

MotifNet: a Motif-based GCNN for Directed Graphs

Federico Monti^{1,2}, Karl Otness⁵, Michael M. Bronstein^{1,2,3,4,5}



USI
Switzerland



Fabula AI
UK



TAU
Israel



Intel
Israel



Harvard
USA

MotifNet: a Motif-based GCNN for Directed Graphs

Federico Monti^{1,2}, Karl Otness⁵, Michael M. Bronstein^{1,2,3,4,5}



USI
Switzerland



Fabula AI
UK



TAU
Israel



Intel
Israel



Harvard
USA


Different formulations of CNN on graphs



Spectral domain^{1,2,3}



Spatial domain^{4,5,6}



Embedding domain^{7,8}

¹Bruna et al. 2014; ²Henaff, Bruna, LeCun 2015; ³Defferrard, Bresson, Vandergheynst 2016; ⁴Masci et al. 2015; ⁵Boscaini et al. 2016; ⁶Monti et al. 2017; ⁷Sinha, Bai, Ramani 2016; ⁸Maron et al. 2017

Directed graphs




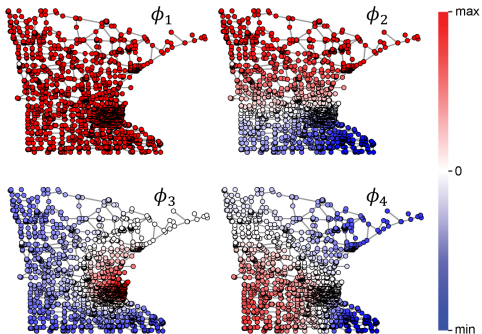




Figure: Monti, Otness, Bronstein 2018; data: Bojchevski, Günnemann 2017

Graph Convolutional Neural Networks

Laplacian eigenfunctions

$$\Delta e^{j\omega x} = \frac{d^2}{dx^2} e^{j\omega x} = -\omega^2 e^{j\omega x}$$

$$\Delta = \mathbf{D} - \mathbf{W} = \Phi \Lambda \Phi^T$$



Convolution Theorem

- Given two functions $f, h : Z \rightarrow \mathbb{R}$ their **convolution** is a function

$$(f \star h)(x) = \sum_{x'=-K/2}^{K/2} f(x-x')h(x') = \mathcal{F}^{-1}\{\mathcal{F}\{f\} \cdot \mathcal{F}\{h\}\}$$

Convolution Theorem

- Given two functions $f, h : Z \rightarrow \mathbb{R}$ their **convolution** is a function

$$(f \star h)(x) = \sum_{x'=-K/2}^{K/2} f(x-x')h(x') = \mathcal{F}^{-1}\{\mathcal{F}\{f\} \cdot \mathcal{F}\{h\}\}$$

- Given two functions $f, h : \mathcal{V} \rightarrow \mathbb{R}$ their **convolution** is a function




Convolution Theorem

- Given two functions $f, h : Z \rightarrow \mathbb{R}$ their **convolution** is a function

$$(f \star h)(x) = \sum_{x'=-K/2}^{K/2} f(x-x')h(x') = \mathcal{F}^{-1}\{\mathcal{F}\{f\} \cdot \mathcal{F}\{h\}\}$$

- Given two functions $f, h : \mathcal{V} \rightarrow \mathbb{R}$ their **convolution** is a function




Convolution Theorem

- Given two functions $f, h : Z \rightarrow \mathbb{R}$ their **convolution** is a function

$$(f \star h)(x) = \sum_{x'=-K/2}^{K/2} f(x-x')h(x') = \mathcal{F}^{-1}\{\mathcal{F}\{f\} \cdot \mathcal{F}\{h\}\}$$

- Given two functions $f, h : \mathcal{V} \rightarrow \mathbb{R}$ their **convolution** is a function



Spectral graph CNN

Convolutional layer expressed in the **spectral domain**

$$\mathbf{g}_l = \xi \left(\sum_{l'=1}^p \Phi \begin{bmatrix} \hat{h}_1^{(l,l')} \\ \vdots \\ \hat{h}_N^{(l,l')} \end{bmatrix} \Phi^\top \mathbf{f}_{l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

where q is the number of features in output and p in input.

Spectral graph CNN

Convolutional layer expressed in the **spectral domain**

$$\mathbf{g}_l = \xi \left(\sum_{l'=1}^p \Phi \begin{bmatrix} \hat{h}_1^{(l,l')} \\ \vdots \\ \hat{h}_N^{(l,l')} \end{bmatrix} \Phi^\top \mathbf{f}_{l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

where q is the number of features in output and p in input.

- ☹ $\mathcal{O}(N)$ parameters per layer
- ☹ $\mathcal{O}(N^2)$ computation of forward and inverse Fourier transforms Φ^\top, Φ (no FFT on graphs)
- ☹ No guarantee of spatial localization of filters

Spectral graph CNN with polynomial filters

Represent spectral transfer function as a **polynomial** of order r

$$\tau_{\alpha}(\lambda) = \sum_{j=0}^r \alpha_j \lambda^j$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^{\top}$ is the vector of filter parameters

Spectral graph CNN with polynomial filters

Represent spectral transfer function as a **Chebyshev polynomial** of order r

$$\tau_{\alpha}(\tilde{\lambda}) = \sum_{j=0}^r \alpha_j T_j(\tilde{\lambda})$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^T$ is the vector of filter parameters,

$$T_j(\tilde{\lambda}) = 2\tilde{\lambda}T_{j-1}(\tilde{\lambda}) - T_{j-2}(\tilde{\lambda}) \quad T_0(\tilde{\lambda}) = 1, \quad T_1(\tilde{\lambda}) = \tilde{\lambda}$$

and $-1 \leq \tilde{\lambda} \leq 1$ is normalized frequency.

Spectral graph CNN with polynomial filters

Represent spectral transfer function as a **Chebyshev polynomial** of order r

$$\tau_{\alpha}(\tilde{\lambda}) = \sum_{j=0}^r \alpha_j T_j(\tilde{\lambda})$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^T$ is the vector of filter parameters,

$$T_j(\tilde{\lambda}) = 2\tilde{\lambda}T_{j-1}(\tilde{\lambda}) - T_{j-2}(\tilde{\lambda}) \quad T_0(\tilde{\lambda}) = 1, \quad T_1(\tilde{\lambda}) = \tilde{\lambda}$$

and $-1 \leq \tilde{\lambda} \leq 1$ is normalized frequency.

Application of the filter to scaled Laplacian $\tilde{\Delta} = \mathbf{D}^{-0.5} \mathbf{A} \mathbf{D}^{-0.5} - \mathbf{I} = \Phi \tilde{\Lambda} \Phi^T$:

$$\mathbf{g} = \Phi \left(\sum_{j=0}^r \alpha_j T_j(\tilde{\Lambda}) \right) \Phi^T \mathbf{f} = \sum_{j=0}^r \alpha_j T_j(\tilde{\Delta}) \mathbf{f}$$

Spectral graph CNN with polynomial filters

Represent spectral transfer function as a **Chebyshev polynomial** of order r

$$\tau_{\alpha}(\tilde{\lambda}) = \sum_{j=0}^r \alpha_j T_j(\tilde{\lambda})$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^\top$ is the vector of filter parameters,

$$T_j(\tilde{\lambda}) = 2\tilde{\lambda}T_{j-1}(\tilde{\lambda}) - T_{j-2}(\tilde{\lambda}) \quad T_0(\tilde{\lambda}) = 1, \quad T_1(\tilde{\lambda}) = \tilde{\lambda}$$

and $-1 \leq \tilde{\lambda} \leq 1$ is normalized frequency.

Application of the filter to scaled Laplacian $\tilde{\Delta} = \mathbf{D}^{-0.5} \mathbf{\Delta} \mathbf{D}^{-0.5} - \mathbf{I} = \Phi \tilde{\Lambda} \Phi^T$:

$$\mathbf{g} = \Phi \left(\sum_{j=0}^r \alpha_j T_j(\tilde{\Lambda}) \right) \Phi^T \mathbf{f} = \sum_{j=0}^r \alpha_j T_j(\tilde{\Delta}) \mathbf{f}$$


- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Filters have guaranteed r -hops support
- ☺ No explicit computation of $\Phi^\top, \Phi \Rightarrow \mathcal{O}(r|\mathcal{E}|)$ computational complexity

Multi-Graph CNN

Multi-graph spectral convolutional layer

$$\mathbf{Y}_l = \xi \left(\sum_{l'=1}^p \sum_{j,j'=0}^r \theta_{jj'ul'} T_j(\tilde{\Delta}_r) \mathbf{X}_{l'} T_{j'}(\tilde{\Delta}_c) \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

applied to p input channels ($m \times n$ matrices $\mathbf{X}_1, \dots, \mathbf{X}_p$) and producing q output channels ($m \times n$ matrices $\mathbf{Y}_1, \dots, \mathbf{Y}_q$)



Multi-Graph CNN

Multi-graph spectral convolutional layer


$$\mathbf{Y}_l = \xi \left(\sum_{l'=1}^p \sum_{j,j'=0}^r \theta_{jj'l'l'} T_j(\tilde{\Delta}_r) \mathbf{X}_{l'} T_{j'}(\tilde{\Delta}_c) \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

applied to p input channels ($m \times n$ matrices $\mathbf{X}_1, \dots, \mathbf{X}_p$) and producing q output channels ($m \times n$ matrices $\mathbf{Y}_1, \dots, \mathbf{Y}_q$)

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Filters have guaranteed r -hops support on both graphs
- ☺ $\mathcal{O}(nm)$ computational complexity

Dealing with Directed Graphs

Directed graphs




Directed graph

Asymmetric adjacency matrix


$$\begin{bmatrix} . & 1 & 1 & 1 & 1 & 1 & . & . & . & . \\ 1 & . & 1 & 1 & 1 & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & 1 & . & . & . & . & 1 & . & . & . \\ . & 1 & . & . & . & 1 & . & . & . & . \\ . & 1 & . & . & . & 1 & . & . & 1 & 1 \\ . & . & . & . & . & 1 & 1 & 1 & . & 1 \\ . & . & . & . & . & . & . & . & . & . \end{bmatrix}$$

Motif-based graph analysis




Directed graph

Asymmetric adjacency matrix

$$\begin{bmatrix} . & 1 & 1 & 1 & 1 & 1 & . & . & . & . \\ 1 & . & 1 & 1 & 1 & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & 1 & . & . & . & . & 1 & . & . & . \\ . & 1 & . & . & . & 1 & . & . & . & . \\ . & 1 & . & . & . & 1 & . & . & 1 & 1 \\ . & . & . & . & . & 1 & 1 & 1 & . & 1 \\ . & . & . & . & . & . & . & . & . & . \end{bmatrix}$$


Graph motifs


Motif-based graph analysis



Directed graph


Motif adjacency matrix for M_7

.	3	1	1	1
3	.	1	1	1	1	1	.	.	.
1	1
1	1
1	1
.	1	1	1	1	.
.	1	.	.	.	1	.	.	2	1
.	1	.	2	.	1
.	1	1	.



Graph motifs

Motif Laplacians



Undirected weighted graph

Motif adjacency matrix for M_7

$$\begin{bmatrix} . & 3 & 1 & 1 & 1 & . & . & . & . & . \\ 3 & . & 1 & 1 & 1 & 1 & 1 & . & . & . \\ 1 & 1 & . & . & . & . & . & . & . & . \\ 1 & 1 & . & . & . & . & . & . & . & . \\ 1 & 1 & . & . & . & . & . & . & . & . \\ . & 1 & . & . & . & . & 1 & 1 & 1 & . \\ . & 1 & . & . & . & 1 & . & . & . & . \\ . & . & . & . & . & 1 & . & . & 2 & 1 \\ . & . & . & . & . & 1 & . & 2 & . & 1 \\ . & . & . & . & . & . & . & 1 & 1 & . \end{bmatrix}$$

Motif Laplacian for motif $k = 1, \dots, K$

$$\tilde{\Delta}_k = \mathbf{I} - \tilde{\mathbf{D}}_k^{-1/2} \tilde{\mathbf{W}}_k \tilde{\mathbf{D}}_k^{-1/2}$$

MotifNet: multivariate polynomial

Apply K -variate polynomial of order r to the motif Laplacians

$$\tau_{\boldsymbol{\theta}}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \theta_0 \mathbf{I} + \sum_{j=1}^r \sum_{k_1, \dots, k_j \in \{1, \dots, K\}} \theta_{k_1, \dots, k_j} \tilde{\Delta}_{k_1} \cdots \tilde{\Delta}_{k_j}$$

where $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_{K, \dots, K})$ is the vector of filter parameters

MotifNet: multivariate polynomial

Apply K -variate polynomial of order r to the motif Laplacians

$$\tau_{\boldsymbol{\theta}}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \theta_0 \mathbf{I} + \sum_{j=1}^r \sum_{k_1, \dots, k_j \in \{1, \dots, K\}} \theta_{k_1, \dots, k_j} \tilde{\Delta}_{k_1} \cdots \tilde{\Delta}_{k_j}$$

where $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_{K, \dots, K})$ is the vector of filter parameters

😊 Explicitly accounts for directed graph structures

MotifNet: multivariate polynomial

Apply K -variate polynomial of order r to the motif Laplacians

$$\tau_{\boldsymbol{\theta}}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \theta_0 \mathbf{I} + \sum_{j=1}^r \sum_{k_1, \dots, k_j \in \{1, \dots, K\}} \theta_{k_1, \dots, k_j} \tilde{\Delta}_{k_1} \cdots \tilde{\Delta}_{k_j}$$

where $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_{K, \dots, K})$ is the vector of filter parameters

- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels

MotifNet: multivariate polynomial

Apply K -variate polynomial of order r to the motif Laplacians

$$\tau_{\boldsymbol{\theta}}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \theta_0 \mathbf{I} + \sum_{j=1}^r \sum_{k_1, \dots, k_j \in \{1, \dots, K\}} \theta_{k_1, \dots, k_j} \tilde{\Delta}_{k_1} \cdots \tilde{\Delta}_{k_j}$$

where $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_{K, \dots, K})$ is the vector of filter parameters

- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels
- ☺ Filters have guaranteed r -hops support

MotifNet: multivariate polynomial

Apply K -variate polynomial of order r to the motif Laplacians

$$\tau_{\theta}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \theta_0 \mathbf{I} + \sum_{j=1}^r \sum_{k_1, \dots, k_j \in \{1, \dots, K\}} \theta_{k_1, \dots, k_j} \tilde{\Delta}_{k_1} \cdots \tilde{\Delta}_{k_j}$$

where $\theta = (\theta_0, \theta_1, \dots, \theta_{K, \dots, K})$ is the vector of filter parameters

- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels
- ☺ Filters have guaranteed r -hops support
- ☺ $\mathcal{O}(1)$ parameters per layer

MotifNet: multivariate polynomial

Apply K -variate polynomial of order r to the motif Laplacians

$$\tau_{\theta}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \theta_0 \mathbf{I} + \sum_{j=1}^r \sum_{k_1, \dots, k_j \in \{1, \dots, K\}} \theta_{k_1, \dots, k_j} \tilde{\Delta}_{k_1} \cdots \tilde{\Delta}_{k_j}$$

where $\theta = (\theta_0, \theta_1, \dots, \theta_{K, \dots, K})$ is the vector of filter parameters

- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels
- ☺ Filters have guaranteed r -hops support
- ☹ $\frac{1-K^{r+1}}{1-K}$ parameters per layer, intractable in practice

MotifNet: multivariate polynomial

Apply K -variate polynomial of order r to the motif Laplacians

$$\tau_{\theta}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \theta_0 \mathbf{I} + \sum_{j=1}^r \sum_{k_1, \dots, k_j \in \{1, \dots, K\}} \theta_{k_1, \dots, k_j} \tilde{\Delta}_{k_1} \cdots \tilde{\Delta}_{k_j}$$

where $\theta = (\theta_0, \theta_1, \dots, \theta_{K, \dots, K})$ is the vector of filter parameters

- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels
- ☺ Filters have guaranteed r -hops support
- ☹ $\frac{1-K^{r+1}}{1-K}$ parameters per layer, intractable in practice
- ☹ $\mathcal{O}(n \frac{1-K^{r+1}}{1-K})$ computational complexity

MotifNet: simplified multivariate polynomial

Two possible simplifications:

- We consider only $K = 2$ simple motifs corresponding to incoming and outgoing edges (MotifNet-d):

$$\tau_{\theta} = \theta_0 \mathbf{I} + \theta_1 \tilde{\Delta}_1 + \theta_2 \tilde{\Delta}_2 + \theta_{11} \tilde{\Delta}_1^2 + \dots + \theta_{22} \tilde{\Delta}_2^2 + \dots \quad (1)$$

MotifNet: simplified multivariate polynomial

Two possible simplifications:

- We consider only $K = 2$ simple motifs corresponding to incoming and outgoing edges (MotifNet-d):

$$\tau_{\theta} = \theta_0 \mathbf{I} + \theta_1 \tilde{\Delta}_1 + \theta_2 \tilde{\Delta}_2 + \theta_{11} \tilde{\Delta}_1^2 + \dots + \theta_{22} \tilde{\Delta}_2^2 + \dots \quad (1)$$

- We consider a simplified version of multivariate polynomials (MotifNet-m) defined recursively as:

MotifNet: simplified multivariate polynomial

Two possible simplifications:

- We consider only $K = 2$ simple motifs corresponding to incoming and outgoing edges (MotifNet-d):

$$\tau_{\theta} = \theta_0 \mathbf{I} + \theta_1 \tilde{\Delta}_1 + \theta_2 \tilde{\Delta}_2 + \theta_{11} \tilde{\Delta}_1^2 + \dots + \theta_{22} \tilde{\Delta}_2^2 + \dots \quad (1)$$

- We consider a simplified version of multivariate polynomials (MotifNet-m) defined recursively as:

$$\mathbf{P}_{\Theta}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{j=0}^r \theta_j \mathbf{P}_j; \quad (2)$$

$$\mathbf{P}_j(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{k=1}^K \alpha_{k,j} \tilde{\Delta}_k \mathbf{P}_{j-1}, \quad j = 1, \dots, p$$

$$\mathbf{P}_0 = \mathbf{I},$$

$$\sum_{k=1}^K \alpha_{k,j} = 1 \quad j = 1, \dots, p$$

MotifNet-m: simplified multivariate polynomial

Apply **attention-based recursive polynomial** of order r

$$\mathbf{P}_{\Theta}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{j=0}^r \theta_j \mathbf{P}_j;$$
$$\mathbf{P}_j(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{k=1}^K \alpha_{k,j} \tilde{\Delta}_k \mathbf{P}_{j-1}, \quad \mathbf{P}_0 = \mathbf{I},$$

where $\Theta = (\theta_0, \dots, \theta_r, \alpha_{1,1}, \dots, \alpha_{K,r})$ is the vector of filter parameters

- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels
- ☺ Filters have guaranteed r -hops support

MotifNet-m: simplified multivariate polynomial

Apply **attention-based recursive polynomial** of order r

$$\mathbf{P}_{\Theta}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{j=0}^r \theta_j \mathbf{P}_j;$$
$$\mathbf{P}_j(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{k=1}^K \alpha_{k,j} \tilde{\Delta}_k \mathbf{P}_{j-1}, \quad \mathbf{P}_0 = \mathbf{I},$$

where $\Theta = (\theta_0, \dots, \theta_r, \alpha_{1,1}, \dots, \alpha_{K,r})$ is the vector of filter parameters

- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels
- ☺ Filters have guaranteed r -hops support
- ☺ $Kr + 1$ parameters per layer

MotifNet-m: simplified multivariate polynomial

Apply **attention-based recursive polynomial** of order r


$$\mathbf{P}_{\Theta}(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{j=0}^r \theta_j \mathbf{P}_j;$$
$$\mathbf{P}_j(\tilde{\Delta}_1, \dots, \tilde{\Delta}_K) = \sum_{k=1}^K \alpha_{k,j} \tilde{\Delta}_k \mathbf{P}_{j-1}, \quad \mathbf{P}_0 = \mathbf{I},$$

where $\Theta = (\theta_0, \dots, \theta_r, \alpha_{1,1}, \dots, \alpha_{K,r})$ is the vector of filter parameters


- ☺ Explicitly accounts for directed graph structures
- ☺ Anisotropic kernels
- ☺ Filters have guaranteed r -hops support
- ☺ $Kr + 1$ parameters per layer
- ☺ $\mathcal{O}(nr)$ computational complexity

Directed CORA

- Citations network representing papers (vertices) and citations (directed edges)
- Goal: vertex-wise classification (paper topic)
- 19,793 documents, 70 classes
- Training/validation/test set: 10% samples (1979 vertices)




Example: directed citation networks (CORA)



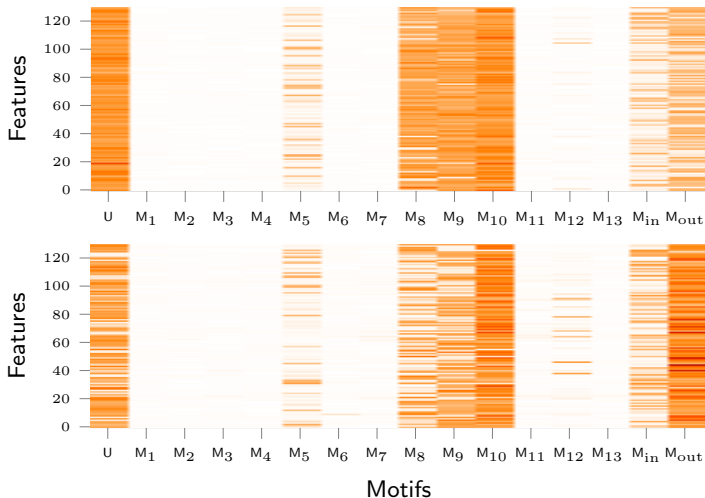
Classification accuracy obtained with ChebNet on undirected graph (blue) and MotifNet-m (orange).

Example: directed citation networks (CORA)



Classification accuracy obtained with ChebNet applied with directed adjacency matrix \mathbf{W} (blue) / \mathbf{W}^T (red), MotifNet-d (green) and MotifNet-m (orange).

Example: directed citation networks (CORA)



Attention scores α obtained with MotifNet-m ($r = 1$) from 1st and 2nd graph convolutional layers. Dark/bright colors represent high/low probabilities.

Conclusions

- We presented MotifNet, a motif-based GCN for directed graphs.
- Our solution realizes anisotropic filters on the provided domain generalizing and enriching classic undirected GCNs.
- Thanks to the proposed simplifications, MotifNet-m requires a number of parameters and a computational complexity comparable with the ones of ChebNet.
- Experimental evaluations on the directed version of CORA show the superior performance of our approach.

Thank You

Bibliography

M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, P. Vandergheynst, "Geometric deep learning: going beyond Euclidean data", arXiv:1611.08097, 2016. **First review paper of geometric deep learning**

M. Gori, G. Monfardini, F. Scarselli, "A new model for learning in graph domains", *Proc. IJCNN* 2005. **First neural net on graphs (GNN framework)**

F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, G. Monfardini, "The graph neural network model", *IEEE Trans. Neural Networks* 20(1):61–80, 2009.

J. Bruna, W. Zaremba, A. Szlam, Y. LeCun, "Spectral networks and locally connected networks on graphs", *Proc. ICML* 2014. **First Spectral CNN on graphs**

M. Henaff, J. Bruna, and Y. LeCun, "Deep convolutional networks on graph-structured data", arXiv:1506.05163, 2015. **Spectral CNN with smooth multipliers**

J. Masci, D. Boscaini, M. M. Bronstein, P. Vandergheynst, "Geodesic convolutional neural networks on Riemannian manifolds", arXiv:1501.06297v2, 2015. **First spatial CNN on manifolds (Geodesic CNN framework)**

Bibliography

M. Defferrard, X. Bresson, P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering”, *Proc. NIPS* 2016. **Spectral CNN with Chebychev polynomial filters (ChebNet)**

T. N. Kipf, M. Welling, “Semi-supervised classification with graph convolutional networks”, arXiv:1609.02907, 2016. **Graph convolutional network (GCN) framework, a simplification of ChebNet**

R. Levie, F. Monti, X. Bresson, M. M. Bronstein, “CayleyNets: Graph convolutional neural networks with complex rational spectral filters”, arXiv:1705.07664, 2017. **Spectral CNN with complex rational filters (CayleyNet)**

J. Atwood and D. Towsley, “Diffusion-convolutional neural networks”, arXiv:1511.02136v2, 2016. **Diffusion CNN framework**

F. Monti, K. Otness, M. M. Bronstein, “MotifNet: a motif-based Graph Convolutional Network for directed graphs”, arXiv:1802.01572, 2018. **MotifNet**

Bibliography

F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, M. M. Bronstein, "Geometric deep learning on graphs and manifolds using mixture model CNNs", *Proc. CVPR* 2017. **Mixture Model network (MoNet) framework**

D. Boscaini, J. Masci, S. Melzi, M. M. Bronstein, U. Castellani, P. Vandergheynst, "Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks", *Computer Graphics Forum* 34(5):13-23, 2015. **Localized Spectral CNN framework**

D. Boscaini, J. Masci, E. Rodolà, M. M. Bronstein, "Learning shape correspondence with anisotropic convolutional neural networks", *Proc. NIPS* 2016. **Anisotropic CNN framework**

F. Monti, M. M. Bronstein, X. Bresson, "Geometric matrix completion with recurrent multi-graph neural networks", arXiv:1704.06803, 22 April 2017. **CNNs on multiple graphs and application to matrix completion**

Bibliography

A. Sinha, J. Bai, K. Ramani, "Deep learning 3D shape surfaces using geometry images", Proc. ECCV 2016. **CNN on spherical authalic parametrization**

H. Maron, M. Galun, N. Aigerman, M. Trope, N. Dym, E. Yumer, V. G. Kim, Y. Lipman, "Convolutional neural networks on surfaces via seamless toric covers", Trans. Graphics 36(4), 2017. **CNN on planar flat-torus parametrization**

L. Yi, H. Su, X. Guo, L. Guibas, "SyncSpecCNN: Synchronized Spectral CNN for 3D Shape Segmentation", Proc. CVPR 2017. **Spectral transformer networks**

C. R. Qi, H. Su, K. Mo, L. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation", Proc. CVPR 2017. **PointNet**

Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, J. M. Solomon, "Dynamic Graph CNN for learning on point clouds", arXiv:1712.00268, 2018. **Edge convolution**

Bibliography

O. Litany, T. Remez, E. Rodolà, A. M. Bronstein, M. M. Bronstein, “Deep functional maps: Structured prediction for dense shape correspondence”, Proc. ICCV 2017. **Intrinsic structured prediction (FMNet)**

O. Litany, A. M. Bronstein, M. M. Bronstein, A. Makadia, “Deformable shape completion with graph convolutional autoencoders”, Proc. CVPR 2018. **Shape completion with intrinsic variational autoencoder**

Y. LeCun, B. Boser, J. S. Denker, R. E. Howard, W. Hubbard, L. D. Jackel, “Backpropagation applied to handwritten ZIP code recognition”, *Neural Computation* 1(4):541–551, 1989. **Classical Euclidean CNN**