


Geometric Deep Learning

8. Graphs and Sets II*

Jun.-Prof. Dr. Jan Stühmer

Based on the slides by Petar Veličković, AMMI 2022
Available at geometricdeeplearning.com

Recap: Blueprint of Geometric Deep Learning




Recap: Building blocks of Geometric Deep Learning

Let Ω and Ω' be domains, $\mathfrak{G}$ a symmetry group over Ω .
Write $\Omega' \subseteq \Omega$ if Ω' can be considered a compact version of Ω .

We define the following building blocks:

- *Linear $\mathfrak{G}$ -equivariant layer* $B: \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C}')$,
satisfying $B(g.x) = g.B(x)$ for all $g \in \mathfrak{G}$ and $x \in \mathcal{X}(\Omega, \mathcal{C})$.
- *Nonlinearity* $\sigma: \mathcal{C} \rightarrow \mathcal{C}'$ applied element-wise as $(\sigma(x))(u): \sigma(x(u))$.
- *Local pooling (coarsening)* $P: \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C})$, such that $\Omega' \subseteq \Omega$.
- *$\mathfrak{G}$ -invariant layer (global pooling)* $A: \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{Y}$,
satisfying $A(g.x) = A(x)$ for all $g \in \mathfrak{G}$ and $x \in \mathcal{X}(\Omega, \mathcal{C})$.

Recap: The three “flavors” of GNN layers



[Bronstein, Bruna, Cohen, Veličković, 2021]

$$\mathbf{h}_i = \phi \left(\mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} c_{ij} \psi(\mathbf{x}_j) \right)$$

$$\mathbf{h}_i = \phi \left(\mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} a(\mathbf{x}_i, \mathbf{x}_j) \psi(\mathbf{x}_j) \right)$$

$$\mathbf{h}_i = \phi \left(\mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j) \right)$$

Graph Networks

General attributed graphs

- Graphs can convey rich information at all granularities
- We previously ignored non-node data for simplicity
- Now we will assume the following possible set of features:
 - Node features, $\mathbf{x}_u \in \mathbb{R}^k$ (e.g. atom type, charge, nb. of hydrogen bonds)
 - Edge features, $\mathbf{x}_{uv} \in \mathbb{R}^l$ (e.g. bond type, is in a ring?)
 - Graph features $\mathbf{x}_G \in \mathbb{R}^m$ (e.g. molecular weight, fingerprints)
- Analogously defining latents $\mathbf{h}_u, \mathbf{h}_{uv}, \mathbf{h}_G$
- It is possible to extend this further (e.g. *hypergraphs*)
but such input can usually be represented as an instance of the above

Graph Networks

- We will now define a general blueprint for a spatial GNN, which generalizes all the flavors referenced before
- We use the Graph Network (Battaglia et al., 2018) as our basis
This is because it operates over generic attributed graphs
- Dataflow:
 - Update edge features (using graph + relevant nodes)
 - Update node features (using updated relevant edges + graph)
 - Update graph features (using updated nodes + edges)
- and extensive usage of *skip connections*!

Graph Networks

- Update edge features (using graph + relevant nodes)

$$\mathbf{h}_{uv} = \psi(\mathbf{x}_u, \mathbf{x}_v, \mathbf{x}_{uv}, \mathbf{x}_{\mathcal{G}})$$


- Update node features (using updated relevant edges + graph)

$$\mathbf{h}_u = \phi\left(\mathbf{x}_u, \bigoplus_{u \in \mathcal{N}_v} \mathbf{h}_{vu}, \mathbf{x}_{\mathcal{G}}\right)$$


- Update graph features (using updated nodes + edges)

$$\mathbf{h}_{\mathcal{G}} = \rho\left(\bigoplus_{u \in \mathcal{V}} \mathbf{h}_u, \bigoplus_{(u,v) \in \mathcal{E}} \mathbf{h}_{uv}, \mathbf{x}_{\mathcal{G}}\right)$$

Graph Networks, visualized




Graph Networks, visualized




$$\mathbf{h}_{uv} = \psi(\mathbf{x}_u, \mathbf{x}_v, \mathbf{x}_{uv}, \mathbf{x}_g)$$

Graph Networks, visualized




$$\mathbf{h}_u = \phi \left(\mathbf{x}_u, \bigoplus_{v \in \mathcal{N}_v} \mathbf{h}_{vu}, \mathbf{x}_G \right)$$

Graph Networks, visualized



$$\mathbf{h}_G = \rho \left(\bigoplus_{u \in \mathcal{V}} \mathbf{h}_u, \bigoplus_{(u,v) \in \mathcal{E}} \mathbf{h}_{uv}, \mathbf{x}_G \right)$$

Graph Networks, visualized



Equivariant and **invariant** layers

Latent Graph Inference

- So far, we've the “ground-truth” graph to be given
- However, this it not always the case
- This leads to laten graph inference, a key emerging area of graph representation learning (and an area of high interest to *science*)
- Further, studying this area will reveal surprising connections between state-of-the-art models and graph neural networks

In the absence of a graph ...

- Even when the input graph is absolutely correct, it does not mean it is **optimal** for the task at hand
- Hence, *rewiring* might be a good idea
- Taken to the extreme, *what to do when there is no graph?* (but we assume an underlying graph may exist)
- Assume we are given a node feature matrix, $\mathbf{X}$, but no adjacency
- We have seen a solution for this already ...

Option 1: Empty edge set

- Assume the graph is completely disconnected
Let $\mathbf{A} = \mathbf{I}$, or, equivalently, $\mathcal{N}_u = \{u\}$ for all nodes u

- Each element is processed independently:

$$\mathbf{h}_u = \psi(\mathbf{x}_u)$$

- This, as we have seen, results in a permutation equivariant function over the set
- **Deep Sets** [Zaheer *et al.*, Neurips 2017] implements such an architecture

Option 2: Complete graph

- Assume the graph is *fully connected*
Let $\mathbf{A} = \mathbf{1}\mathbf{1}^T$, or, equivalently, $\mathcal{N}_u = \mathcal{V}$ for all nodes u
- If we do not model any information about the edges, i.e. we do not have access to any coefficients depending on the vertices interacting along the edges, any convolutional GNN is **equivalent** to Deep Sets:

$$\mathbf{h}_u = \phi \left(\mathbf{x}_u, \bigoplus_{v \in \mathcal{V}} \psi(\mathbf{x}_v) \right)$$

- This is because the edges are indistinguishable, and the complete edge set of the graph permutation invariant

Option 3: Complete graph (with attention)

- Assume the graph is *fully connected*

Let $\mathbf{A} = \mathbf{1}\mathbf{1}^T$, or, equivalently, $\mathcal{N}_u = \mathcal{V}$ for all nodes u


- This motivates the use of a more expressive GNN flavor, the attentional,

$$\mathbf{h}_u = \phi \left(\mathbf{x}_u, \bigoplus_{v \in \mathcal{V}} a(\mathbf{x}_u, \mathbf{x}_v) \psi(\mathbf{x}_v) \right)$$

which yields the *self attention* operator, the core of the Transformer architecture [Vaswani *et al.* Neurips 2017]

A Note on Transformers

- Transformers **are** Graph Neural Networks:
 - Fully-connected graph
 - Attentional “flavor”
- The sequential structural information is injected through the **positional embeddings**
- Dropping them yields a fully-connected Graph Attention Network (GAT)
- Attention amounts to inferring **soft adjacency**



Option 4: Complete graph with message passing

- When nodes interact in more complex ways (e.g. physics, relational reasoning), the message passing flavor may be suitable
- For physics modelling: Interaction Nets [Battaglia *et al.*, NeurIPS'16]
treat every interacting body as a node in a fully connected graph
- For visual reasoning: Relational Networks [Santoro *et al.*, NeurIPS'17]
treat every pixel representation (after CNN as feature extractor) as a node in a fully connected graph

Empty edge set vs Fully connected graph


- Empty graph ignores a potential wealth of information
- Full graph can be hard to scale (quadratic space, large neighborhoods)
- Ideally, want to infer the adjacency A , commonly termed “latent graph inference”
- Five recently proposed prominent ways to do so:
 - Variational [Kipf, Fetaya *et al.*, ICML’18] NRI
 - No learning [Wang *et al.*, ACM TOG’18] DGCNN
 - Reinforcement Learning [Kazi *et al.*, 2020] DGM
 - Supervised [Veličković *et al.*, NeurIPS’20] PGN
 - Self-supervised [Fatemi *et al.*, NeurIPS’21] SLAPS

How powerful are graph neural networks?

- GNNs are a powerful tool for processing real-world graph data
But they won't solve *any* task specified on a graph accurately!
- Canonical example: deciding *graph isomorphism*
- Can the GNN **distinguish** two *non*-isomorphic graphs? $\mathbf{h}_{\mathcal{G}_1} \neq \mathbf{h}_{\mathcal{G}_2}$
If not, any kind of task discriminating them is hopeless
- Assess power of GNNs by which graphs they are able to distinguish.
- To formalise this, we will study a popular graph isomorphism test

Weisfeiler-Lehman Test


- Simple but powerful way of distinguishing: pass random hashes of sums along edges
- Iterate until hashes don't change
"Possibly isomorphic" if histograms are same



A. Lehman




B. Weisfeiler




Weisfeiler, Lehman 1968; Portraits: Ihor Gorskiy

Weisfeiler-Lehman Test




Weisfeiler, Lehman 1968

Weisfeiler-Lehman Test



Weisfeiler, Lehman 1968


Weisfeiler-Lehman Test



Weisfeiler, Lehman 1968




Weisfeiler-Lehman Test



Weisfeiler, Lehman 1968

Weisfeiler-Lehman Test

- The connection to GNNs was made already early, e.g. in GCN
- **Untrained** GNNs can work very well: Random features are equivalent to hash in Weisfeiler-Lehman test
- The test only provides a necessary but insufficient condition for graph isometry and can fail, e.g. for



Expressivity of GNNs


- For discrete node features, GNNs can only be **as powerful** as the Weisfeiler-Lehman test
- Condition for maximal expressivity is an *injective* aggregation (e.g. sum)
The proof is similar to one used to show Deep Sets universality
- Graph isomorphism networks [Xu *et al.* ICLR 2019] propose a simple maximally-expressive GNN

$$\mathbf{h}_v = \phi \left((1 + \epsilon) \mathbf{h}_v + \sum_{u \in \mathcal{N}_v} \mathbf{h}_u \right)$$


The real valued coefficient ϵ results in a node-individual feature

Augmented Message Passing

- We can improve GNNs by analyzing the **failure cases** of Weisfeiler-Lehman




- For example, like Weisfeiler-Lehman, GNNs cannot detect closed triangles



Augmented Message Passing

- We can improve GNNs by analyzing the **failure cases** of Weisfeiler-Lehman
- Augment nodes with randomized features [Sato *et al.* 2021]
- Now a node can “see itself” k hops away!



(b) Random Features.

Augmented Message Passing

- We can improve GNNs by analyzing the **failure cases** of Weisfeiler-Lehman
- Augment nodes with randomized features [Sato *et al.* 2021]
See also: RP-GNN [Murphy *et al.* ICML'19], P-GNN [You *et al.* ICML'19]
- These methods belong to *augmented message passing*
- They can be broadly categorized into three groups
 - Modifying **features** (as above)
 - Modifying the **message-passing rule**, e.g. DGN [Beaini *et al.* ICML'20]
 - Modifying the **graph** structure, e.g. 1-2-3-GNNs [Morris *et al.* AAAI'19]

What have we covered?

- A deep dive into graph neural networks (GNNs)
- The Graph Network architecture over generic attributed graphs
- How do *Transformers* and *Deep Sets* fit within the GNN flavors
- (Breaking) the limits of GNN expressive power: *graph isomorphism testing, augmented message passing*