

Autoregressive Models for Sequences of Graphs

IJCNN 2019

Daniele Zambon^{*†}, Daniele Grattarola^{*†}, Lorenzo Livi[‡], Cesare Alippi^{†§}

* Equal contribution

† Università della Svizzera italiana

‡ University of Manitoba & University of Exeter

§ Politecnico di Milano



Università
della
Svizzera
italiana

Sequences of graphs

A graph of N nodes is a tuple $g = (\mathcal{V}, \mathcal{E}) \in \mathcal{G}$ s.t.:

$$\mathcal{V} = \{v_i \in \mathbb{R}^F\}_{i=1, \dots, N} \quad \mathcal{E} = \{e_{ij} \in \mathbb{R}^S\}_{v_i, v_j \in \mathcal{V}}$$

Sequences of graphs

A graph of N nodes is a tuple $g = (\mathcal{V}, \mathcal{E}) \in \mathcal{G}$ s.t.:

$$\mathcal{V} = \{v_i \in \mathbb{R}^F\}_{i=1, \dots, N} \quad \mathcal{E} = \{e_{ij} \in \mathbb{R}^S\}_{v_i, v_j \in \mathcal{V}}$$

A graph stochastic process (GSP) is a sequence of random variables $\{g_t \in \mathcal{G}\}_{t=0, \dots, T}$




Sequences of graphs

A graph of N nodes is a tuple $g = (\mathcal{V}, \mathcal{E}) \in \mathcal{G}$ s.t.:

$$\mathcal{V} = \{v_i \in \mathbb{R}^F\}_{i=1, \dots, N} \quad \mathcal{E} = \{e_{ij} \in \mathbb{R}^S\}_{v_i, v_j \in \mathcal{V}}$$

A graph stochastic process (GSP) is a sequence of random variables $\{g_t \in \mathcal{G}\}_{t=0, \dots, T}$



Problem formulation: graph autoregression

Traditional AR model

$$f : \mathbb{R}^{p \times d} \rightarrow \mathbb{R}^d$$

$$x_{t+1} = f(\mathbf{x}_t^p) + \epsilon$$

$$f(\mathbf{x}_t^p) = \mathbb{E}[f(\mathbf{x}_t^p) + \epsilon]$$

$$\text{Var}[\epsilon] = \sigma^2 < \infty$$

Graph AR model

Problem formulation: graph autoregression

Traditional AR model

$$f : \mathbb{R}^{p \times d} \rightarrow \mathbb{R}^d$$

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t^p) + \epsilon$$

$$f(\mathbf{x}_t^p) = \mathbb{E}[f(\mathbf{x}_t^p) + \epsilon]$$

$$\text{Var}[\epsilon] = \sigma^2 < \infty$$

Graph AR model

$$\phi : \mathcal{G}^p \rightarrow \mathcal{G}$$

Problem formulation: graph autoregression

Traditional AR model

$$f : \mathbb{R}^{p \times d} \rightarrow \mathbb{R}^d$$

$$x_{t+1} = f(\mathbf{x}_t^p) + \epsilon$$

$$f(\mathbf{x}_t^p) = \mathbb{E}[f(\mathbf{x}_t^p) + \epsilon]$$

$$\text{Var}[\epsilon] = \sigma^2 < \infty$$

Graph AR model

$$\phi : \mathcal{G}^p \rightarrow \mathcal{G}$$

$$\mathbf{g}_{t+1} = H(\phi(\mathbf{g}_t^p), \eta)$$

Problem formulation: graph autoregression

Traditional AR model

$$f : \mathbb{R}^{p \times d} \rightarrow \mathbb{R}^d$$

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t^p) + \epsilon$$

$$f(\mathbf{x}_t^p) = \mathbb{E}[f(\mathbf{x}_t^p) + \epsilon]$$

$$\text{Var}[\epsilon] = \sigma^2 < \infty$$

Graph AR model

$$\phi : \mathcal{G}^p \rightarrow \mathcal{G}$$

$$\mathbf{g}_{t+1} = H(\phi(\mathbf{g}_t^p), \eta)$$

$$\phi(\mathbf{g}_t^p) \in \mathbb{E}_\eta^f[H(\phi(\mathbf{g}_t^p), \eta)]$$

Where:

$$\mathbb{E}^f[\mathbf{g}] := \arg \min_{\mathbf{g}' \in \mathcal{G}} \int_{\mathcal{G}} d(\mathbf{g}, \mathbf{g}')^2 dQ_{\mathbf{g}}(\mathbf{g})$$

Problem formulation: graph autoregression

Traditional AR model

$$f : \mathbb{R}^{p \times d} \rightarrow \mathbb{R}^d$$

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t^p) + \epsilon$$

$$f(\mathbf{x}_t^p) = \mathbb{E}[f(\mathbf{x}_t^p) + \epsilon]$$

$$\text{Var}[\epsilon] = \sigma^2 < \infty$$

Graph AR model

$$\phi : \mathcal{G}^p \rightarrow \mathcal{G}$$

$$\mathbf{g}_{t+1} = H(\phi(\mathbf{g}_t^p), \eta)$$

$$\phi(\mathbf{g}_t^p) \in \mathbb{E}_\eta^f[H(\phi(\mathbf{g}_t^p), \eta)]$$


$$\text{Var}^f[\eta] < \infty$$

Where:

$$\mathbb{E}^f[g] := \arg \min_{g' \in \mathcal{G}} \int_{\mathcal{G}} d(g, g')^2 dQ_g(g)$$

$$\text{Var}^f[g] := \min_{g' \in \mathcal{G}} \int_{\mathcal{G}} d(g, g')^2 dQ_g(g)$$


Implementing $\phi(\mathbf{g}_t^p)$



Implementing $\phi(\mathbf{g}_t^p)$

Graph embedding

$\phi_{emb}(\cdot; \theta_{emb}) : \mathcal{G} \rightarrow \mathbb{R}^l$, applied in parallel to each graph in the regressor.



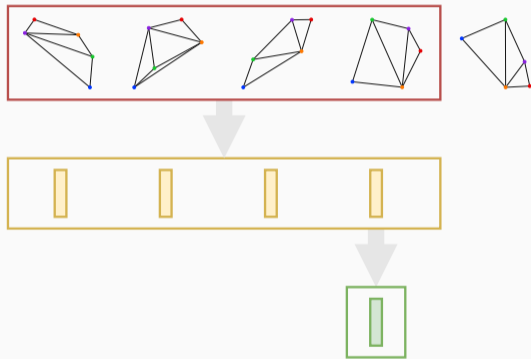
Implementing $\phi(\mathbf{g}_t^p)$

Graph embedding

$\phi_{emb}(\cdot; \theta_{emb}) : \mathcal{G} \rightarrow \mathbb{R}^l$, applied in parallel to each graph in the regressor.

Vector prediction

$\phi_{prd}(\cdot; \theta_{prd}) : \mathbb{R}^{k \times l} \rightarrow \mathbb{R}^l$, captures temporal dynamics in embedding space.



Implementing $\phi(\mathbf{g}_t^p)$

Graph embedding


$\phi_{emb}(\cdot; \theta_{emb}) : \mathcal{G} \rightarrow \mathbb{R}^l$, applied in parallel to each graph in the regressor.

Vector prediction

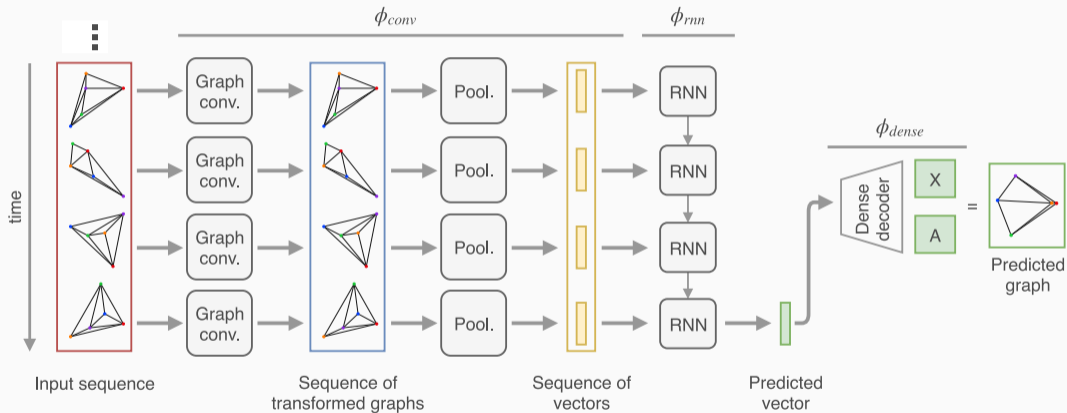
$\phi_{prd}(\cdot; \theta_{prd}) : \mathbb{R}^{k \times l} \rightarrow \mathbb{R}^l$, captures temporal dynamics in embedding space.

Graph decoder

$\phi_{dec}(\cdot; \theta_{dec}) : \mathbb{R}^l \rightarrow \mathcal{G}$, outputs graph from predicted vector.



Implementing $\phi(\mathbf{g}_t^p)$ with a GNN




Experiments: rotational model

Points moving in a 2D plane with delayed rotations:

$$f(\mathbf{x}_t^p) = R(\mathbf{x}_t^p) \cdot \mathbf{x}_t$$

$$R_n(\mathbf{x}_t^p) = \begin{bmatrix} \cos(\omega) & \sin(\omega) \\ -\sin(\omega) & \cos(\omega) \end{bmatrix}$$

$$\omega = c_n + \alpha \cos \left(\sum_{i=0}^{p-1} x_{t-i,2n-1} + x_{t-i,2n} \right)$$



Experiments: Partially Masked Linear Dynamical System (PMLDS)


Given a c -dimensional oscillating linear system:

$$x_{t+1} = Rx_t$$

$$x_t \in \mathbb{R}^c \text{ and } R \in \mathbb{R}^{c \times c}$$

we only observe the top $N \cdot F$ components.

This results in a dynamical system of order $p \propto (c - N \cdot F)$ [1].



Mean

Assumes GSP is stationary, predicts the mean graph: $\hat{g}_{t+1} = \mathbb{E}^f[g]$.

Baselines

Mean

Assumes GSP is stationary, predicts the mean graph: $\hat{g}_{t+1} = \mathbb{E}^f[g]$.

Martingale

Assumes GSP is a martingale, predicts the previous graph: $\hat{g}_{t+1} = g_t$.

Baselines

Mean

Assumes GSP is stationary, predicts the mean graph: $\hat{g}_{t+1} = \mathbb{E}^f[g]$.

Martingale

Assumes GSP is a martingale, predicts the previous graph: $\hat{g}_{t+1} = g_t$.

Moving average

Predicts the mean graph from the k preceding observations:

$$\hat{g}_{t+1} = \mathbb{E}^f[\mathbf{g}_t^k] = \arg \min_{g'} \sum_{g_i \in \mathbf{g}_t^k} d(g', g_i)^2.$$

Baselines

Mean

Assumes GSP is stationary, predicts the mean graph: $\hat{g}_{t+1} = \mathbb{E}^f[g]$.

Martingale

Assumes GSP is a martingale, predicts the previous graph: $\hat{g}_{t+1} = g_t$.

Moving average

Predicts the mean graph from the k preceding observations:


$$\hat{g}_{t+1} = \mathbb{E}^f[\mathbf{g}_t^k] = \arg \min_{g'} \sum_{g_i \in \mathbf{g}_t^k} d(g', g_i)^2.$$

Vector AR

Graphs are represented as vectors $u_t = [\text{vec}(\mathcal{V}_t)^\top, \text{vec}(\mathcal{E}_t)^\top]^\top \in \mathbb{R}^{N \cdot F + N^2}$, and then a linear autoregressive model is applied:

$$\hat{u}_{t+1} = B_0 + \sum_{i=1}^k B_i \cdot u_{t-i+1}.$$

Results






Figure 1: Graph edit distance between predicted and true graphs.




Figure 2: Loss v. complexity of the problem.



Spektral



A library for **relational representation learning**:

- Keras API, TensorFlow backend;
- Includes state-of-the-art GNN layers [2]–[6];
- [dani el egrattarola.github.io/spektral](https://github.com/dani-egrattarola/spektral).

Key ideas:

1. Formalized AR models for graphs;
2. GNNs as AR functions from \mathcal{G}^p to \mathcal{G} ;

Future works:

- Skip graph embedding, work directly in \mathcal{G} ;
- Improve definition of H .

Code: <https://github.com/dan-zam/NGAR>

D. Grattarola - grattd@usi.ch - Twitter: [@ricreasphait](https://twitter.com/ricreasphait)

D. Zambon - zambod@usi.ch





Università
della
Svizzera
italiana

- [1] E. Ott, *Chaos in Dynamical Systems*. Cambridge University Press, 2002.
- [2] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, “Graph attention networks,” *arXiv preprint arXiv:1710.10903*, 2017.
- [3] F. M. Bianchi, D. Grattarola, C. Alippi, and L. Livi, “Graph neural networks with convolutional arma filters,” *arXiv preprint arXiv:1901.01343*, 2019.
- [4] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *International Conference on Learning Representations (ICLR)*, 2016.
- [5] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3844–3852.

- [6] M. Simonovsky and N. Komodakis, “Dynamic edge-conditioned filters in convolutional neural networks on graphs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [7] D. I. Shuman, M. J. Faraji, and P. Vandergheynst, “A multiscale pyramid transform for graph signals,” *IEEE Transactions on Signal Processing*, vol. 64, no. 8, pp. 2119–2134, 2016.
- [8] R. Ying, J. You, C. Morris, X. Ren, W. L. Hamilton, and J. Leskovec, “Hierarchical graph representation learning with differentiable pooling,” *arXiv preprint arXiv:1806.08804*, 2018.
- [9] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, “Gated graph sequence neural networks,” *arXiv preprint arXiv:1511.05493*, 2015.


Notes on complexity



- **Graph network** blocks: $\mathcal{G} \rightarrow \mathcal{G}$.
- **Graph pooling** blocks: $\mathcal{G}_N \rightarrow \mathcal{G}_{M < N}$.
- **Differentiable**: can be integrated in any deep learning model.
- **State of the art** in graph representation learning.


Standard convolution

Convolutional neural networks exploit the **spatial locality** of pixels.




Graph convolution

Generalise convolutional layers to **arbitrary neighbours**.




Graph convolution




Node representation

Graph convolution

Generalise convolutional layers to **arbitrary neighbours**.




Graph convolution




Node representation

Graph convolution

Generalise convolutional layers to **arbitrary neighbours**.




Graph convolution




Node representation

Graph convolution

Generalise convolutional layers to **arbitrary neighbours**.

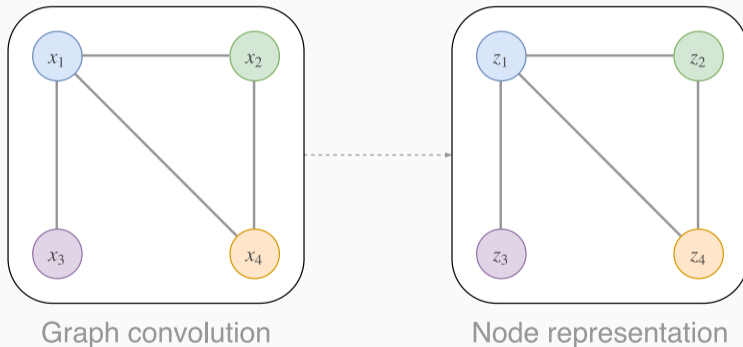


Graph convolution



Node representation

Generalise convolutional layers to **arbitrary neighbours**.



Graph convolution

The output for each node is given by [4]:

$$Z_i = \sum_{j \in \mathcal{N}_i} \tilde{A}_{ij} X_j W + b$$

Can also be written in matrix form:

$$Z = \tilde{A} \cdot X \cdot W + b$$

Note

\tilde{A} can be computed from the adjacency matrix in several ways (e.g., normalised Laplacian [4], row average [6], attention [2], ...)

[2] P. Veličković *et al.*, "Graph Attention Networks," 2017.

[4] T. N. Kipf *et al.*, "Semi-supervised classification with graph convolutional networks," 2016.

[6] M. Simonovsky *et al.*, "Dynamic edge-conditioned filters in convolutional neural networks on graphs," 2017.

We can take edge features into account by replacing W with a neural network that **transforms edge features into convolutional kernels** [6]:


$$Z_i = \sum_{j \in \mathcal{N}(i)} \tilde{A}_{ij} \cdot X_j \cdot W + b$$

becomes

$$Z_i = \sum_{j \in \mathcal{N}(i)} \tilde{A}_{ij} \cdot X_j \cdot f(E_{ji}) + b$$

[6] M. Simonovsky et al., "Dynamic edge-conditioned filters in convolutional neural networks on graphs," 2017.

Aggregate node features to obtain a lower-resolution graph [5], [7], [8].




[5] M. Defferrard *et al.*, "Convolutional neural networks on graphs with fast localized spectral filtering," 2016.

[7] D. I. Shuman *et al.*, "A multiscale pyramid transform for graph signals," 2016.

[8] R. Ying *et al.*, "Hierarchical Graph Representation Learning with Differentiable Pooling," 2018.

Graph pooling

Global pooling is the final step to reduce a graph to a vector, e.g., $Z = \sum_i \alpha_i X_i$.



α_i can be learned with neural network $f(X) : \mathbb{R}^{N \times F} \rightarrow [0, 1]^N$ to compute the relative importance of each node [9].