD}}\right)^{2m !!} :$$

The dynamics is $_{HFD}$ for a.e. $F(0)$ and $F(m) = jj F(m) jj ! \quad F_1$ s.t. $f_1^{r} = f_1^{r}$.

Conversely, if G is not bipartite, then for a.e. $F(0)$ the system $F(t +) = AF(t)W$, with W symmetric, is $_{LFD}$ independent of the spectrum of W .

HFD dynamics via negative eigenvalues

- linear discrete gradient flows can be HFD due to the negative eigenvalues of W
  - Differently from previous results[25], no bound on spectral radius of W coming from the graph topology as long as τ is small enough
    - Recall that previous over-smoothing results required W to have sufficiently small singular values depending on the spectrum of L
    - If we have symmetry and control the spectrum of W we can avoid over-smoothing (and in fact be HFD) in terms of positive vs negative eigenvalues of W

[25] Nt and Maehara (2019); Oono and Suzuki (2020); Cai and Wang (2020)

The role of the residual connection

- Without a residual term the dynamics is ill-conditioned for a.e. $F(0)$ independently of the sign and magnitude of the eigenvalues of W
  - provides a justification for the residual connection in terms of the spectrum of W
  - explains via induced dynamics and spectral analysis the `expressivity' results in Chen et al. (2020)

Reversing time and sign of the edge weights

Let $\{\frac{W}{r}\}$ be the spectrum of W with orthonormal eigenvectors $\{\frac{W}{r}\}$ and $= U U^{>}$

[26] Similar effect as in Bo et al. (2021); Yan et al. (2021)

Reversing time and sign of the edge weights

Let $\{ {}^{W}_{r} g$ be the spectrum of W with orthonormal eigenvectors $\{ {}^{W}_{r} g$ and $= U U^{>}$

Introduce $z^{r}(t) : V ! R$ defined by $z^{r}_{i}(t) = hf_{i}(t); {}^{W}_{r} i$, then gradient flow becomes:

$$z^{r}(t +) = U(I + {}^{W}_{r}(I)) U^{>} z^{r}(t) = z^{r}(t) + {}^{W}_{r} A z^{r}(t)$$

Along ${}^{W}_{r}$ if ${}^{W}_{r} < 0$ then the dynamics is equivalent to flipping the sign of the edges [26].

[26] Similar effect as in Bo et al. (2021); Yan et al. (2021)

GNNs as Gradient Flows part 4: ablation studies and experiments

Structure of the framework

General ingredients of the framework GRAFF (Gradient Flow Framework)

- Encoding block $_{EN} : R^{n \quad p} ! \quad R^{n \quad d}$ is used to process input features $E_S 2 R^{n \quad p}$

Structure of the framework

General ingredients of the framework GRAFF (Gradient Flow Framework)

- Encoding block $_{EN} : R^{n \ p} ! \ R^{n \ d}$ is used to process input features $E_0 2 R^{n \ p}$
- Symmetric channel-mixing matrices $; W \ 2 \ R^{d \ d}$ that are shared across the layers

Structure of the framework

General ingredients of the framework GRAFF (Gradient Flow Framework)

- Encoding block $_{EN} : R^{n \ p} ! \ R^{n \ d}$ is used to process input features s $2 \ R^{n \ p}$
- Symmetric channel-mixing matrices ; $W \ 2 \ R^{d \ d}$ that are shared across the layers
- Decoding block $_{DE} : R^{n \ d} ! \ R^{n \ k}$, where k is the number of label classes

Structure of the framework

General ingredients of the framework GRAFF (Gradient Flow Framework)

- Encoding block $_{EN} : R^{n\ p} ! R^{n\ d}$ is used to process input features $F_0 2 R^{n\ p}$
- Symmetric channel-mixing matrices $; W\ 2\ R^{d\ d}$ that are shared across the layers
- Decoding block $_{DE} : R^{n\ d} ! R^{n\ k}$, where k is the number of label classes

$$F(t +) = F(t) + F(t) + AF(t)W + F(0) ; F(0) = _{EN}(F_0);$$

Different choices for W

- Sum-variant: $W = W^{0} + W^{\triangleright}$! `no-control' on spectrum

[27] Provides justi cation to Chen et al. (2020)

Different choices for W

- Sum-variant: $W = W^{0} + W^{0>}$! `no-control' on spectrum
- (Neg)-Prod $W = W^{0>}W^{0}$! signed eigenvalues

[27] Provides justi cation to Chen et al. (2020)

Different choices for W

- Sum-variant: $W = W^0 + W^{0\top}$! `no-control' on spectrum
- (Neg)-Prod $W = W^{0\top} W^0$! signed eigenvalues
- W diagonally-dominant (DD) : take W^0 symmetric with zero diagonal and $w \in \mathbb{R}^d$ de ned by $w = q \sum |W^0| + r$, and set $W = \mathrm{diag}(w) + W^0$! by Gershgorin Theorem the model `can' easily re-distribute mass in the spectrum on $\mathbb{R}$ [27].

[27] Provides justi cation to Chen et al. (2020)

Complexity and number of parameters

GRAFF scales as $O(|V|pd + |E|d)$, where p and d are input feature and hidden dimension

! our model is faster than GCN with small number of parameters $pd + d^2 + 3d + dk$

Figure 4: Runtime ablation for inference on Cora dataset

Ablation and synthetic experiments: setting

Recall our claims about role of `channel-mixing' W:

- Positive eigenvalues of W induce attraction in a residual convolutional model

Ablation and synthetic experiments: setting

Recall our claims about role of `channel-mixing' W:

- Positive eigenvalues of W induce attraction in a residual convolutional model
- Negative eigenvalues of W induce repulsion in a residual convolutional model

Ablation and synthetic experiments: setting

Recall our claims about role of `channel-mixing' W:

- Positive eigenvalues of W induce attraction in a residual convolutional model
- Negative eigenvalues of W induce repulsion in a residual convolutional model
- A non-residual convolutional model is always dominated by low-frequencies independent of the spectrum of the W

Ablation and synthetic experiments: setting

Recall our claims about role of `channel-mixing' W:

- Positive eigenvalues of W induce attraction in a residual convolutional model
- Negative eigenvalues of W induce repulsion in a residual convolutional model
- A non-residual convolutional model is always dominated by low-frequencies independent of the spectrum of the W

To investigate our claims we use the synthetic Cora dataset of Zhu et al. (2020)

! graphs are generated for target levels of homophily via preferential attachment: we expect LFD to be better than HFD with high homophily and vice-versa for low homophily

Ablation and synthetic experiments: part 1

Goal: Explain performance wrt homophily in terms of the spectrum of W

- Neg-prod is better than prod on low-homophily graphs, confirms HFD dynamics

Ablation and synthetic experiments: part 1

Goal: Explain performance wrt homophily in terms of the spectrum of W

- Neg-prod is better than prod on low-homophily graphs, confirms HFD dynamics
- prod (attraction-only) struggles in low-homophily even with residual connection

Ablation and synthetic experiments: part 1

Goal: Explain performance wrt homophily in terms of the spectrum of W

- Neg-prod is better than prod on low-homophily graphs, confirming HFD dynamics
- prod (attraction-only) struggles in low-homophily even with residual connection
- `neutral' variants like sum and (DD) are more flexible and outperform GCN confirming that non-residual convolutional models are LFD irrespectively of the spectrum of W

Ablation and synthetic experiments: part 2

Goal: Use homophily to assess if the evolution is smoothing: compute homophily of the prediction (cross) and compare with that read from the encoding (i.e. $\mathbf{z}$ evolution)

Ablation and synthetic experiments: part 2

Goal: Use homophily to assess if the evolution is smoothing (i.e. compute homophily of the prediction (cross) and compare with that read from the encoding (no evolution))

- neg-prod: homophily decreases after evolution while with prod the prediction is smoother than the true homophily

Ablation and synthetic experiments: part 2

Goal: Use homophily to assess if the evolution is smoothing ($\mathcal{H}$ compute homophily of the prediction (cross) and compare with that read from the encoding (i.e. $\mathcal{H}$ no evolution)

- neg-prod: homophily decreases after evolution while with prod the prediction is smoother than the true homophily
- (DD) and sum variants adapt better to the true homophily

Ablation and synthetic experiments: part 2

Goal: Use homophily to assess if the evolution is smoothing ($\mathcal{H}_{\text{evolution}}$) compute homophily of the prediction (cross) and compare with that read from the encoding (i.e. $\mathcal{H}_{\text{encoding}}$)

- neg-prod: homophily decreases after evolution while with prod the prediction is smoother than the true homophily
- (DD) and sum variants adapt better to the true homophily
- The encoding compensates when the spectrum of W has a sign

Conclusions and where to next?

What was the message then?

- Framework where the MPNNs equations minimize a multi-particle learnable energy

What was the message then?

- Framework where the MPNNs equations minimize a multi-particle learnable energy
- Analysis of the interaction between the spectrum of the graph and the spectrum of th
  `channel-mixing'! when and why the dynamics is low (high) frequency dominated

What was the message then?

- Framework where the MPNNs equations minimize a multi-particle learnable energy
- Analysis of the interaction between the spectrum of the graph and the spectrum of th `channel-mixing'! when and why the dynamics is low (high) frequency dominated
- Re ned existing asymptotic analysis of MPNNs to account for the role of the spectrum of the channel-mixing

What was the message then?

- Framework where the MPNNs equations minimize a multi-particle learnable energy
- Analysis of the interaction between the spectrum of the graph and the spectrum of th `channel-mixing'! when and why the dynamics is low (high) frequency dominated
- Re ned existing asymptotic analysis of MPNNs to account for the role of the spectrum of the channel-mixing
- From a practical perspective, our framework allows for `educated' choices resulting in a simple, more explainable convolutional model: our results refute the folklore of graph convolutional models being too `simple' for complex benchmarks.

Limitations and future directions

We restricted to a constant bilinear form W, how about non-constant alternatives $W(F;t)$ that are aware of the features? requirement for local `heterogeneity' with efficiency

Limitations and future directions

We restricted to a constant bilinear form W, how about non-constant alternatives $W(F;t)$ that are aware of the features? requirement for local `heterogeneity' with efficiency

What can we say about dynamics that are neither LFD nor HFD?

Limitations and future directions

We restricted to a constant bilinear form W, how about non-constant alternatives $W(F;t)$ that are aware of the features? requirement for local `heterogeneity' with efficiency

What can we say about dynamics that are neither LFD nor HFD?

The energy formulation points to new models more `physics' inspired

Thank you!

For any question/complaint/video-game recommendation do not hesitate to contact me! :

fdigiovanni (at) twitter com

@Francesco_dgv

References

Bo, D., Wang, X., Shi, C., and Shen, H. (2021). Beyond low-frequency information in graph convolutional networks. In AAAI. AAAI Press.

Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. (2017). Geometric deep learning: going beyond euclidean data. IEEE Signal Processing Magazine, 34(4):18–42.

Cai, C. and Wang, Y. (2020). A note on over-smoothing for graph neural networks. arXiv preprint arXiv:2006.13318.

Chamberlain, B., Rowbottom, J., Gorinova, M. I., Bronstein, M., Webb, S., and Rossi, E. (2021). Grand: Graph neural diffusion. In International Conference on Machine Learning, pages 1407–1418. PMLR.

Chen, M., Wei, Z., Huang, Z., Ding, B., and Li, Y. (2020). Simple and deep graph convolutional networks. In International Conference on Machine Learning, pages 1725–1735. PMLR.

Chen, R. T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. (2018). Neural ordinary differential equations. Advances in neural information processing systems, 31.

Eells, J. and Sampson, J. H. (1964). Harmonic mappings of riemannian manifolds. American journal of mathematics, 86(1):109–160.

Eliasof, M., Haber, E., and Treister, E. (2021). Pde-gcn: Novel architectures for graph neural networks motivated by partial differential equations. Advances in Neural Information Processing Systems, 34.

Haber, E. and Ruthotto, L. (2018). Stable architectures for deep neural networks. Inverse Problems, 34.

Kimmel, R., Sochen, N., and Malladi, R. (1997). From high energy physics to low level vision. In International Conference on Scale-Space Theories in Computer Vision, pages 236–247. Springer.

Kipf, T. N. and Welling, M. (2017). Semi-Supervised Classi cation with Graph

Convolutional Networks. In *Proceedings of the 5th International Conference on Learning Representations*, ICLR '17.

Nt, H. and Maehara, T. (2019). Revisiting graph neural networks: All we have is low-pass filters. *arXiv preprint arXiv:1905.09550*.

Oono, K. and Suzuki, T. (2020). Graph neural networks exponentially lose expressive power for node classification. In *International Conference on Learning Representations*.

Perona, P. and Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *PAMI*, 12(7):629–639.

Rusch, T. K., Chamberlain, B. P., Rowbottom, J., Mishra, S., and Bronstein, M. M. (2022). Graph-coupled oscillator networks. In *International Conference on Machine Learning*.

Thorpe, M., Nguyen, T. M., Xia, H., Strohmer, T., Bertozzi, A., Osher, S., and Wang, B. (2021). Grand++: Graph neural diffusion with a source term. In *International Conference on Learning Representations*.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. (2018). Graph attention networks. In *International Conference on Learning Representations*.

Wu, F., Souza, A., Zhang, T., Fifty, C., Yu, T., and Weinberger, K. (2019). Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR.

Xhonneux, L.-P., Qu, M., and Tang, J. (2020). Continuous graph neural networks. In *International Conference on Machine Learning*, pages 10432–10441. PMLR.

Yan, Y., Hashemi, M., Swersky, K., Yang, Y., and Koutra, D. (2021). Two sides of the same coin: Heterophily and oversmoothing in graph convolutional neural networks. *arXiv preprint arXiv:2102.06462*.

Zhou, D. and Schölkopf, B. (2005). Regularization on discrete spaces. In *Joint Pattern Recognition Symposium*, pages 361–368. Springer.

Zhu, J., Yan, Y., Zhao, L., Heimann, M., Akoglu, L., and Koutra, D. (2020). Beyond

homophily in graph neural networks: Current limitations and effective designs. *Advances in Neural Information Processing Systems*, 33:7793–7804.