Graph Neural Networks

Fernando Gama

Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

CosmoStat seminar on Machine Learning for Astrophysics March 5th, 2021

Thanks: D. Owerko, E. Tolstaya, L. Ruiz, Q. Li, T.-K. Hu, A. Prorok, A. G. Marques, Z. Wang, E. Isufi, G. Leus, J. Bruna, S. Sojoudi and A. Ribeiro



▶ Graphs are generic models of signal structure that can help to learn in several practical problems

Authorship Attribution

Recommendation Systems



Gama, Ribeiro, Bruna, "Diffusion Scattering Transforms on Graphs", ICLR 2019

Gama, Bruna, Ribeiro, "Stability of Graph Neural Networks to Relative Perturbations", IEEE ICASSP 2020



▶ Graphs are generic models of signal structure that can help to learn in several practical problems

Decentralized Control of Autonomous Systems

Smart Grids



Tolstaya, Gama, Paulos, Pappas, Kumar, Ribeiro, "Learning Decentralized Controllers for Robot Swarms with Graph Neural Networks", CoRL 2019

Owerko, Gama, Ribeiro, "Optimal Power Flow Using Graph Neural Networks", IEEE ICASSP 2020



► Successful machine learning leverages structure ⇒ Convolutional neural networks (CNNs) We are good at learning over this
Challenge is we want to learn on this





- Scales, exploits data structure, and has an efficient implementation (distributed)
- ▶ Graph Convolutions \Rightarrow Graph Signal Processing \Rightarrow Graph Filtering



Graph Neural Networks

Equivariance and Stability

Distributed Collaborative Intelligent Systems

- $\blacktriangleright \text{ Graph signal processing } \Rightarrow \text{Mathematical framework}$
- Graph convolutions \Rightarrow Local, distributed
 - \Rightarrow Generalize time convolutions
- ▶ Permutation equivariance \Rightarrow Exploit structure
- Stability to changes in the underlying network
- ► Transferability and scalability
- ▶ Team of agents \Rightarrow Collaborate to accomplish global task
- Autonomy of agents \Rightarrow Decentralized actions
- Global objective vs. Local actions



Outline

Graph Neural Networks fgama@berkeley.edu

Graph Neural Networks

Equivariance and Stability Properties

Robotics

Conclusions



Graph Neural Networks

Graph Neural Networks fgama@berkeley.edu

Graph Neural Networks

Equivariance and Stability Properties

Robotics

Conclusions



Graph Convolutions

• Graph convolution \Rightarrow Linear combination of shifted versions of the signal

$$\mathbf{x} *_{\mathbf{S}} \mathbf{h} = \sum_{k=0}^{K-1} h_k \mathbf{S}^k \mathbf{x}$$

 $\blacktriangleright \text{ Notion of shift } \mathbf{S} \ \Rightarrow \text{Matrix description of graph } \Rightarrow \mathbf{Sx \ shifts \ the \ signal \ x}$





Gama, Isufi, Leus, Ribeiro, "Graphs, Convolutions and Neural Networks: From Graph Filters to Graph Neural Networks", IEEE SPM, 2020



Graph Convolutions

Graph Neural Networks fgama@berkeley.edu

• Graph convolution \Rightarrow Linear combination of shifted versions of the signal

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

- ▶ Notion of shift $\mathbf{S} \Rightarrow$ Matrix description of graph (adjacency, Laplacian)
- Linear combination of neighboring signal \Rightarrow Local operation

	1					



Gama, Isufi, Leus, Ribeiro, "Graphs, Convolutions and Neural Networks: From Graph Filters to Graph Neural Networks", IEEE SPM, 2020



Nonlinear Graph Signal Processing

- Traditional signal processing
 - \Rightarrow Best linear filter that exploits structure

 $\min_{\{h_k\}} \mathsf{J}(\mathbf{z}_1) = \min_{\{h_k\}} \mathsf{J}(\mathbf{H}(\mathbf{S})\mathbf{x})$

- ▶ Linear models ⇒ Limited representation
 ⇒ Nonlinear graph signal processing
- ► Graph perceptron \Rightarrow Nonlinear processing \Rightarrow Graph filter \Rightarrow Pointwise nonlinearity \Rightarrow Learn graph filter $\{h_k\} \Rightarrow \min_{\{h_k\}} J(\mathbf{x}_1)$
- Basic nonlinear description of models
 - \Rightarrow Increase representation power \Rightarrow Repeat

Gama, Isufi, Leus, Ribeiro, "Graphs, Convolutions and Neural Networks: From Graph Filters to Graph Neural Networks", IEEE SPM, 2020





Graph Neural Networks

- \blacktriangleright Cascade of L layers
 - \Rightarrow Graph convolutions with filters $\mathcal{H} = \{\mathbf{h}_{\ell}\}$
 - \Rightarrow Pointwise nonlinearity (activation functions)
- ► The GNN $\Phi(\mathbf{x}; \mathbf{S}, \mathcal{H})$ depends on the filters \mathcal{H} ⇒ Learn filter taps \mathcal{H} from training data ⇒ Also depends on the graph **S**
- ► Nonlinear mapping $\Phi(\mathbf{x}; \mathbf{S}, \mathcal{H})$
 - \Rightarrow Exploit underlying graph structure ${\bf S}$
 - \Rightarrow Local information
 - \Rightarrow Distributed implementation



Gama, Marques, Leus, Ribeiro, "Convolutional Neural Network Architectures for Signals Supported on Graphs", IEEE TSP, 2019



Equivariance and Stability Properties of Graph Neural Networks (Stability Properties of Graph Neural Networks) (Stability Networks) (

Graph Neural Networks

Equivariance and Stability Properties

Robotics

Conclusions



Graph Neural Networks: Why?

- Time convolutions are intuitive. Graph convolutions not so much.
 ⇒ Local information, distributed implementation
- ▶ CNNs are good at machine learning \Rightarrow Translation equivariant, stable [Mallat '12]
- ▶ Permutation equivariance \Rightarrow Exploit internal symmetries of the graph
- ▶ Stability to graph perturbations \Rightarrow Similar graphs yield similar outputs
- ▶ Permutation Equivariance + Stability \Rightarrow Scalability and transferability

Gama, Ribeiro, Bruna, "Diffusion Scattering Transforms on Graphs", ICLR 2019 Gama, Bruna, Ribeiro, "Stability of Graph Scattering Transforms", NeurIPS 2019



Permutation Equivariance

• Consider the graph convolution operator $\mathbf{H}(\mathbf{S})\mathbf{x} = \sum_{k=0}^{\infty} h_k \mathbf{S}^k \mathbf{x}$

▶ Depends on filter parameters $\mathbf{h} = \{h_k\}_{k=0}^{\infty}$ and shift operator S; applied to the input signal x

Theorem (Gama, Bruna, Ribeiro)

Graph convolutions are equivariant to permutations. For graphs with permuted shift operators $\hat{\mathbf{S}} = \mathbf{P}^{\mathsf{T}} \mathbf{S} \mathbf{P}$ and permuted graph signals $\hat{\mathbf{x}} = \mathbf{P}^{\mathsf{T}} \mathbf{x}$ it holds

 $\mathbf{H}(\hat{\mathbf{S}})\hat{\mathbf{x}} = \mathbf{P}^{\mathsf{T}}\mathbf{H}(\mathbf{S})\mathbf{x}$

$$\mathbf{Proof} \Rightarrow \mathbf{H}(\hat{\mathbf{S}})\hat{\mathbf{x}} = \sum_{k=0}^{\infty} h_k \, \hat{\mathbf{S}}^k \hat{\mathbf{x}} = \sum_{k=0}^{\infty} h_k \, (\mathbf{P}^\mathsf{T} \mathbf{S} \mathbf{P})^k \mathbf{P}^\mathsf{T} \mathbf{x} = \mathbf{P}^\mathsf{T} \left(\sum_{k=0}^{\infty} h_k \, \mathbf{S}^k \mathbf{x} \right) = \mathbf{P}^\mathsf{T} \mathbf{H}(\mathbf{S}) \mathbf{x}$$

Gama, Bruna, Ribeiro, "Stability of Graph Scattering Transforms", NeurIPS 2019



GNNs Inherit Permutation Equivariance from Graph Filters

▶ GNN is a cascade of layers

 \Rightarrow Graph filters and pointwise nonlinearities

- Pointwise operation ⇒ No mixing of node values
 ⇒ Independent of the graph
- ▶ GNN retains permutation equivariance





Theorem (Gama, Ribeiro, Bruna)

GNNs are equivariant to permutations. For graphs with permuted shift operators $\hat{\mathbf{S}} = \mathbf{P}^{\mathsf{T}} \mathbf{S} \mathbf{P}$ and permuted graph signals $\hat{\mathbf{x}} = \mathbf{P}^{\mathsf{T}} \mathbf{x}$ it holds

$\Phi(\hat{\mathbf{x}}; \hat{\mathbf{S}}, \mathcal{H}) = \mathbf{P}^{\mathsf{T}} \Phi(\mathbf{x}; \mathbf{S}, \mathcal{H})$

where $\Phi(\hat{\mathbf{x}}; \hat{\mathbf{S}}, \mathcal{H})$ is the output of processing $\hat{\mathbf{x}}$ on $\hat{\mathbf{S}}$ with GNN \mathcal{H} and $\Phi(\mathbf{x}; \mathbf{S}, \mathcal{H})$ is the output of processing \mathbf{x} on \mathbf{S} with the same GNN \mathcal{H} .

Signal Processing with Graph Neural Networks is independent of labeling

Gama, Bruna, Ribeiro, "Stability Properties of Graph Neural Networks", TSP, 2020



Equivariance to Permutations is More Valuable than Apparent

Graph Neural Networks fgama@berkeley.edu

- ▶ Invariance to node relabelings allows GNNs to exploit internal symmetries of graph signals
- Although different, signals on (a) and (b) are permutations of one other
 - \Rightarrow Permutation equivariance means that the GNN can learn to process (b) from seeing (a)



▶ Permutation Equivariance is not a good idea in all problems \Rightarrow Edge-Variant GNNs





Equivariance to Permutations is a Property of Graph Filters

▶ Permutation equivariance is a property of graph convolutions inherited to GNNs

- \Rightarrow Exploits data structure (internal symmetries of the graph)
- ▶ Why choose GNNs over graph convolutions?
 - \Rightarrow Q1: What is good about pointwise nonlinearities?
 - \Rightarrow Q2: What is wrong with linear graph convolutions?
- ▶ A2: They can be unstable to perturbations of the graph if we push their discriminative power
- ▶ A1: They make GNNs stable to perturbations while retaining discriminability
- ▶ These questions can be answered with an analysis in the **spectral domain**



Graph Convolutions in the Frequency Domain

• Graph convolution is a polynomial on the shift operator
$$\Rightarrow \mathbf{y} = \sum_{k=0}^{\infty} h_k \mathbf{S}^k \mathbf{x}$$

▶ Decompose operator as $\mathbf{S} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{\mathsf{H}}$ to write the spectral representation of the graph convolution

$$\mathbf{V}^{\mathsf{H}}\mathbf{y} = \mathbf{V}^{\mathsf{H}}\sum_{k=0}^{\infty}h_{k}(\mathbf{V}\mathbf{\Lambda}\mathbf{V}^{\mathsf{H}})^{k}\mathbf{x} \Rightarrow \tilde{\mathbf{y}} = \sum_{k=0}^{\infty}h_{k}\mathbf{\Lambda}^{k}\mathbf{\tilde{x}}$$

where we have used the graph Fourier transform (GFT) definitions $\mathbf{\tilde{x}}=\mathbf{V}^{\mathsf{H}}\mathbf{x}$ and $\mathbf{\tilde{y}}=\mathbf{V}^{\mathsf{H}}\mathbf{y}$

- Graph convolution is a pointwise operation in the spectral domain $\Rightarrow \tilde{y}_i = \tilde{h}(\lambda_i)\tilde{x}_i$
 - \Rightarrow Determined by the (graph) frequency response $\Rightarrow \sum_{k=0}^{\infty} h_k \lambda_i^k = \tilde{h}(\lambda_i)$



Graph Frequency Response

• We can reinterpret the frequency response as a polynomial on continuous $\lambda \Rightarrow \tilde{h}(\lambda) = \sum_{k=0}^{\infty} h_k \lambda^k$



Frequency response is the same no matter the graph \Rightarrow It's instantiated on its particular spectrum



Restricting the Class of Allowable Filters

Let $h(\lambda)$ be the frequency response of filter **H**. We say **H** is integral Lipschitz if $|\lambda h'(\lambda)| \leq C$



▶ Integral Lipschitz filters have to be wide for large λ ; but they can be thin for low λ



$Stability \ of \ Graph \ Neural \ Networks \ with \ Integral \ Lipschitz \ Filter S^{\rm raph \ Neural \ Networks}_{\ fgama@berkeley.edu}$

▶ Relative distance between S and $\hat{S} \Rightarrow$ Smallest matrix E that maps S into a permutation of \hat{S}

$$\mathcal{E} = \left\{ \mathbf{E} : \mathbf{P}^{\mathsf{T}} \hat{\mathbf{S}} \mathbf{P} = \mathbf{S} + \mathbf{E}^{\mathsf{T}} \mathbf{S} + \mathbf{S} \mathbf{E} \right\} \quad \Rightarrow \quad d(\mathbf{S}, \hat{\mathbf{S}}) = \min_{\mathbf{E} \in \mathcal{E}} \left\| \mathbf{E} \right\| \quad \approx \frac{\|\mathbf{S} - \mathbf{S}\|}{\|\mathbf{S}\|}$$

Theorem

Consider a GNN with L layers having integral Lipschitz filter \mathbf{H}_{ℓ} with constant C. Graphs S and $\hat{\mathbf{S}}$ satisfy $d(\mathbf{S}, \hat{\mathbf{S}}) \leq \varepsilon/2$. The matrix E that achieves minimum distance satisfies $\|\mathbf{E}/\|\mathbf{E}\| - \mathbf{I}\| \leq \varepsilon$. It holds that for all signals \mathbf{x}

$$\min_{\mathbf{P}\in\mathcal{P}} \| \boldsymbol{\Phi}(\mathbf{x}; \hat{\mathbf{S}}, \mathcal{H}) - \mathbf{P}^{\top} \boldsymbol{\Phi}(\mathbf{x}; \mathbf{S}, \mathcal{H}) \| \leq CL \, \varepsilon + \mathcal{O}(\varepsilon^2)$$

▶ GNNs can be made stable to graph perturbations if filters are integral Lipschitz





Insights: Edge Dilations

- ▶ Obtain valuable inisights \Rightarrow Consider particular case of edge dilation $\Rightarrow \hat{\mathbf{S}} = (1 + \varepsilon)\mathbf{S}$
- ► An edge dilation just produces a spectrum dilation $\Rightarrow \hat{\lambda}_i = (1 + \varepsilon)\lambda_i$, $\mathbf{E} = (\varepsilon/2)\mathbf{I}$



b Small deformations may result in large filter variations for large λ if filter is not integral Lipschitz



Insights: Edge Dilations

- ▶ Obtain valuable inisights \Rightarrow Consider particular case of edge dilation $\Rightarrow \hat{\mathbf{S}} = (1 + \varepsilon)\mathbf{S}$
- ► An edge dilation just produces a spectrum dilation $\Rightarrow \hat{\lambda}_i = (1 + \varepsilon)\lambda_i$, $\mathbf{E} = (\varepsilon/2)\mathbf{I}$



▶ Integral Lipschitz is always stable \Rightarrow Eigenvalue does not move or filter does not move



Discriminative Graph Filter Banks are Unstable

- ▶ Q2: What is wrong with linear graph convolutions?
- ▶ Cannot be simultaneously stable to deformations and discriminate features at large eigenvalues



▶ Limits their value in machine learning problems where features at large eigenvalues are important



Discriminative Graph Filter Banks are Unstable

- ▶ Q2: What is wrong with linear graph convolutions?
- ▶ Cannot be simultaneously stable to deformations and discriminate features at large eigenvalues



▶ Limits their value in machine learning problems where features at large eigenvalues are important



Nonlinearities Create Low Frequency Components

Graph Neural Networks fgama@berkeley.edu

- ▶ Q1: What is good about pointwise nonlinearities?
- ▶ Preserve permutation equivariance while generating low graph frequency components
 - \Rightarrow Which we can discriminate with stable filters



Spectrum of rectified graph signal $\mathbf{x}_{relu} = max(\mathbf{x}, 0)$

▶ The nonlinearity demodulates. It creates low frequency content that is stable



Nonlinearities Create Low Frequency Components

- ▶ Q1: What is good about pointwise nonlinearities?
- ▶ Preserve permutation equivariance while generating low graph frequency components
 - \Rightarrow Which we can discriminate with stable filters

GNNs are **stable** and **selective** information processing architectures

▶ The nonlinearity demodulates. It creates low frequency content that is stable



Distributed Control: Robotics

Graph Neural Networks fgama@berkeley.edu

Graph Neural Networks

Equivariance and Stability Properties

Robotics

Conclusions



Distributed Coordination of a Robot Swarm

▶ We want the team to coordinate on their individual velocities without colliding with each other

- ▶ This is a very easy problem to solve if we allow for centralized coordination $\Rightarrow \mathbf{u}_i = \sum_{i=1}^N \mathbf{v}_i$
- ▶ But it is very difficult to solve if we do do not allow for centralized coordination \Rightarrow $\mathbf{u}_i = \dots$

Tolstaya, Gama, Paulos, Pappas, Kumar, Ribeiro, "Learning Decentralized Controllers for Robot Swarms with Graph Neural Networks", CoRL 2019



Information Structure on Distributed Systems

Graph Neural Networks fgama@berkeley.edu

▶ The challenge in designing behaviors for distributed systems is the partial information structure



Tolstaya, Gama, Paulos, Pappas, Kumar, Ribeiro, "Learning Decentralized Controllers for Robot Swarms with Graph Neural Networks", CoRL 2019



- Optimal centralized actions act directly on all states $\pi^{\star}(\mathbf{x}(t))$ and can be readily computed
- ▶ Distributed actions can only depend on information history $\Rightarrow \mathcal{X}_i(t) = \bigcup_{k=0}^{K-1} \left\{ \mathbf{x}_j(t-k) : j \in \mathcal{N}_i^k \right\}$
 - \Rightarrow Optimal distributed actions $\mathbf{u}(t) = \Phi^{\star}(\mathcal{X}_i(t); \mathcal{G})$ are famously difficult to find [Witsenhausen '68]
 - \Rightarrow When optimal solutions are out of reach we resort to heuristics \Rightarrow data driven heuristics

Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Imitation Learning of Distributed Actions

- Parametrize $\Phi(\mathcal{X}_i(t); \mathcal{G})$ with a graph neural network $\Phi(\mathbf{x}(t); \mathbf{S}(t), \mathcal{H})$ \Rightarrow Adopt learning architecture that naturally processes distributed partial information $\mathcal{X}_i(t)$
- Train $\Phi(\mathbf{x}(t); \mathbf{S}(t), \mathcal{H})$ by using **imitation learning**
- Optimal centralized policy $\pi^{\star}(\mathbf{x}(t))$ can be computed during training time
- Find the parameters \mathcal{H} that make $\Phi(\mathbf{x}(t); \mathbf{S}(t), \mathcal{H})$ closer to $\pi^{\star}(\mathbf{x}(t))$

$$\mathcal{H}^{\star} = \operatorname*{arg\,min}_{\mathcal{H}} \mathbb{E}_{\pi^{\star}} \Big[\mathsf{J} \Big(\Phi \big(\mathbf{x}(t); \mathbf{S}(t), \mathcal{H} \big), \pi^{\star} (\mathbf{x}(t)) \Big) \Big]$$

Centralized policy required at train time but not at test time



GNN Properties Play a Key Role

► GNN controllers $\mathbf{u}(t) = \Phi(\mathbf{x}(t); \mathbf{S}(t), \mathcal{H})$ respect the partial information structure $\mathcal{X}_i(t)$

 \Rightarrow Graph filters $\mathbf{H}(\mathcal{X}_i(t)) = \sum_{k=0}^{K-1} h_k \mathbf{S}(t) \cdots \mathbf{S}(t-(k-1)) \mathbf{x}(t-k)$

Permutation equivariance

- \Rightarrow If two agents observe the same input
- \Rightarrow Their k-hop neighbors observe the same inputs
- \Rightarrow And the local neighborhood structures of the graph are the same
- ▶ Then the output of the control policy is the same at both nodes
- ▶ Stability \Rightarrow If graphs are similar, GNN outputs will be similar
- ▶ These properties are a necessity for offline training



Offline Training vs Online Execution

- ▶ If we want to train offline and execute online we can't assume the graph is the same
- ▶ Train offline on a graph like this



And execute online on a graph like this



- ▶ Graph convolutions ⇒ Learn parameters \mathcal{H} ⇒ Run on any S: $\mathbf{H}(\mathbf{S}) = \sum_{k=0}^{K-1} h_k \mathbf{S}^k$
- ▶ Permutation equivariance \Rightarrow Exploit symmetries, reuse data
- Stability \Rightarrow Similar graphs yield similar outputs

Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Team Dynamics

- Team of N agents with positions $\mathbf{r}_i(t)$, velocities $\mathbf{v}_i(t)$ and acceleration $\mathbf{u}_i(t)$
- ▶ System dynamics \Rightarrow Constant acceleration during each interval T_s

$$\begin{cases} \mathbf{r}_i(t+1) = \mathbf{u}_i(t)T_s^2/2 + \mathbf{v}_i(t)T_s + \mathbf{r}_i(t) \\ \mathbf{v}_i(t+1) = \mathbf{u}_i(t)T_s + \mathbf{v}_i(t) \end{cases}$$

- ▶ Control acceleration $\mathbf{u}_i(t) \Rightarrow$ Design sequence of acceleration values $\{\mathbf{u}_i(t)\}$
- ▶ Proof of concept \Rightarrow More realistic scenarios are available [Tolstaya et al '19, Li et al '20, Hu et al '21]

Tolstaya, Gama, Paulos, Pappas, Kumar, Ribeiro, "Learning Decentralized Controllers for Robot Swarms with GNNs", CoRL 2019. Li, Gama, Ribeiro, Prorok, "Graph Neural Networks for Decentralized Multi-Robot Path Planning", IRoS 2020. u, Gama, Wang, Ribeiro, Sadler "VGAI: A Vision-Based Decentralized Controller Learning Framework for Robot Swarms", ICASSP 2021 Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Problem Statement

- Coordinate the velocities $\mathbf{v}_i(t)$ of all agents to be the same while avoiding collisions
- ▶ Measure the cost by computing the velocity variation across the team for the entire trajectory

$$\mathsf{J}[\{\mathbf{v}_i(t)\}] = \frac{1}{N} \sum_t \sum_{i=1}^N ||\mathbf{v}_i(t) - \bar{\mathbf{v}}_j(t)||^2 , \ \bar{\mathbf{v}}_j = \frac{1}{N} \sum_{j=1}^N \mathbf{v}_j(t)$$

► Control accelerations $\mathbf{u}_i(t)$ such that $\min_{\mathbf{u}_i(t), t \ge 0} \mathsf{J}[\{\mathbf{v}_i(t)\}] \Rightarrow \text{Velocity } \mathbf{v}_i(t+1) = \mathbf{u}_i(t)T_s + \mathbf{v}_i(t)$



Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Communication Network

- ▶ Decentralized setting \Rightarrow Agents can only communicate if $\|\mathbf{r}_i(t) \mathbf{r}_j(t)\| \le R$
- ▶ Communication graph $\mathbf{S}(t) \Rightarrow \text{Edge } [\mathbf{S}(t)]_{ij}$ if agents can communicate
- ▶ Use GNNs that respect the communicaton graph \Rightarrow Local and distributed operations

$$\mathbf{U}(\mathcal{X}_i(t),\mathcal{H}) = \mathbf{\Phi}(\mathbf{X}(t);\mathbf{S}(t),\mathcal{H}) , \ \mathbf{H}(\mathcal{X}_i(t)) = \sum_{k=0}^{K-1} h_k \ \mathbf{S}(t) \cdots \mathbf{S}(t-(k-1))\mathbf{x}(t-k)$$



Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Imitation Learning

► Train GNN $\Phi(\mathbf{X}(t); \mathbf{S}(t), \mathcal{H})$ by using imitation learning \Rightarrow Find \mathcal{H} that make Φ closer to $\mathbf{U}^{\star}(t)$

$$\mathcal{H}^{\star} = \operatorname*{arg\,min}_{\mathcal{H}} \mathbb{E}_{\mathbf{U}^{\star}} \Big[\| \Phi(\mathbf{X}(t); \mathbf{S}(t), \mathcal{H}) - \mathbf{U}^{\star}(t) \| \Big]$$

▶ The optimal centralized solution $\mathbf{U}^{\star}(t)$ that avoids collisions is given by

$$\mathbf{u}_{i}^{\star}(t) = -\sum_{j=1}^{N} \left(\mathbf{v}_{i}(t) - \mathbf{v}_{j}(t) \right) \underbrace{-\sum_{j=1}^{N} \nabla_{\mathbf{r}_{i}(t)} P(\mathbf{r}_{i}(t), \mathbf{r}_{j}(t))}_{\text{collision avoidance}}$$

► The state $\mathbf{x}_i(t)$ for each agent is set to $\nabla_{\mathbf{r}_i(t)} P(\mathbf{r}_i(t), \mathbf{r}_j(t)) \Rightarrow$ Local to each agent

 \blacktriangleright Centralized controller \mathbf{U}^{\star} is required at train time **but not** at testing time



Flocking: Transferring



Offline optimal trajectory

Online learned trajectory

Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Flocking: Scalability



Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Flocking: Transferring at Scale

▶ Train on 50 agents. Test on 100 agents.

Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



- ► GNNs learn controllers to drive a team to fly at the same velocity while avoiding collisions ⇒ Control actions taken by each agent based only on outdated neighboring information
- ▶ Transfer \Rightarrow The communication network changes from time to time \Rightarrow GNNs work
- ▶ Scale \Rightarrow Learn to control teams of increasing number of agents \Rightarrow GNNs work
- ▶ Transfer at scale \Rightarrow Train on a small team, test on a big team \Rightarrow GNNs work

Tolstaya, Gama, Paulos, Pappas, Kumar, Ribeiro, "Learning Decentralized Controllers for Robot Swarms with Graph Neural Networks", CoRL 2019 Gama, Li, Tolstaya, Prorok, Ribeiro, "Decentralized Control with Graph Neural Networks", arxiv.org/abs/2012.14906



Conclusions

Graph Neural Networks fgama@berkeley.edu

Graph Neural Networks

Equivariance and Stability Properties

Robotics

Conclusions



- $\blacktriangleright \ {\rm Successful \ learning \ on \ graphs \ } \Rightarrow {\rm Scalability, \ exploit \ data \ structure, \ distributed \ implementation}$
- ▶ Graph neural networks (GNNs) \Rightarrow Graph convolutions followed by pointwise nonlinearities
- ▶ GNNs are permutation equivariant and stable to changes in the graph \Rightarrow Scale, transfer
- ▶ Graph convolutions are either stable or selective, but cannot be both
- \blacktriangleright Nonlinearities \Rightarrow GNNs are both stable and selective information processing architectures
- ► Learning decentralized controllers ⇒ Distributed systems ⇒ Partial information structure ⇒ Parametrize controller with GNN ⇒ Naturally adapts to partial information structure
- ▶ Flocking \Rightarrow Coordinate a team of robots to fly at the same velocities \Rightarrow Transfers at scale

Thank you!

