

# Designing and explaining graph neural networks

**Presenter:** Simone Scardapane



SAPIENZA  
UNIVERSITÀ DI ROMA



intelligent signal processing  
and multimedia lab

**Baker Hughes seminar, 2023**

# Not just me!



Aurelio Uncini  
**Full professor**



Massimo Panella  
**Full professor**



Indro Spinelli  
**Researcher  
(comp. science)**



Paolo Di Lorenzo  
**Associate professor**



Filippo Bianchi  
**Associate professor (UiT)**




Simone Scardapane  
**Tenure-track Assistant Professor**



<https://www.scardapane.it/>



[https://twitter.com/s\\_scardapane](https://twitter.com/s_scardapane)



Alessio Verdone  
**PhD**



Lev Telyatnikov  
**PhD**



Alessio Devoto  
**PhD**



Alessandro Baiocchi  
**PhD**



Gaetano Saurio  
**PhD**



Michele Guerra  
**PhD (UiT)**




Claudio Battiloro  
**PhD**

# Introduction

On the importance of **graphs** in  
**deep learning**


# Data ingestion in deep learning

Audio



WaveNet, Wav2Vec, ...

Images



CNNs, Vision Transformers, ...


Texts

Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt. Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt.

Word embeddings, Transformers, ...

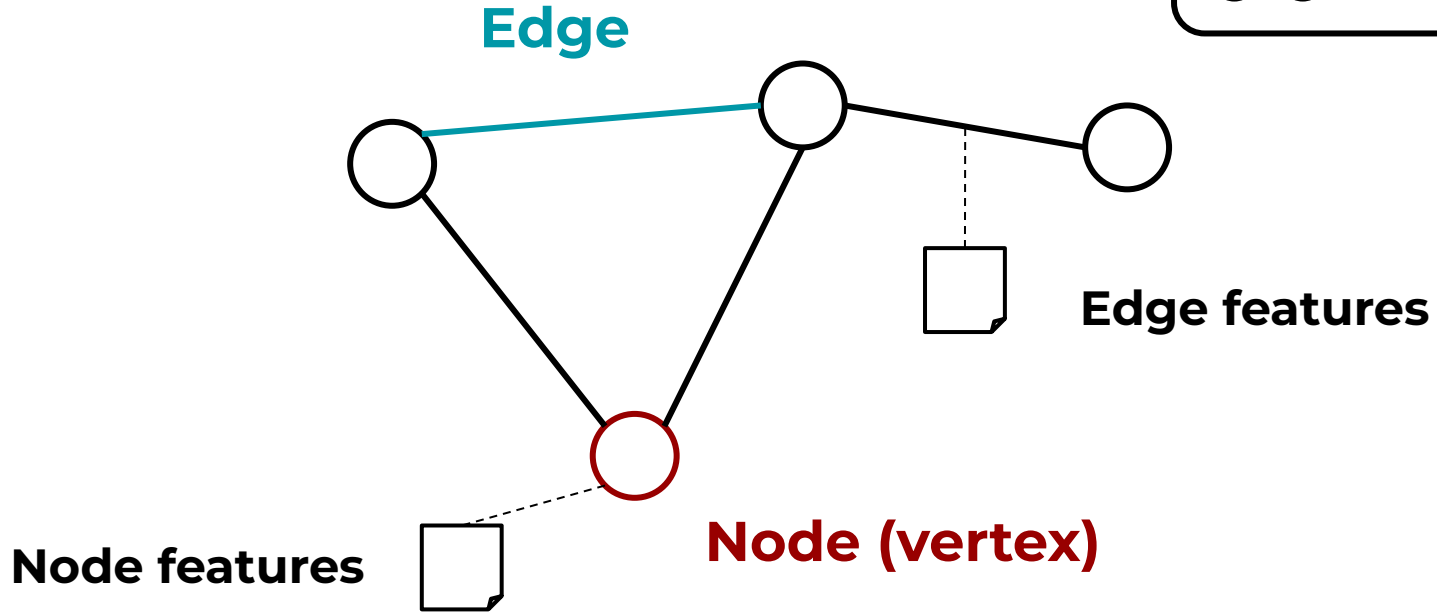

???

Graphs



# Defining a graph

Other graphs!

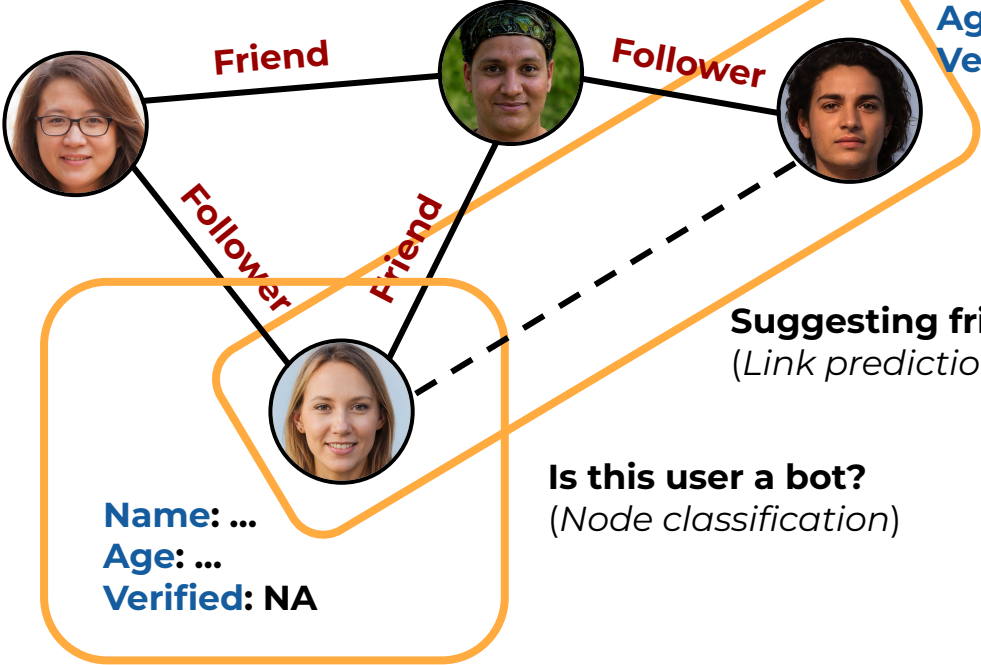


# Applications (1/3)

Name: ...  
Age: ...  
Verified: yes

Name: ...  
Age: ...  
Verified: yes

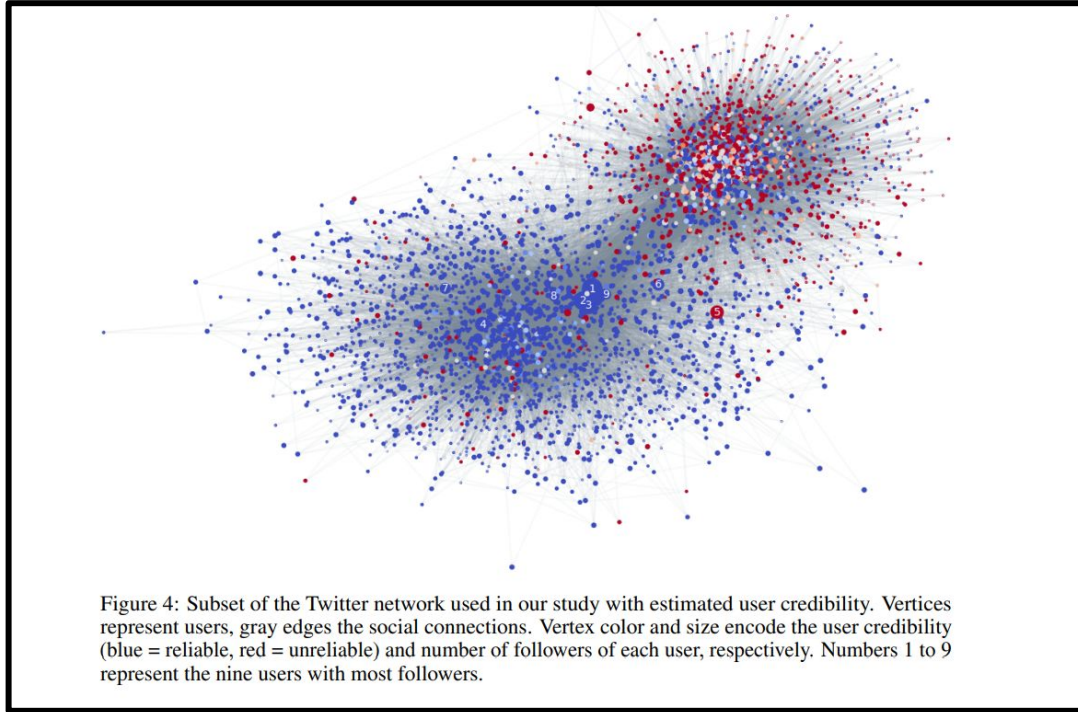
Name: ...  
Age: ...  
Verified: NA



Name: ...  
Age: ...  
Verified: NA

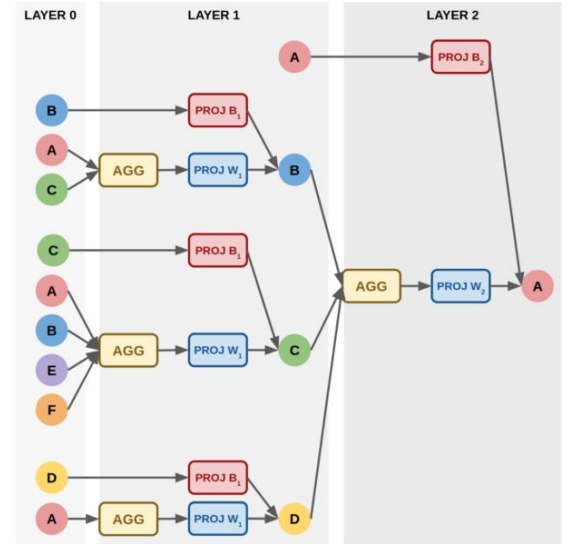


Is this user a bot?  
(Node classification)

# Fake news detection on Twitter



Monti, F., Frasca, F., Eynard, D., Mannion, D. and Bronstein, M.M., 2019. [Fake News Detection on Social Media using Geometric Deep Learning](#). *arXiv preprint arXiv:1902.06673*.


# Recommending systems in Uber



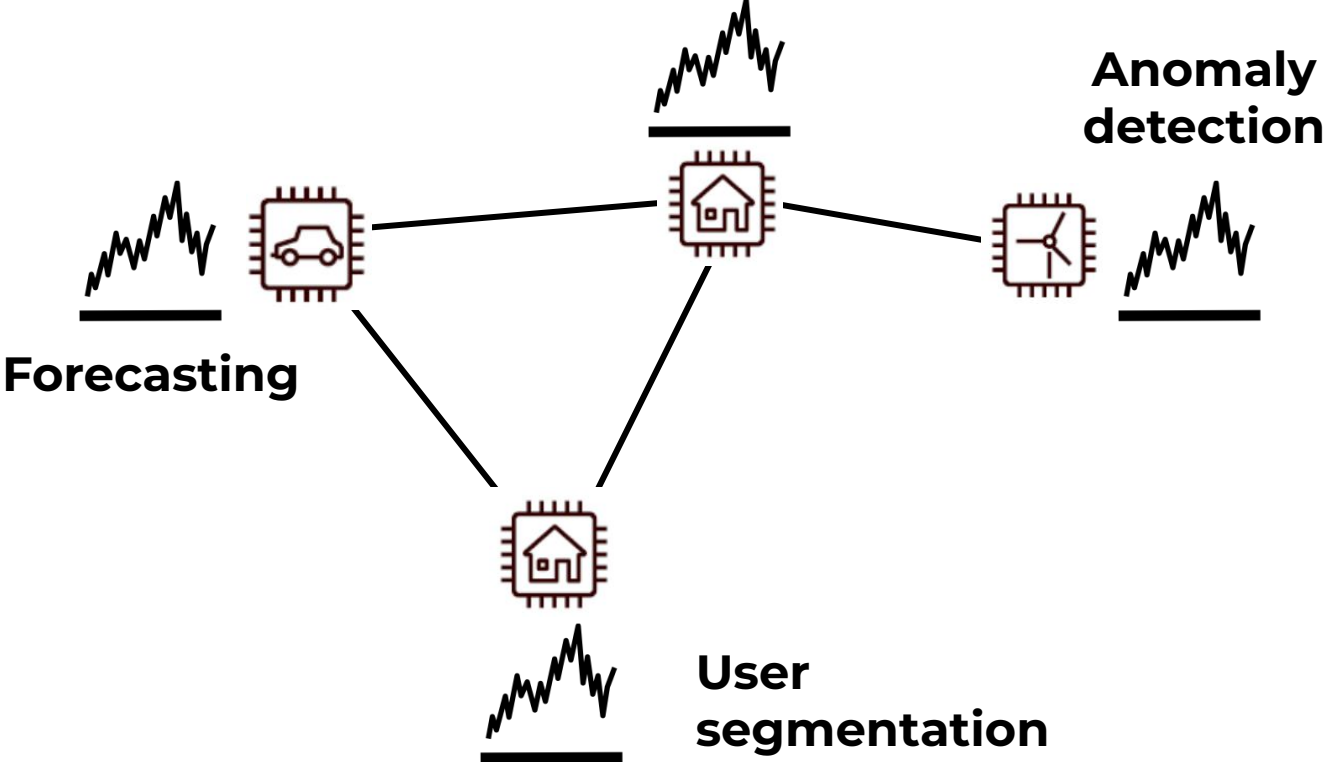
[Food Discovery with Uber Eats: Using Graph Learning to Power Recommendations](#)




# Applications (2/3)



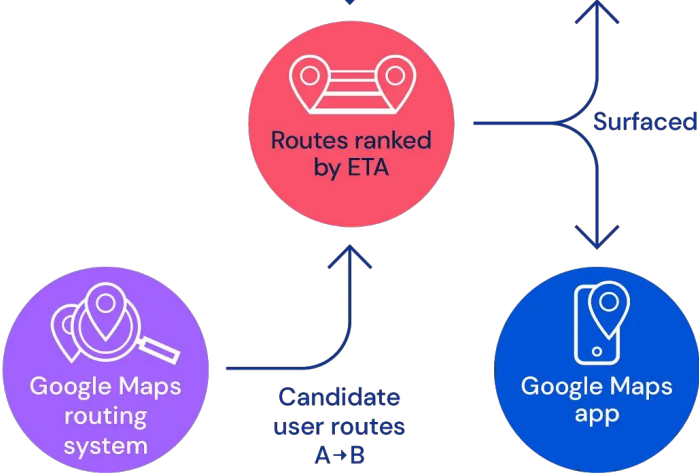
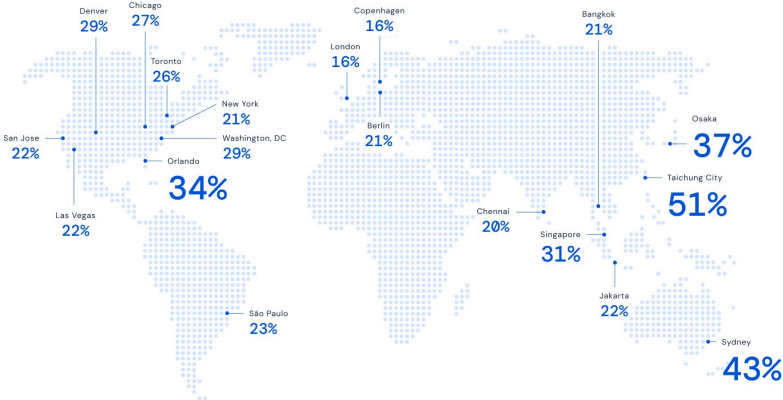
# Applications (3/3)



# Traffic prediction on Google Maps




Google Maps ETA Improvements Around the World



Traffic prediction with advanced Graph Neural Networks

# Distributed energy grids



What happens when our graph is also **physically distributed**?

# Energy forecasting on smart-grids


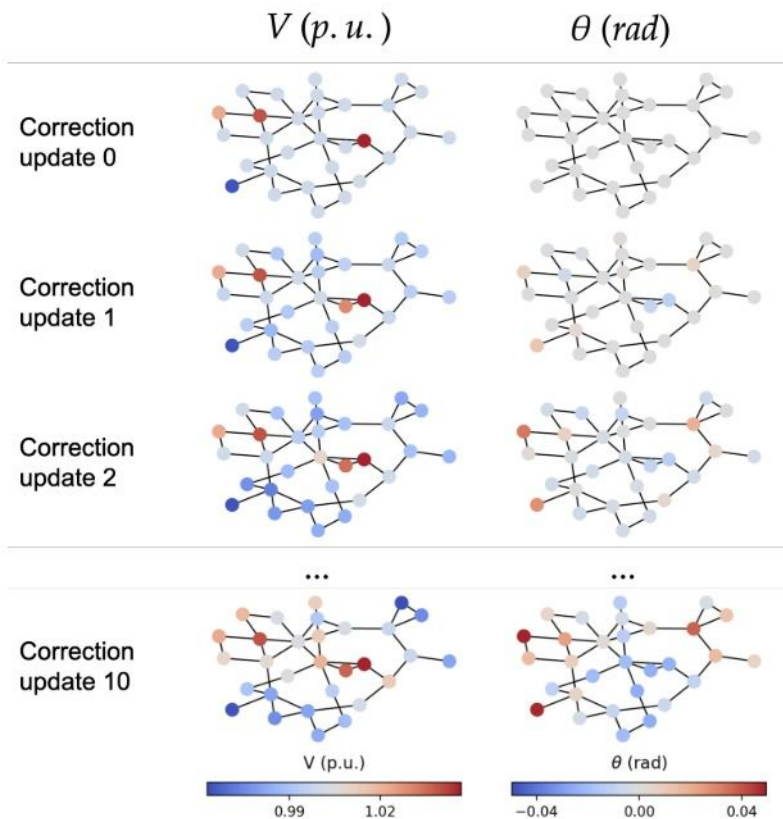
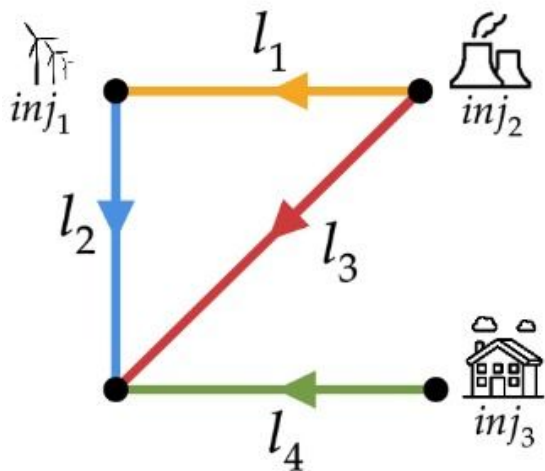


Fig. 2. Map showing the distribution of 31 simulated photovoltaic plants.

TABLE I  
1-DAY AHEAD FORECASTING ERROR OF MODELS TESTED ON UNIVARIATE REAL AND SYNTHETIC DATASETS

Model	PV 4		PV 31		PV 10	
	MSE	MAE	MSE	MAE	MSE	MAE
LSTM-FC	$0.0085 \pm 0.0002$	$0.0491 \pm 0.0011$	$0.0243 \pm 0.0002$	$0.0861 \pm 0.0008$	$0.0116 \pm 0.0007$	$0.0598 \pm 0.0046$
CNN-FC	$0.0144 \pm 0.0003$	$0.0673 \pm 0.0009$	$0.0237 \pm 0.0003$	$0.0855 \pm 0.0010$	$0.0109 \pm 0.0056$	$0.0578 \pm 0.0019$
GNN	$0.0063 \pm 0.0002$	$0.0412 \pm 0.0003$	$0.0108 \pm 0.0004$	$0.0559 \pm 0.0013$	$0.0043 \pm 0.0008$	$0.0355 \pm 0.0003$

# Power systems



Donon, B., Donnot, B., Guyon, I. and Marot, A., 2019. [Graph neural solver for power systems](#). In 2019 International Joint Conference on Neural Networks (IJCNN) (pp. 1-8). IEEE.

# Fully distributed GCNs





Fig. 1. Illustration of the proposed approach. In step 1, nodes communicate to perform inference. In step 2, a symmetric communication phase is executed to compute local gradients. In step 3, agents exchange local variables to asymptotically reach agreement. For steps 1-2, a representative active node is shown in red. Directed arrows show the flow of messages.

# Graph networks in classic deep learning




- Translation?
- Sentiment?
- Next word?
- Part-of-speech tags?

[\[1911.12247\] Contrastive Learning of Structured World Models](#)

**Transformers** are basically GNNs on fully-connected graphs!

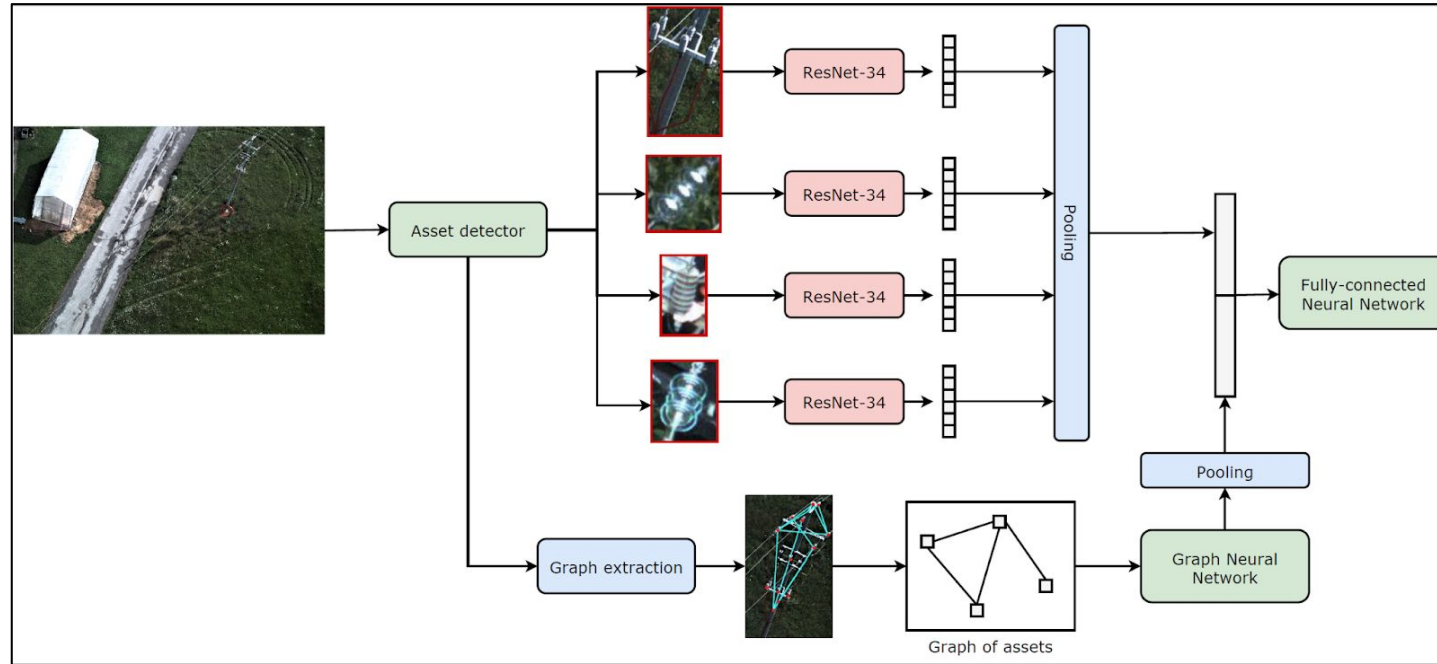
<https://thegradient.pub/transformers-are-graph-neural-networks/>

GNNs can be used to include **relational reasoning** in classical models!





# Integrating GNNs in other networks




Devoto, A., et al., 2022. **Re-identification of objects from aerial photos with hybrid siamese neural networks**. *IEEE Transactions on Industrial Informatics*, in press.

# Graph and learning

**Graphs** and **matrices**

# Graphs are (represented by) matrices



$$\begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{pmatrix}$$

$$\begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{pmatrix}$$

$$\begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix}$$

## Adjacency matrix

$$\mathbf{A} \in \mathbb{R}^{n \times n}$$

$n$  vertices in the graph

## Node features

$$\mathbf{X} \in \mathbb{R}^{n \times d}$$

each node has  $d$  features

## Edge features

$$\mathbf{E} \in \mathbb{R}^{e \times f}$$

each edge has  $f$  features

# Other interesting matrices

The degree matrix is a diagonal matrix collecting the degrees:

$$D_{ii} = \sum_j A_{ij}$$

We can use the degree matrix to normalize the adjacency matrix:

$$\mathbf{A}_{\text{norm}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$$


We can also add self-loops to the adjacency matrix:  $\mathbf{A} = \mathbf{A} + \mathbf{I}$

In general, we can replace the adjacency with any matrix with a **corresponding sparsity** structure.

# Diffusion over a graph

We can view the adjacency matrix as an operator that *diffuses* information across the graph:

$$[\mathbf{AX}]_i = \underbrace{\sum_j A_{ij} \mathbf{X}_j}_{\text{sparse sum!}}$$



For a classical adjacency matrix (only 0/1), the above simplifies to a sum, but in general it will be a **weighted sum**.

# Laplacian of a graph

The **Laplacian** is another fundamental graph matrix:

$$\mathbf{L} = \mathbf{D} - \mathbf{A}$$

(Note: we can also build normalized variants by replacing  $\mathbf{A}$  with any variant seen above.)

Any Laplacian acts as a diffusion operator as above, but its eigen-decomposition is fundamental in a number of disciplines:

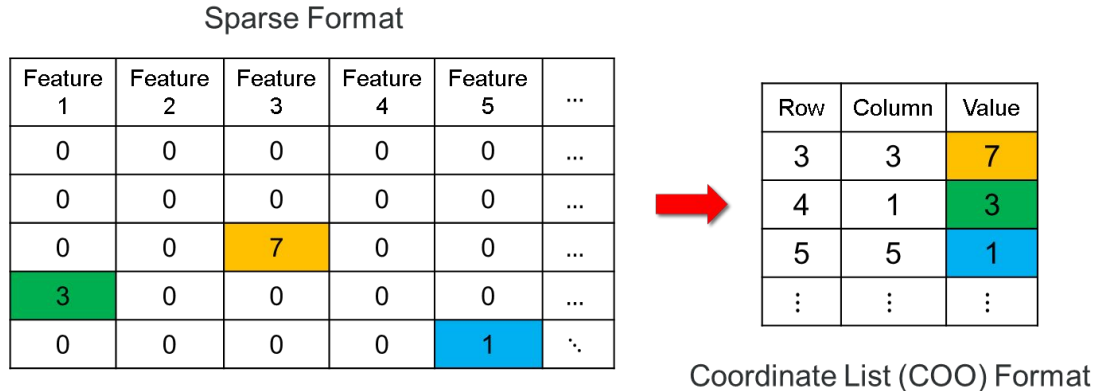
1. **Spectral graph theory**, and associated ML algorithms (e.g., spectral clustering).
2. **Graph signal processing**, where it allows to define equivalents of a Fourier transformation.

# Storing the graph matrices

[Tip: Working with Sparse Data in SAS - SAS Support Communities](#)

Matrix notation is convenient for describing the operations, but **sparse** matrices have dedicated storage formats in most software.

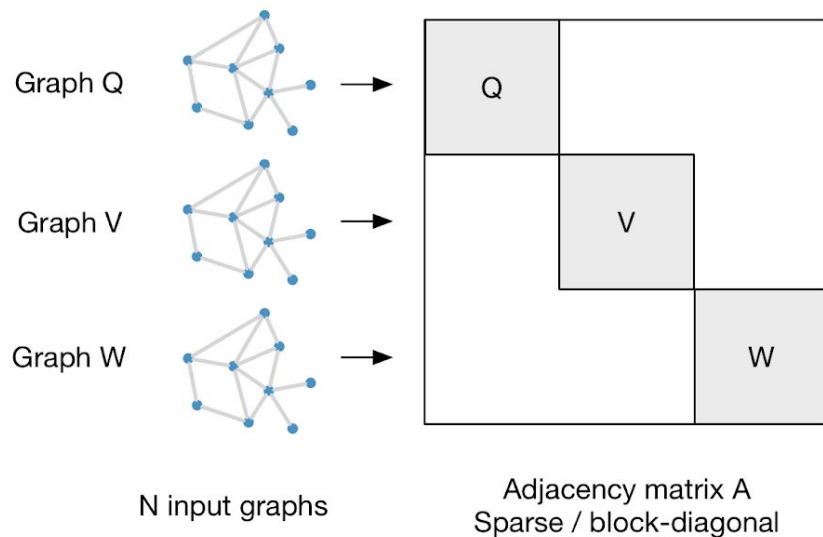
Coordinate list (COO) and its variants (CSR, CSC) is typical:



This also allows highly optimized variants of matrix multiplication.

# Handling multiple graphs

We can handle multiple graphs (e.g., mini-batches) by considering a single graph with several disconnected components:





# Implementing graph NNs

Software and code

# Scaling up to huge graphs

## Open Graph Benchmark

Benchmark datasets, data loaders and evaluators for graph machine learning

GET STARTED

VIEW UPDATES

The Open Graph Benchmark (OGB) is a collection of realistic, large-scale, and diverse benchmark datasets for machine learning on graphs. OGB datasets are automatically downloaded, processed, and split using the [OGB Data Loader](#). The model performance can be evaluated using the [OGB Evaluator](#) in a unified manner.

OGB is a community-driven initiative in active development. We expect the benchmark datasets to




**OPEN GRAPH BENCHMARK**

# Software?

## DeepGraphLibrary

Amazon based, production-ready



PyTorch  
geometric

PyTorch, more research oriented



Spektral


Keras-like, TensorFlow 2.0

tensorflow/gnn

TensorFlow GNN is a library to build Graph Neural Networks on the TensorFlow platform.



Alpha release, poorly documented



JAX-based  
Strong DeepMind adoption

# Notebook time!

[Colab Notebooks and Video Tutorials — pytorch\\_geometric documentation](#)

<https://colab.research.google.com/drive/1nV44NrNqcXC2thU6-zzxnJPnlalo870m?usp=sharing>

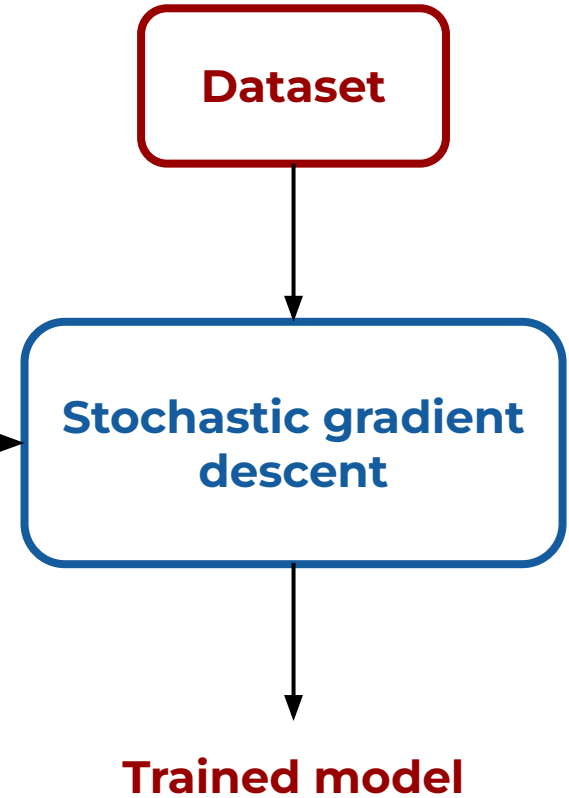
**Building graph layers**

**Graph convolutional layers**


# Neural networks in a single slide

$$f \left( \begin{array}{c} \bigcirc \\ \bigcirc \\ \bigcirc \\ \bigcirc \end{array} \right) = \begin{array}{c} \bigcirc \\ \bigcirc \\ \bigcirc \\ \bigcirc \end{array}$$

$\mathbf{W}\mathbf{x} + \mathbf{b}$

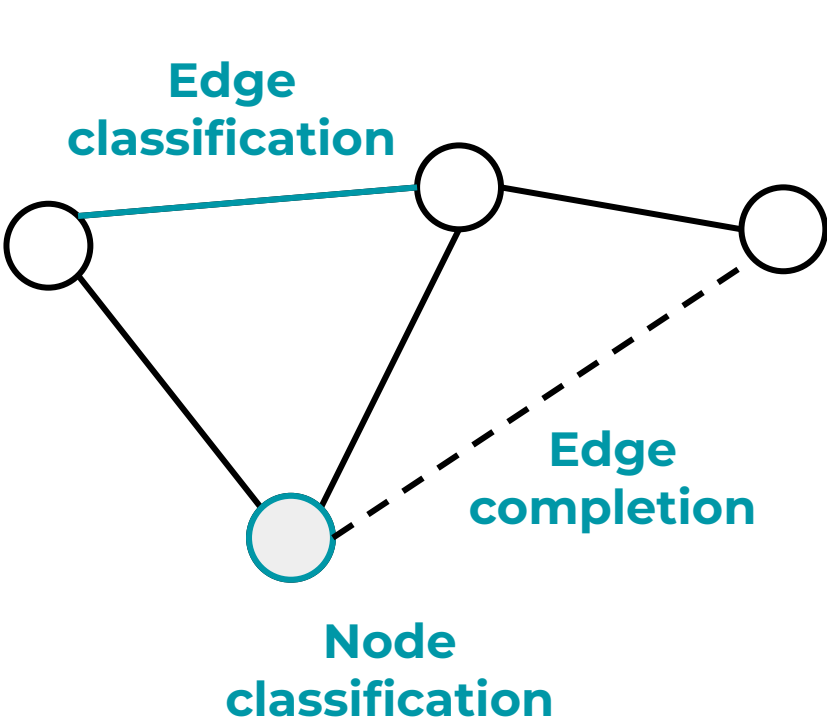


# How can we learn on a graph?



We want to do deep learning, hence  $f$  should be **differentiable**, **composable**, **scalable**.

# Summary: what can we learn on a graph?



Graph classification

Graph generation

Graph summarization





# GNNs before deep learning!

## The Graph Neural Network Model

Franco Scarselli, Marco Gori, *Fellow, IEEE*, Ah Chung Tsoi, Markus Hagenbuchner, *Member, IEEE*, and Gabriele Monfardini

**Abstract**—Many underlying relationships among data in several areas of science and engineering, e.g., computer vision, molecular chemistry, molecular biology, pattern recognition, and data mining, can be represented in terms of graphs. In this paper, we propose a new neural network model, called graph neural network (GNN) model, that extends existing neural network methods for processing the data represented in graph domains. This GNN model, which can directly process most of the practically useful types of graphs, e.g., acyclic, cyclic, directed, and undirected, implements a function  $\tau(\mathbf{G}, n) \in \mathbb{R}^m$  that maps a graph  $\mathbf{G}$  and one of its nodes  $n$  into an  $m$ -dimensional Euclidean space. A supervised learning algorithm is derived to estimate the parameters of the proposed GNN model. The computational cost of the proposed algorithm is also considered. Some experimental results are shown to validate the proposed learning algorithm, and to demonstrate its generalization capabilities.

ples a function  $\tau$  that maps a graph  $G$  and one of its nodes  $n$  to a vector of reals<sup>1</sup>:  $\tau(\mathbf{G}, n) \in \mathbb{R}^m$ . Applications to a graphical domain can generally be divided into two broad classes, called *graph-focused* and *node-focused* applications, respectively, in this paper. In *graph-focused* applications, the function  $\tau$  is independent of the node  $n$  and implements a classifier or a regressor on a graph structured data set. For example, a chemical compound can be modeled by a graph  $\mathbf{G}$ , the nodes of which stand for atoms (or chemical groups) and the edges of which represent chemical bonds [see Fig. 1(a)] linking together some of the atoms. The mapping  $\tau(\mathbf{G})$  may be used to estimate the probability that the chemical compound causes a certain disease [13]. In Fig. 1(b), an image is represented by a region adjacency graph where nodes denote homogeneous regions of intensity of

# Geometric deep learning

GEOMETRIC DEEP LEARNING

Grids, Groups, Graphs, Geodesics, and Gauges

Michael M. Bronstein, Joan Bruna, Taco Cohen, Petar Veličković

[Read the paper](#) [Read the blog post](#) [Watch the ICLR'21 keynote](#)

[Watch the Erlangen keynote](#) [Follow the Lectures \(AMMI 2021\)](#)

[Contact the authors](#)

Bronstein, M.M., Bruna, J., Cohen, T. and Veličković, P., 2021. **Geometric deep learning: Grids, groups, graphs, geodesics, and gauges**. *arXiv preprint arXiv:2104.13478*.



# A zoo of techniques...

TABLE III: Summary of RecGNNs and ConvGNNs. Missing values (“-”) in pooling and readout layers indicate that the method only experiments on node-level/edge-level tasks.

Approach	Category	Inputs	Pooling	Readout	Time Complexity
GNN* (2009) [15]	RecGNN	$A, X, X^e$	-	a dummy super node	$O(m)$
GraphESN (2010) [16]	RecGNN	$A, X$	-	mean	$O(m)$
GGNN (2015) [17]	RecGNN	$A, X$	-	attention sum	$O(m)$
SSE (2018) [18]	RecGNN	$A, X$	-	-	-
Spectral CNN (2014) [19]	Spectral-based ConvGNN	$A, X$	spectral clustering+max pooling	max	$O(n^3)$
Henaff et al. (2015) [20]	Spectral-based ConvGNN	$A, X$	spectral clustering+max pooling		$O(n^3)$
ChebNet (2016) [21]	Spectral-based ConvGNN	$A, X$	efficient pooling	sum	$O(m)$
GCN (2017) [22]	Spectral-based ConvGNN	$A, X$	-	-	$O(m)$
CayleyNet (2017) [23]	Spectral-based ConvGNN	$A, X$	mean/graclus pooling	-	$O(m)$
AGCN (2018) [40]	Spectral-based ConvGNN	$A, X$	max pooling	sum	$O(n^2)$
DualGCN (2018) [41]	Spectral-based ConvGNN	$A, X$	-	-	$O(m)$
NN4G (2009) [24]	Spatial-based ConvGNN	$A, X$	-	sum/mean	$O(m)$
DCNN (2016) [25]	Spatial-based ConvGNN	$A, X$	-	mean	$O(n^2)$

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C. and Yu, P.S., 2019. [A comprehensive survey on graph neural networks](#). arXiv preprint arXiv:1901.00596.

# Deep learning is about leveraging structure



[Generate cross-stitch patterns from any image.](#)

# Image convolutions

$$f \left( \begin{array}{cccc} \square & \square & \square & \square \\ \square & \square & \square & \square \\ \square & \square & \square & \square \\ \square & \square & \square & \square \end{array} \right) = \mathbf{a}^T \begin{array}{c} \square \\ \square \\ \square \\ \square \end{array} + \sum_{j \in \begin{array}{ccc} \square & \square & \square \\ \square & \square & \square \\ \square & \square & \square \end{array}} \mathbf{c}_j^T \begin{array}{c} \square \\ \square \\ \square \\ \square \end{array}$$

Local pixel operation

Aggregation

Neighbour pixel operation

# Images vs. graphs

How much of the structure of an image do we find in graphs?

✓ **Locality** (neighbourhood).

✗ Fixed number of neighbours.

✗ Neighbours have a definite ordering.

# Permutation equivariance

Consider a generic permutation matrix of dimension  $n$ :

$$\underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{P}} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix} = \begin{bmatrix} \mathbf{x}_2 \\ \mathbf{x}_1 \\ \mathbf{x}_3 \end{bmatrix}$$

Permuting the nodes of a graph results in an equivalent graph (**isomorphism**):

$$(\mathbf{P}\mathbf{X}, \mathbf{P}\mathbf{A}\mathbf{P}^\top) \sim (\mathbf{X}, \mathbf{A})$$

## Permutation equivariance (2)

Any graph layer must possess a property called **permutation equivariance**:

$$\underbrace{f(\mathbf{P}\mathbf{X}, \mathbf{P}\mathbf{A}\mathbf{P}^\top)}_{\text{permuted input}} = \overbrace{\mathbf{P} f(\mathbf{X}, \mathbf{A})}^{\text{permuted output}}$$

Running the layer with a different node ordering should modify the output ordering **only**.



# Putting everything together (graph convolutions)

$$f \left( \begin{array}{c} \text{graph} \end{array} \right) = A_{ii} \text{ (red circle)} + \sum_{j \in \text{neighbors}} A_{ij} \text{ (blue circle)}$$

Local **vertex** operation

We use the **adjacency matrix** (or similar)

**Note:** the mixing coefficients are no more learnable!  
Here, only local operations can be trainable.

We use the **adjacency matrix**  
(or similar)

# Graph convolutions

Writing it out explicitly for a single node:

$$[\mathbf{H}]_i = \sum_j A_{ij} f(\mathbf{X}_j) + A_{ii} g(\mathbf{X}_i)$$

In the simplest case (this is a classical **graph convolutional layer**):

$$f(\mathbf{x}) = g(\mathbf{x}) = \mathbf{W}\mathbf{x}$$

---

Kipf, T.N. and Welling, M., 2016. **Semi-supervised classification with graph convolutional networks**. arXiv preprint arXiv:1609.02907.

# Properties of a graph convolutional layer

The layer can be understood as:



1. Applying a local operation at every node;
2. **Aggregating** the updated embeddings depending on the graph topology.





(If you are curious: it can be understood as a linear filter in the spectral domain given by the Graph Fourier transform.)

We can write it out compactly as:

$$\mathbf{H} = \mathbf{A}\mathbf{X}\mathbf{W}$$

# Graph convolution visualized




-   Pixel/Node updated
-  Convolution
-  Graph Convolution

**Building graph *networks***

**Stacking** graph layers

# Stacking graph convolutional layers



Performing multiple updates increases the “**receptive field**” of each node.

<https://spindro.github.io/post/gnn/>

# Building deep graph networks

Stacking graph layers works similarly to standard deep networks, since the topology of the graph is unchanged:

$$\mathbf{H} = \phi \left( \mathbf{A} \underbrace{\phi(\mathbf{A} \mathbf{X} \mathbf{W}_1)}_{\text{layer 1}} \mathbf{W}_2 \right)$$

Deep graph networks may suffer from **oversmoothing** (all representations collapse to the same). Some specialized solutions exist (e.g., PairNorm regularization).


# Pooling

Table 1: Pooling methods in the SRC framework. GNN indicates a stack of one or more message-passing layers, MLP is a multi-layer perceptron,  $\mathbf{L}$  is the normalized graph Laplacian,  $\beta$  is a regularization vector (see [42]),  $\mathbf{D}$  is the degree matrix,  $\mathbf{u}_{max}$  is the eigenvector of the Laplacian associated with the largest eigenvalue,  $\mathbf{i}$  is a vector of indices,  $\mathbf{A}_{\mathbf{i},\mathbf{i}}$  selects the rows and columns of  $\mathbf{A}$  according to  $\mathbf{i}$ .

Method	Select	Reduce	Connect
DiffPool [55]	$\mathbf{S} = \text{GNN}_1(\mathbf{A}, \mathbf{X})$ (w/ auxiliary loss)	$\mathbf{X}' = \mathbf{S}^\top \cdot \text{GNN}_2(\mathbf{A}, \mathbf{X})$	$\mathbf{A}' = \mathbf{S}^\top \mathbf{A} \mathbf{S}$
MinCut [6]	$\mathbf{S} = \text{MLP}(\mathbf{X})$ (w/ auxiliary loss)	$\mathbf{X}' = \mathbf{S}^\top \mathbf{X}$	$\mathbf{A}' = \mathbf{S}^\top \mathbf{A} \mathbf{S}$
NMF [3]	Factorize: $\mathbf{A} = \mathbf{W} \mathbf{H} \rightarrow \mathbf{S} = \mathbf{H}^\top$	$\mathbf{X}' = \mathbf{S}^\top \mathbf{X}$	$\mathbf{A}' = \mathbf{S}^\top \mathbf{A} \mathbf{S}$
LaPool [42]	$\begin{cases} \mathbf{V} = \ \mathbf{L}\mathbf{X}\ _d; \\ \mathbf{i} = \{i \mid \forall j \in \mathcal{N}(i) : \mathbf{V}_i > \mathbf{V}_j\} \\ \mathbf{S} = \text{SparseMax} \left( \beta \frac{\mathbf{X}\mathbf{X}_i^\top}{\ \mathbf{X}\  \ \mathbf{X}_i\ } \right) \end{cases}$	$\mathbf{X}' = \mathbf{S}^\top \mathbf{X}$	$\mathbf{A}' = \mathbf{S}^\top \mathbf{A} \mathbf{S}$
Graclus [16]	$\mathcal{S}_k = \left\{ \mathbf{x}_i, \mathbf{x}_j \mid \arg \max_j \left( \frac{\mathbf{A}_{ij}}{\mathbf{D}_{ii}} + \frac{\mathbf{A}_{jj}}{\mathbf{D}_{jj}} \right) \right\}$	$\mathbf{X}' = \mathbf{S}^\top \mathbf{X}$	METIS [26]
NDP [7]	$\mathbf{i} = \{i \mid \mathbf{u}_{max,i} > 0\}$	$\mathbf{X}' = \mathbf{X}_{\mathbf{i}}$	Kron r. [18]
Top- $K$ [24]	$\mathbf{y} = \frac{\mathbf{X}\mathbf{p}}{\ \mathbf{p}\ }$ ; $\mathbf{i} = \text{top}_K(\mathbf{y})$	$\mathbf{X}' = (\mathbf{X} \odot \sigma(\mathbf{y}))_{\mathbf{i}}$	$\mathbf{A}' = \mathbf{A}_{\mathbf{i},\mathbf{i}}$
SAGPool [30]	$\mathbf{y} = \text{GNN}(\mathbf{A}, \mathbf{X})$ ; $\mathbf{i} = \text{top}_K(\mathbf{y})$	$\mathbf{X}' = (\mathbf{X} \odot \sigma(\mathbf{y}))_{\mathbf{i}}$	$\mathbf{A}' = \mathbf{A}_{\mathbf{i},\mathbf{i}}$



# Tackling multiple tasks






1. Node classification:  $\text{sotfmax}(\mathbf{h}_i)$
2. Edge classification:  $\text{sotfmax}(\mathbf{h}_i^T \mathbf{h}_j)$
3. Graph classification:  $\text{sotfmax} \left( \frac{1}{N} \sum_i \mathbf{h}_i \right)$

# Mini-batching in graph networks

There are two types of mini-batching for graph networks:

1. For graph-level tasks, we can create a mini-batch of several graphs.
2. For node/edge-level tasks, we can create a mini-batch by **sampling** nodes and edges from a larger graph.



# Notebook time!

[Colab Notebooks and Video Tutorials — pytorch\\_geometric documentation](#)

<https://colab.research.google.com/drive/1nV44NrNqcXC2thU6-zzxnJPnlalo870m?usp=sharing>

# Beyond graph convolutions

## Message-passing graph layers

# Revisiting the GCN layer

We can write down the GCN layer in a slightly more general form:

$$\mathbf{H}_i = \text{Aggregate} (\{\phi(\mathbf{X}_j) | j \in \mathcal{N}_i\})$$

For each neighbour we build an update, which is the GCN case is weighted by the corresponding element of the adjacency matrix.

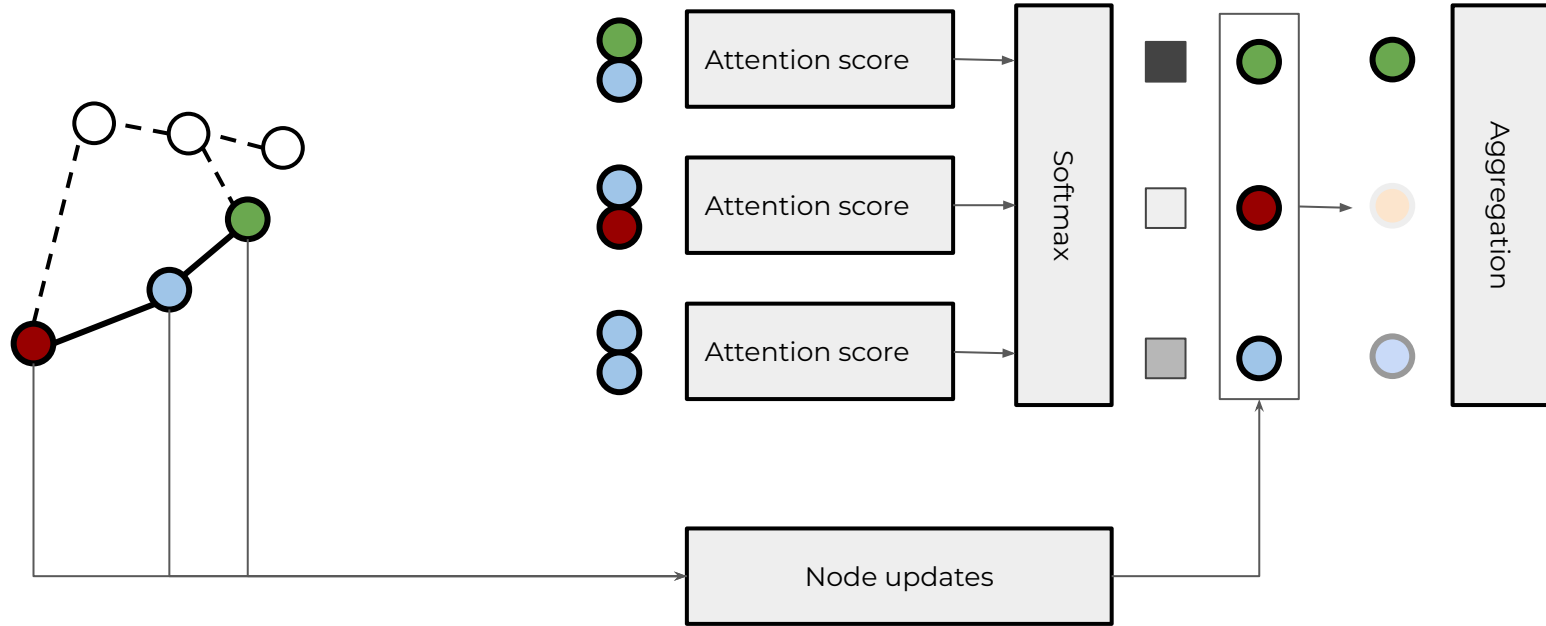
# Attention over nodes

The biggest limit of graph convolutional layers is that the weight node  $i$  gives to node  $j$  is fixed.

In principle, it is possible that one neighbour is more important than another, and we would like to learn this.

This can be achieved with an **attention mechanism**.

# Visualizing graph attention (for node )



## A more general setup (2)


More in detail, in a **graph attention layer**, we compute messages using an attention mechanism:

$$\psi(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{a}^\top \text{LeakyReLU}(\mathbf{W} [\mathbf{x}_i \parallel \mathbf{x}_j])$$

$$\mathbf{m}_i = \sum_j \text{softmax}_j(\mathbf{m}_{ij})(\mathbf{W}\mathbf{x}_j)$$



# Message-passing neural networks






Instead of using directly the adjacency matrix, nodes can exchange **messages** with several mechanisms (e.g., attention models).

---

Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O. and Dahl, G.E., 2017, July. **Neural message passing for quantum chemistry**. In *International Conference on Machine Learning* (pp. 1263-1272). PMLR.

# A more general setup



Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R. and Gulcehre, C., 2018. **Relational inductive biases, deep learning, and graph networks**. arXiv preprint arXiv:1806.01261.

# Notebook time!

[Colab Notebooks and Video Tutorials — pytorch\\_geometric documentation](#)


<https://colab.research.google.com/drive/1nV44NrNqcXC2thU6-zzxnJPnlalo870m?usp=sharing>

# Beyond message-passing


**Transformers** as universal neural models

# The Transformers revolution

Audio



Audio Transformers  
(2020-2021)



Images


Vision Transformers  
(2020-2021)

NLP Transformers  
(2017)

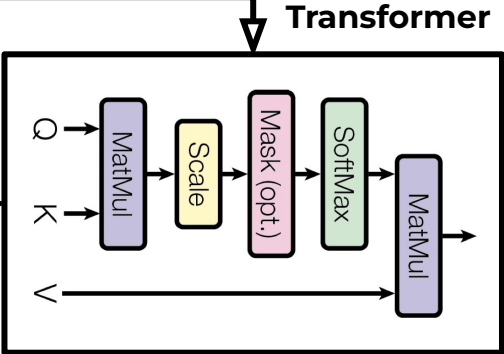
Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt. Dies ist ein Blindtext. An ihm lässt sich vieles über die Schrift ablesen, in der er gesetzt ist. Auf den ersten Blick wird der Grauwert der Schriftfläche sichtbar. Dann kann man prüfen, wie gut die Schrift zu lesen ist und wie sie auf den Leser wirkt.

Texts

Graph Transformers  
(2022)



Graphs



# Scaling laws for Vision Transformers

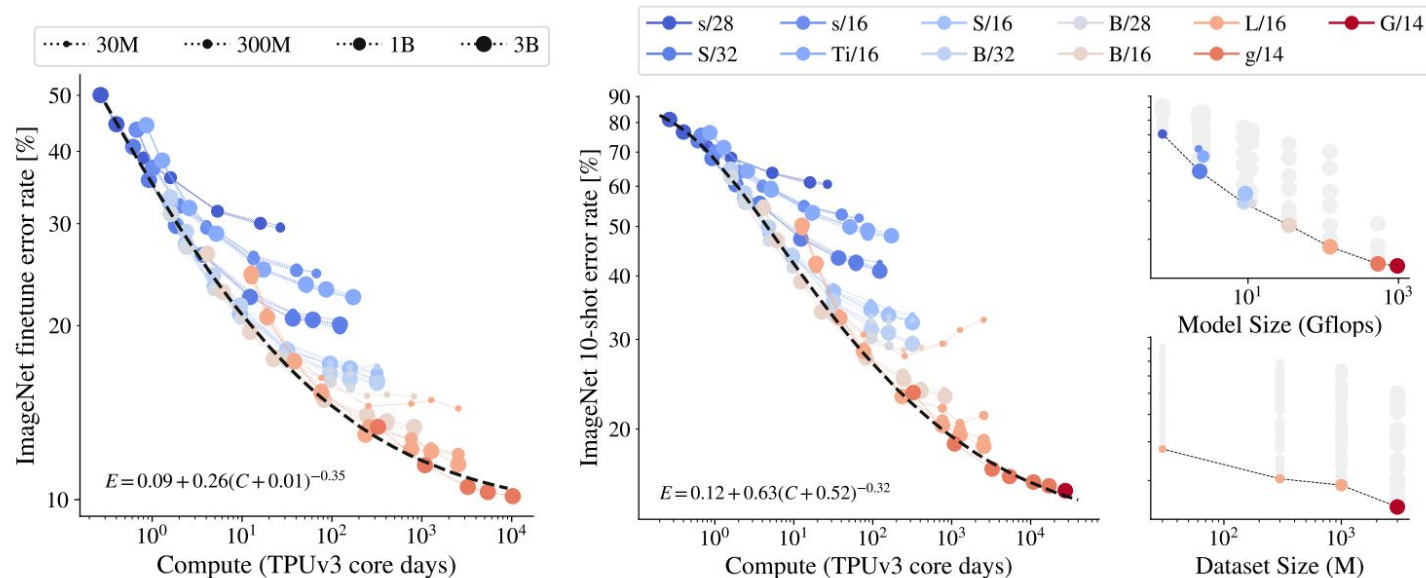


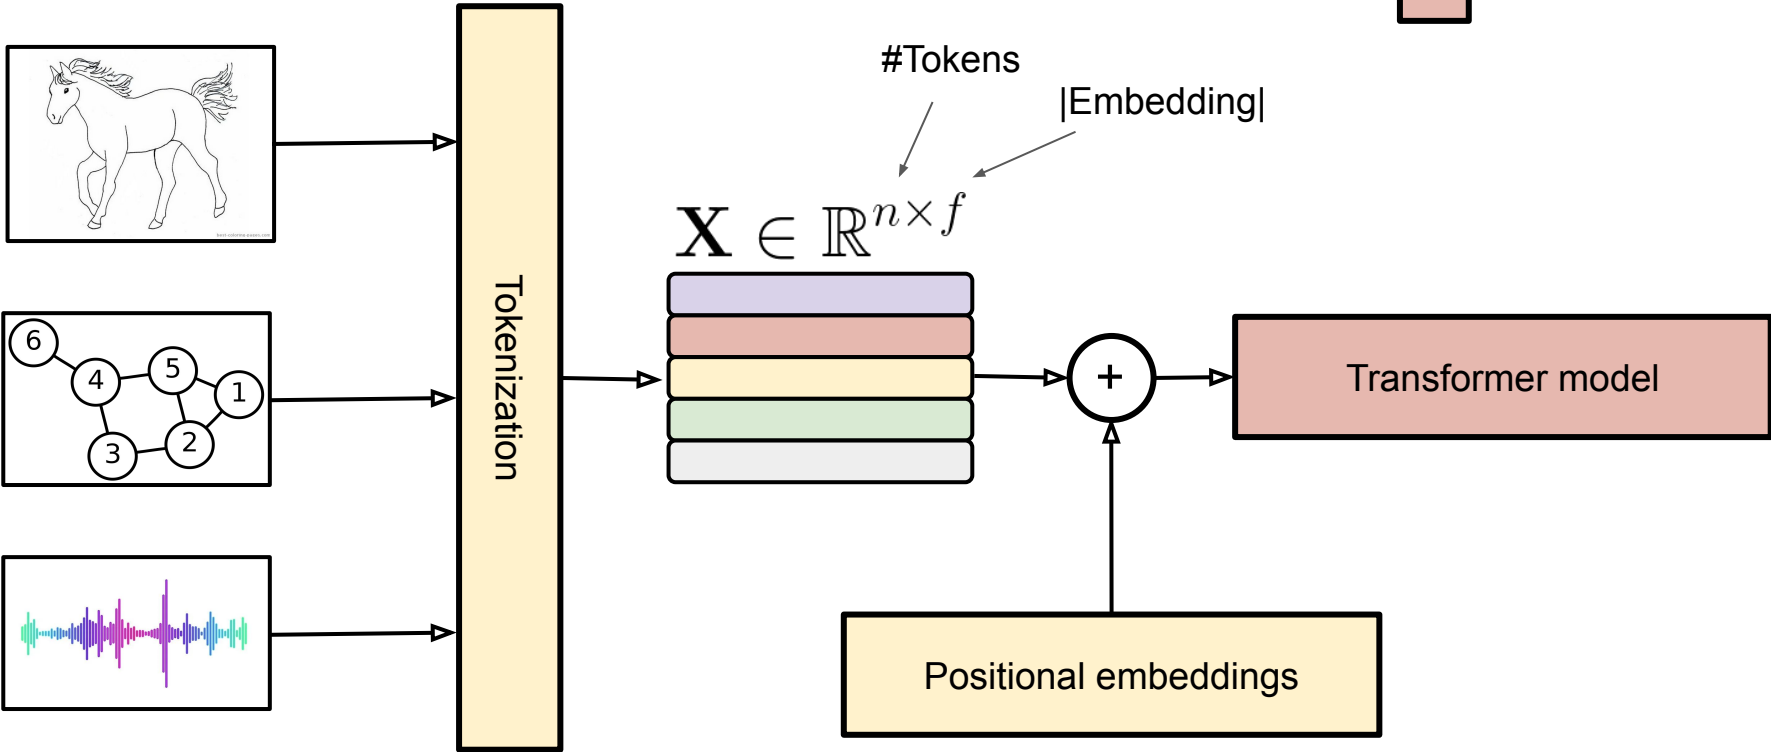



Figure 2. **Left/Center:** Representation quality, measured as ImageNet finetune and linear 10-shot error rate, as a function of total training compute. A saturating power-law approximates the Pareto frontier fairly accurately. Note that smaller models (blue shading), or models trained on fewer images (smaller markers), saturate and fall off the frontier when trained for longer. **Top right:** Representation quality when bottlenecked by model size. For each model size, a large dataset and amount of compute is used, so model capacity is the main bottleneck. Faintly-shaded markers depict sub-optimal runs of each model. **Bottom Right:** Representation quality by datasets size. For each dataset size, the model with an optimal size and amount of compute is highlighted, so dataset size is the main bottleneck.

# Transformers at a glance

-  Input-dependent
-  Input-independent



# Zooming in



A small neural network (e.g., 2 layers) applied to each token independently.  
*Fast, average number of parameters.*

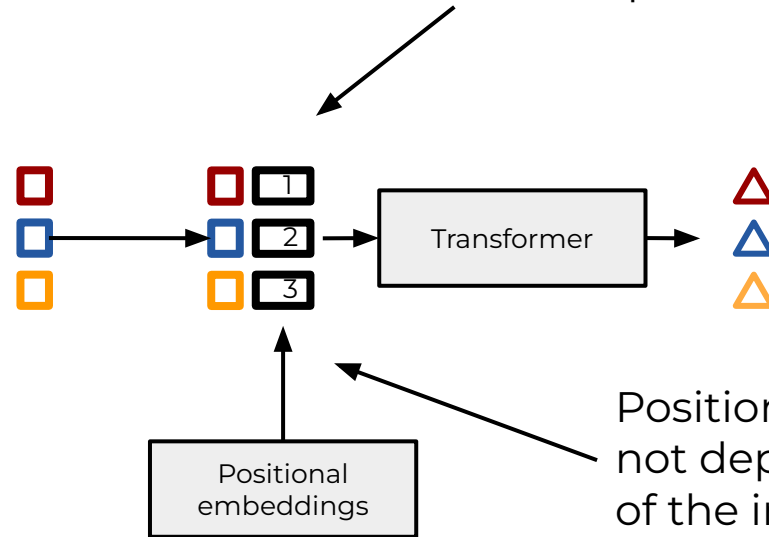
Main component, allowing to combine information *across* different tokens.  
*Quadratic in the number of tokens!*

Normalization (typically layer normalization): helps in stabilizing mean and variance of the embeddings.  
*Very fast, small number of parameters.*



# Positional embeddings


This vector should encode information about the position of the corresponding token.




Positional embeddings do not depend on the *content* of the image or input.

# Sinusoidal embeddings

[Transformer Architecture: The Positional Encoding - Amirhossein Kazemnejad's Blog](#)



# Example of (graph) transformers



# Taxonomy of graph transformers

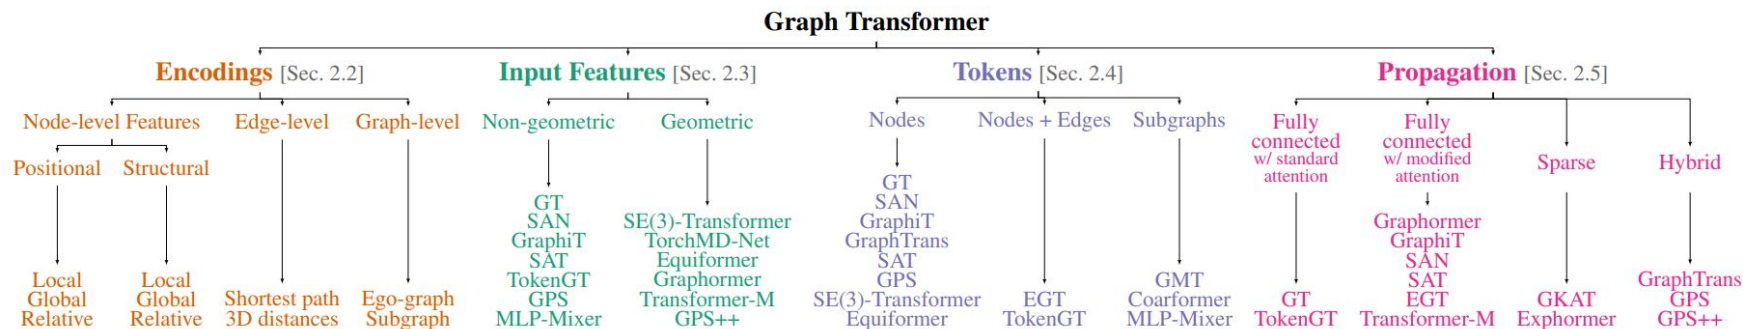


Figure 1: Categorization of graph transformers along four main categories with representative architectures.

[\[2302.04181\] Attending to Graph Transformers](#)

# Graph transformers

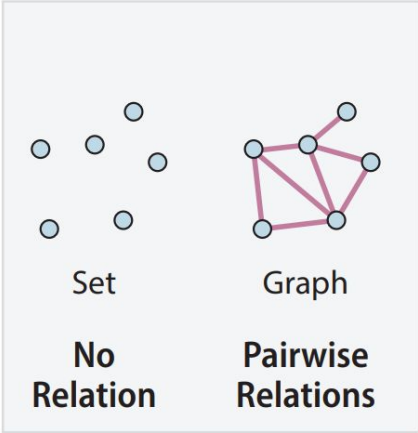
- ✓ Each MHA operation has a global receptive field over the graph (better for long-range interactions).
- ✗ Much harder to design structural embeddings to encode the graph connectivity (e.g., random walks, Laplacian embeddings).
- ✗ Attention is quadratic in the number of tokens.  
  
No clear scaling law for graphs / unclear expressiveness compared to GNNs.

**Beyond graphs**

**Hypergraphs** and **latent inference**

# Scaling to higher-order structures

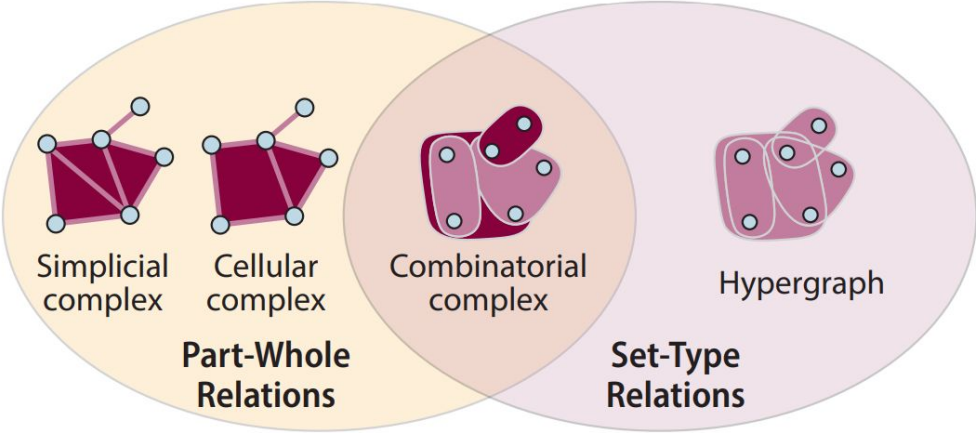
## Traditional Discrete Domains



○ : Nodes

∖ : Edges

## Domains of Topological Deep Learning

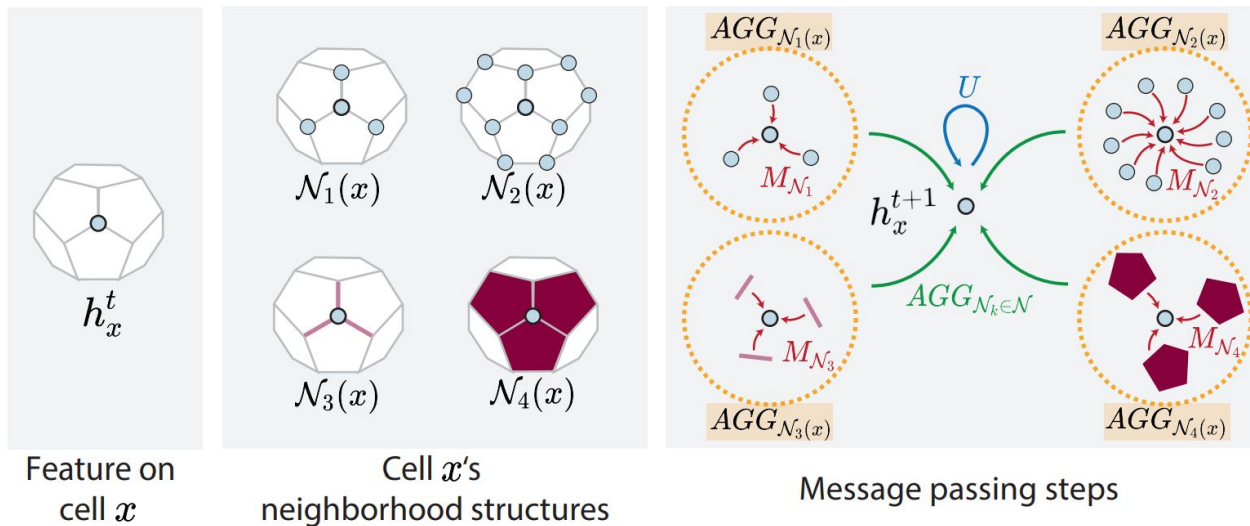


∖ is part of ▲

◐ not necessarily part of ●

Figure 2: Domains: Nodes in blue, (hyper)edges in pink, and faces in dark red. Inspired by Hajij et al. (2022a).


# Topological message-passing



**Figure 8: Message passing steps:** 1: Message (red), 2: Within-neighborhood aggregation (orange), 3: Between-neighborhood aggregation (green), 4: Update (blue). The scheme updates a feature  $\mathbf{h}_x^{t,(r)}$  on a  $r$ -cell  $x$  at layer  $t$  (left column) into a new feature  $\mathbf{h}_x^{t+1,(r)}$  on that same cell at the next layer  $t + 1$  (right column). Here, the scheme uses four neighborhood structures  $\mathcal{N}_k$  for  $k \in \{1, 2, 3, 4\}$  (middle column). Inspired by (Hajij et al., 2023).



# Topological lifting



**Figure 4: Lifting Topological Domains.** (a) A graph is “lifted” to a hypergraph by adding hyperedges that connect groups of nodes. (b) In the process of lifting a graph to a simplicial complex, a pairwise edge must be added in order to form triangular faces. (c) A graph can be converted to a cellular complex by adding faces of any shape. (d) Hyperedges can be added to a cellular complex to lift the structure to a combinatorial complex. Figure adapted from [Hajij et al. \(2023\)](#).

# Learning the latent connectivity

## Differentiable Graph Module (DGM) Graph Convolutional Networks

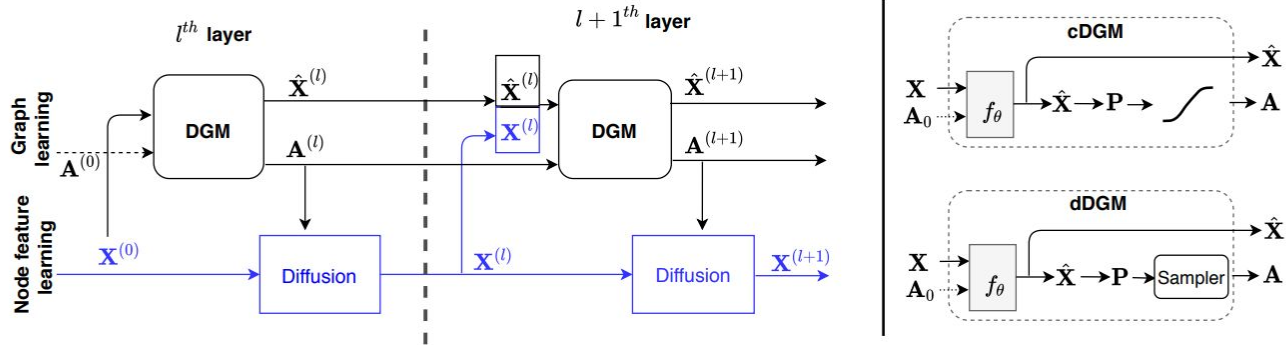


Figure 1. Left: Two-layered architecture including Differentiable Graph Module (DGM) that learns the graph, and Diffusion Module that uses the graph convolutional filters. Right: Details of DGM in its two variants, cDGM and dDGM.

# Learning the *topological* connectivity




Figure 1: The proposed two-step procedure for Latent Topology Inference (LTI) via regular cell complexes. The Differentiable Cell Complex Module (DCM) is a function that first learns a graph describing the pairwise interactions among data points via the  $\alpha$ -Differentiable Graph Module ( $\alpha$ -DGM), and then it leverages the graph as the 1-skeleton of a regular cell complex whose 2-cells (polygons), describing multi-way interactions among data points, are learned via the Polygon Inference Module (PIM). The inferred topology is then used in two message passing networks, at node (Graph Neural Network, GNN) and edge (Cell Complex Neural Network, CCNN) levels to solve the downstream task. The whole architecture is trained in an end-to-end fashion.

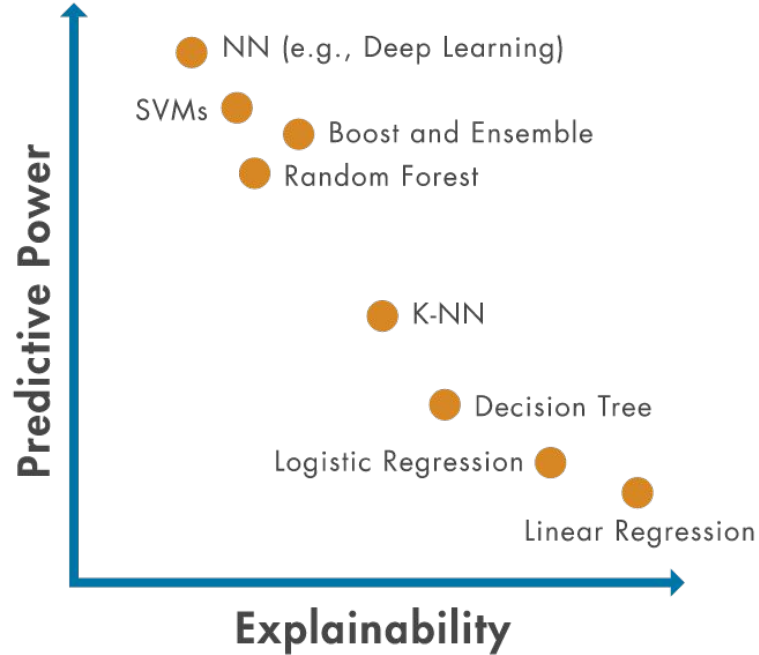
# Introduction

A primer on **explainability**

# What is explainable AI (XAI)?

1. Broad field concerning the development of tools to increase **trust** and **understanding** of a model's predictions.
2. *Common folk dichotomy*: intrinsically interpretable models (e.g., linear regression, decision trees) are orthogonal to models with strong representational power (e.g., deep networks).

# Explainability vs. accuracy?



# Who is explainability for?

1. Most XAI methods are targeted towards **practitioners** of the same methods (i.e., they are akin to debugging tools).
2. It is much harder to target XAI tools towards **end-users** (e.g., clinical staff). Different methods may disagree on the “explanation”, they may not be accurate, and they lack principled **evaluation metrics**.



Zachary Lipton 

@zacharylipton



The precarious state of “interpretable deep learning” is that we should be far more scared upon hearing that a hospital or government deploys any such technique than upon hearing that they haven’t.


[Traduci il Tweet](#)

1:07 AM · 2 feb 2022 · Twitter for iPhone

**73** Retweet **9** Tweet di citazione **520** Mi piace




# XAI categorization




1. XAI tools can be categorized depending on whether they provide **global** or **local** explanations.
2. Some methods are **model-agnostic** (they only need the outputs of the models), other are **model-specific**.
3. Finally, methods can be categorized depending on what type of information they provide in output.

# Global vs. local explanations


1 Complex model and Point of Interest



2 Approximate with Simple Model



3 "Explain" = drivers within area of interest




# Introduction

**Local explanations** for graph NNs

# Explaining a graph NN

Consider a trained graph NN. In the majority of works, a **local explanation** is a small subgraph (and subset of features) which “explain” the prediction.



# Masked predictions

$$f(\mathbf{M}_X \odot \mathbf{X}, \mathbf{M}_A \odot \mathbf{A}) \approx f(\mathbf{X}, \mathbf{A})$$

Binary masks




Original prediction



A “good” explanation should have specific properties (e.g., smallest possible masks).

# Evaluating explanations



**Figure 1: Schematic of attribution task setup and attribution metrics.** A. We create classification and regression tasks for which we have a computable ground-truth. We train GNN models on these labels, and calculate attributions using the graph inputs and attribution methods we adapt to graphs. B. We quantify attribution performance with four metrics. *Accuracy* measures how well an attribution matches ground-truth. *Consistency* measures how accuracy varies across different hyperparameters of a model. *Faithfulness* measures how well the performance of an attribution method matches model performance. *Stability* measures how attributions change when the input is perturbed.

# “Classic” XAI methods

**Attribution** techniques


# Saliency maps

A **saliency map** is an object of the same dimensionality as the input, providing information about which features were most important for a given prediction.

Formally ( $i$  is the index of the class of interest):

$$\text{Saliency map} = \max_{\text{channels}} \left| \frac{\partial f_i(x)}{\partial x} \right|$$





# Saliency maps for graphs

A very similar procedure can be done for a graph NN, to obtain graph saliency maps:

$$\mathbf{M}_A = \frac{\partial f(\mathbf{X}, \mathbf{A})}{\partial \mathbf{A}}$$

(We focus mostly on edge saliency maps, as they are easier to visualize.)

# Example





Figure 8: **Graph 1: nodes 0, 1, 2 and 3 are initially sick; nodes 4 and 11 are immune; the others are healthy.** After one propagation step, the infection reaches nodes 5, 6, 7 and 10. The network predicts the correct label for every node of the graph, following the spread of the infection along non-virtual edges to non-immune nodes. The figures that follow are a visualization of the explanations produced for nodes: 10, 13, 4

# Example



Sensitivity Analysis (SA)

# Limits of saliency maps

Simple saliency maps have issues that balances their simplicity:

1. They are highly **unstable** wrt small changes in the input.
2. They are not well **localized**.
3. They have no formal guarantees.

In particular, they do not respect a property called **sensitivity**: if two inputs differ for a single element but have different predictions, a saliency map is not guaranteed to highlight that pixel.

# Integrated gradients


In the CV field, **integrated gradients** are a powerful alternative to standard saliency maps.

They recover sensitivity by integrating the gradients along a path moving from the empty image to the current one.

We can do something similar by considering an empty adjacency matrix  $\mathbf{A}_0$ :

$$\mathbf{M}_A = (\mathbf{A} - \mathbf{A}_0) \int_{\alpha \in [0,1]} \frac{\partial f(\mathbf{X}, \alpha \mathbf{A} + (1 - \alpha) \mathbf{A}_0)}{\partial \mathbf{A}}$$

# Some benchmarks



**Figure 2: Example ground truth attributions for each task.** The first four graph-classification tasks require a model to identify all nodes (green) in one or more subgraphs (colored lasso) in molecular graphs. Each graph may have multiple positive ground-truths, shown in the *Benzene* task. Ground truth attributions for the *CrippenLogP* regression task take on continuous values. Lower row has node-classification tasks. Relevant subgraphs are circled. Only one neighborhood of the graph is shown.

# Some benchmarks (2)

Graph-level tasks												
	Benzene				Amine AND Ether AND Benzene				CrippenLogP			
	GCN	MPNN	GraphNets	GAT	GCN	MPNN	GraphNets	GAT	GCN	MPNN	GraphNets	GAT
Random Baseline	0.61	0.61	0.61	0.61	0.5	0.5	0.5	0.5	0.13	0.13	0.13	0.13
GradInput	0.72	0.54	0.54	0.56	0.52	0.53	0.55	0.41	0.12	0.09	0.13	0.1
SmoothGrad(GI)	0.71	0.54	0.54	0.53	0.51	0.55	0.59	0.38	0.15	0.11	0.15	0.11
GradCAM-last	0.74	0.72	0.66	0.66	0.54	0.74	0.55	0.46	0.04	0.33	0.24	0.07
GradCAM-all	0.75	0.68	0.84	0.62	0.54	0.62	0.7	0.44	0.05	0.27	0.27	0.09
IG	0.97	0.89	0.94	0.95	0.69	0.59	0.72	0.54	0.31	0.24	0.24	0.27
CAM	0.98	0.96	0.76	0.99	0.75	0.76	0.6	0.65	0.2	0.37	0.28	0.23
Attention Weights	--	--	--	0.51	--	--	--	0.51	--	--	--	-0.06

Node-level tasks												
	BA-Shapes				BA-Community				Tree-Grid			
	GCN	MPNN	GraphNets	GAT	GCN	MPNN	GraphNets	GAT	GCN	MPNN	GraphNets	GAT
Random Baseline	0.27	0.27	0.27	0.27	0.38	0.38	0.38	0.38	0.62	0.62	0.62	0.62
GradInput	0.58	0.64	0.39	0.72	0.52	0.51	0.5	0.5	0.65	0.71	0.66	0.67
SmoothGrad(GI)	0.58	0.64	0.39	0.72	0.52	0.51	0.51	0.49	0.65	0.71	0.66	0.67
GradCAM-last	0.79	0.84	0.86	0.8	0.7	0.67	0.68	0.61	0.7	0.77	0.81	0.7
GradCAM-all	0.67	0.78	0.65	0.76	0.67	0.71	0.73	0.57	0.68	0.7	0.67	0.68
IG	--	--	--	--	0.81	0.75	0.72	0.62	--	--	--	--
Attention Weights	--	--	--	0.5	--	--	--	0.5	--	--	--	0.49

**Figure 3: Attribution method accuracy across tasks and model architectures.** Colors are used to distinguish two metric types — attribution AUROC for attribution on classification tasks, and attribution Kendall’s tau on the regression task. CAM and IG perform consistently well across tasks and models. For error bars, please see Figure S1.



# Tailored XAI methods

**GNNExplainer**

# GNNExplainer

For graphs, it is fundamental that the resulting subgraph is small, since visualizing and interpreting it is complex.

For this reason, tailored methods have been proposed to force sparsity as the main concern.

**GNNExplainer** works by optimizing the masks using the following criteria:

1. Keeping the original prediction consistent;
2. Having small masks (l1 regularization);

# GNNExplainer cost function

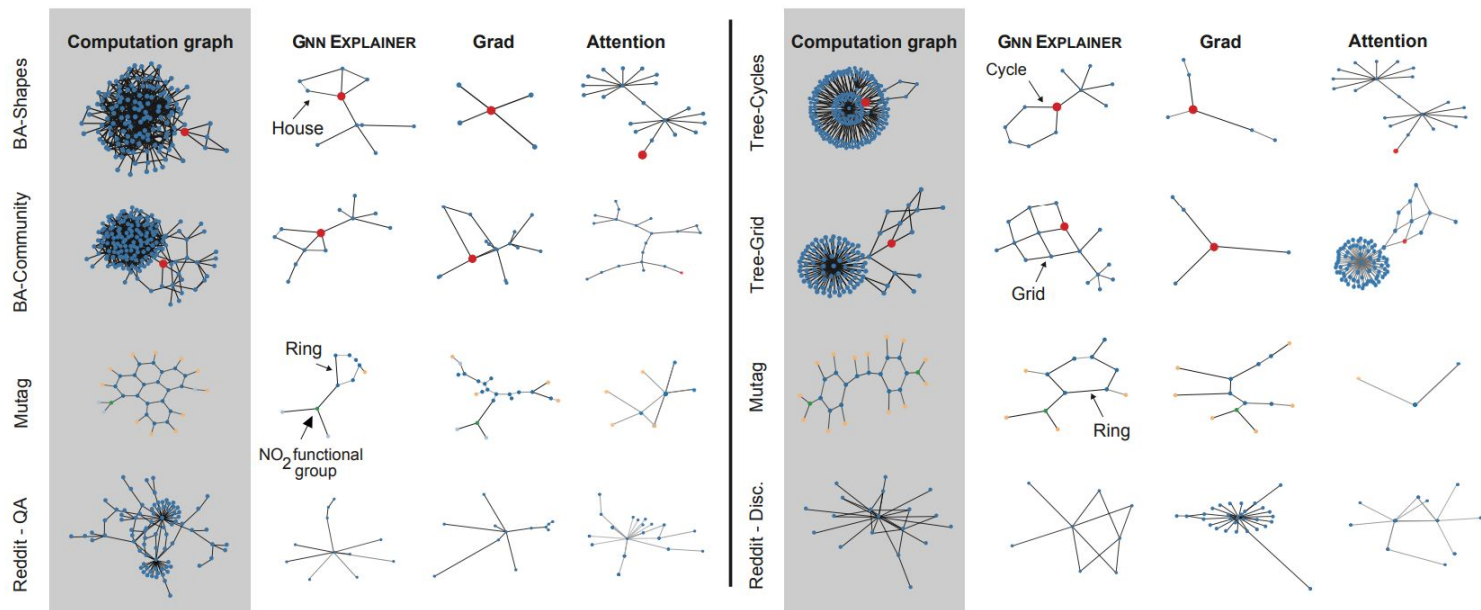
$$\mathbf{M}_A = \arg \min_{\mathbf{M}} \text{CE}(f(\mathbf{X}, \mathbf{A}), f(\mathbf{X}, \mathbf{M} \odot \mathbf{A})) + \alpha \|\mathbf{M}\|_1 + \beta \text{H}[\mathbf{M}]$$

Prediction should stay consistent

The mask should be **sparse**

Weights should be as close as possible to 0 or 1

# Some results



**Figure 3: Examples of single-instance important subgraphs. The red node is the explained node.**

# PGExplainer




Figure 2: Illustration of PGExplainer for explaining GNNs on graph classification. (1) The left part demonstrates the explanation network. It takes node representations  $\mathbf{Z}$  as well as the original graph  $G_o$  as inputs to compute  $\Omega$ , the latent variables in edge distributions. Edge distributions are served as the explanation. In case that an explanatory subgraph is wanted, we select top-ranked edges according to latent variables  $\Omega$ . (2) A random graph  $\hat{G}_s$  is sampled from edge distributions and then feed to the trained GNN model to get the prediction  $\hat{Y}_s$ . (3) Parameter  $\Psi$  in the explanation network is optimized with cross-entropy between the original prediction  $Y_o$  and the updated prediction  $\hat{Y}_s$ .


# Notebook time!

[Colab Notebooks and Video Tutorials — pytorch\\_geometric documentation](#)

<https://colab.research.google.com/drive/1nV44NrNqcXC2thU6-zzxnJPnlalo870m?usp=sharing>

**Selected topics**

**Data influence**





# Limits of saliency maps

1. All the previous methods estimated the influence of single **features** (e.g., edges) on the prediction.
2. An alternative class of methods explore the influence of single **data points** on the prediction, e.g., how much training on a certain graph (or node) has influenced the prediction on a separate graph (or node).
3. This is a more complex scenario, since the influence has to be computed across the entire training run.

# Gradient tracing

Consider an idealized training procedure where at iteration  $t$  we update the parameter vector as:

$$w_{t+1} = w_t - \eta \nabla l(w_t, z_t)$$

The **influence** of point  $z$  on point  $z'$  is defined as:

$$\text{TracInIdeal}(z, z') = \sum_{t: z_t = z} l(w_t, z') - l(w_{t+1}, z')$$

# Gradient tracing

By first-order approximation, it can be shown that:

$$\text{TracInIdeal}(z, z') \approx \sum_{t: z_t = z} \eta \nabla l(w_t, z) \cdot \nabla l(w_t, z')$$

This can be approximated by storing  $k$  checkpoints during training and computing:

$$\text{TracInIdeal}(z, z') \approx \sum_{i=1}^k \eta \nabla l(w_i, z) \cdot \nabla l(w_i, z')$$

# Results on CV

		<u>Proponents</u>					<u>Opponents</u>				
		3 (3)	3 (5)	3 (3)	3 (8)	3 (3)	5 (5)	5 (3)	5 (5)	5 (5)	5 (5)
	<b>TracIn</b>										
		3 (5)	3 (7)	3 (9)	3 (8)	3 (8)	7 (3)	8 (3)	2 (3)	8 (3)	8 (3)
	<b>Representer</b>										
		3 (5)	3 (9)	3 (5)	3 (5)	3 (5)	8 (9)	8 (3)	4 (4)	9 (9)	3 (5)
	<b>Influence</b>										

(a) Correctly classified 3.

# Results on CV

		<u>Proponents</u>					<u>Opponents</u>				
	<b>TracIn</b>	6 (5)	6 (5)	6 (5)	6 (5)	6 (5)	5 (5)	5 (6)	5 (5)	5 (6)	5 (6)
	<b>Representer</b>	6 (5)	6 (5)	6 (0)	6 (0)	6 (0)	0 (6)	0 (6)	5 (6)	8 (6)	8 (6)
	<b>Influence</b>	6 (5)	6 (5)	6 (5)	6 (5)	6 (6)	0 (6)	5 (6)	5 (6)	5 (6)	5 (5)

(b) Incorrectly classified 6

# Results on CV



Figure 5: CIFAR-10 results: Proponents and opponents examples of a correctly classified cat for influence functions, representer point, and TracIn. (Predicted class in brackets)

# Notebook time!

[Colab Notebooks and Video Tutorials — pytorch\\_geometric documentation](#)

<https://colab.research.google.com/drive/1nV44NrNqcXC2thU6-zzxnJPnlalo870m?usp=sharing>

# Selected topics

**New directions** for improving XAI



# Challenges for current XAI on graphs

- 1. Most interpretability methods are trained post-hoc.**
2. Lack of datasets with an interpretability gold standard.
3. Defining proper metrics for assessing the results is non trivial.

# Meta-learning for enhancing XAI

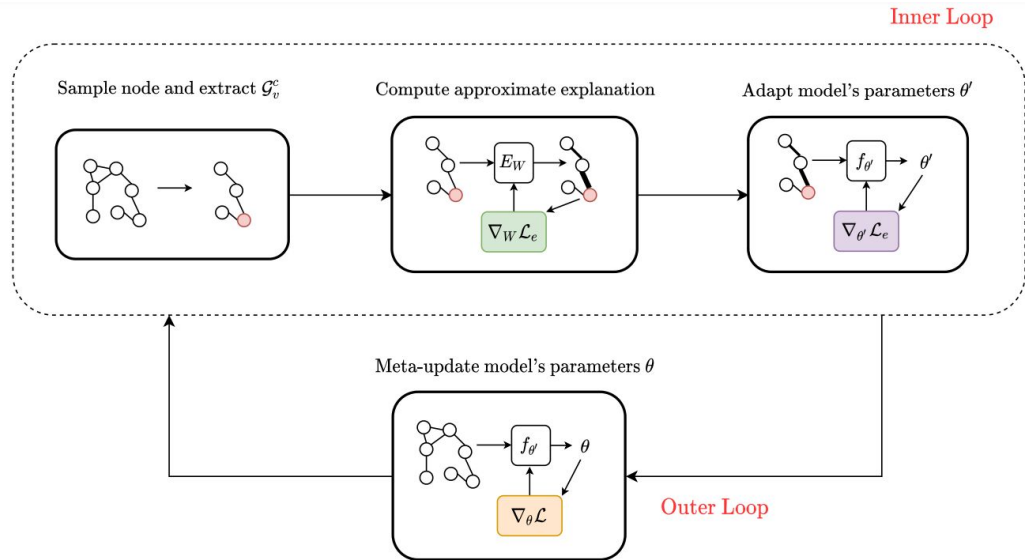


Fig. 3. Schematics of our meta-learning framework for improving GNN's explainability at training time. MATE steers the optimization procedure toward more interpretable minima in the inner loop, meanwhile optimizing for the original task in the outer one. The inner loop adapts the model's parameters to a single "explanation task." It starts with the sampling of a random node and its computational subgraph. Then, we train GNNExplainer to explain the current model's prediction. Afterward, we can adapt the model's parameters to the "explanation task" ending in a new model's state. Finally, we meta-update the original parameters minimizing the cross-entropy loss computed with the adapted parameters.

# Some results

TABLE III  
 VISUALIZATION OF THE EXPLANATION SUBGRAPHS FOR THE NODE CLASSIFICATION TASK. NODE COLORS REPRESENT NODE LABELS. DARKNESS OF THE EDGES SIGNALS IMPORTANCE FOR CLASSIFICATION. THE GROUND-TRUTH MOTIF IS PRESENTED IN THE FIRST ROW

	BA-shapes	BA-community	Tree-cycles	Tree-grids
Motif				
GNNExp				
MATE+GNNExp				
PGExp				
MATE+PGExp				

# Prototype-based GNNs

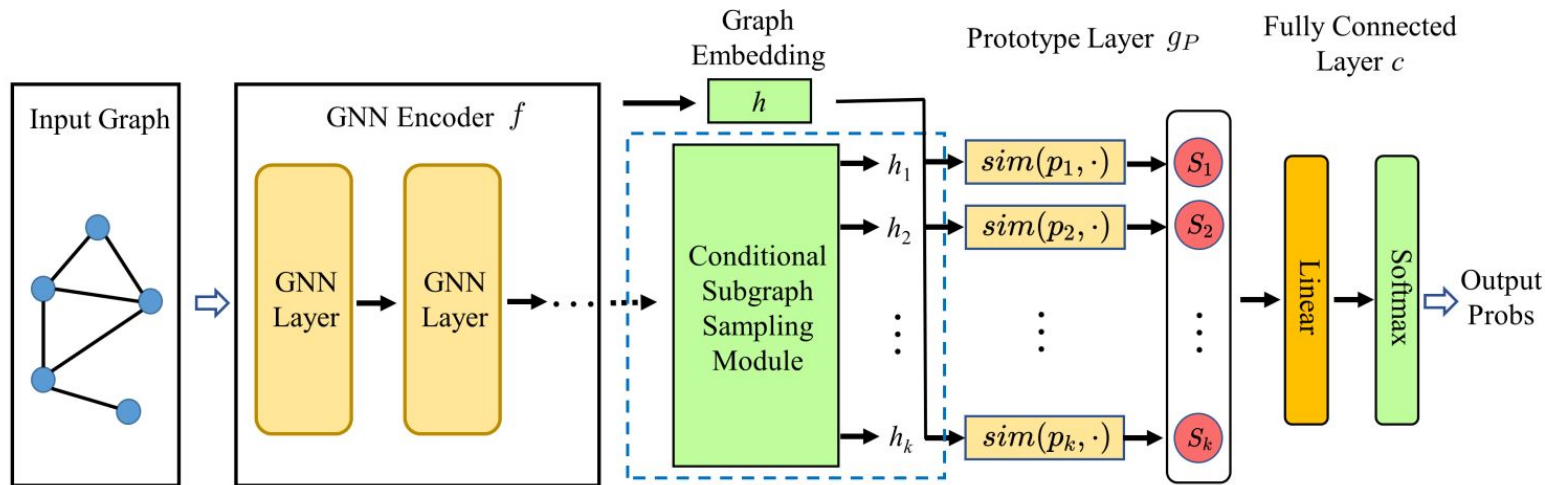


Figure 1: The architecture of our proposed ProtGNN/ProtGNN+. The model mainly consists of three parts: GNN encoder  $f$ , prototype layer  $g_P$ , and the fully connected layer  $c$  appended by softmax to output probabilities in multi-class classification tasks. ProtGNN calculates the similarity score ( $\text{sim}(p_k, \cdot)$  in the illustration) between the graph embedding and the learned prototypes in the prototype layer. For further interpretability, the conditional subgraph sampling module (in the dashed bounding box) is incorporated in ProtGNN+ to output subgraphs most similar to each learned prototype.

# Thanks! Questions?



Simone Scardapane  
**Tenure-track Assistant Professor**



<https://www.scardapane.it/>



[https://twitter.com/s\\_scardapane](https://twitter.com/s_scardapane)