

Outline

- We now move to Sets to Graphs
- Sets can be considered as "node-only" Graphs without edges
- Graphs are a very versatile abstraction used ubiquitously in sciences
- For many years, Geometric Deep Learning was nearly synonymous with Graph Representation Learning

any scientific fields study data with an underlying structure that is non-Euclidean. Some examples include social networks in computational social sciences, sensor networks in computational social sciences, sensor networks in computer graphics. In many applications, such geometric data are large and complex (in the case of social networks, on the scale of billions) and are natural targets for machine-learning techniques. In particular, we would like to use deep neural networks, which have recently proven to be powerful tools for a broad range of problems from computer vision, natural-language processing, and audio analysis. However, these tools have been most successful on data with an underlying Euclidean or grid-like structure and in cases where the invariances of these

Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst

structures are built into networks used to model them. *Geometric deep learning* is an umbrella term for emerging techniques attempting to generalize (structured) deep neural models to non-Euclidean domains, such as graphs and manifolds. The purpose of this article is to overview different examples of geometric deep-learning problems and present available solutions, key difficulties, applications, and future research directions in this nascent field.

Overview of deep learning

Deep learning refers to learning complicated concepts by building them from simpler ones in a hierarchical or multilayer manner. Artificial neural networks are popular realizations of such deep multilayer hierarchies. In the past few years, the growing computational power of modern graphics processing unit (GPU)-based computers and the availability of large training data sets have allowed successfully training neural networks with many layers and degrees of freedom (DoF) [1]. This has led to qualitative breakthroughs on a wide variety of tasks, from speech recognition [2], [3] and machine translation [4] to image analysis and computer vision [5]–[11] (see [12]

Geometric Deep Learning

Going beyond Euclidean data

Graphs = *Systems of Relations and Interactions*



Molecules

Interactomes

Social networks

Why Graphs?

- Simple mathematical abstraction for systems of relations or interactions
- Very well developed theory
- Many successful applications from chemistry to social media to even pure math
- In many ways, graphs are the main modality of data we receive from Nature
- Graph representation learning is likely critical on the path to Artificial General Intelligence (AGI).

Why Graphs?

"The image of the world around us, which we carry in our head, is just a *model*. Nobody in his head imagines all the world, government or country. He has only selected *concepts*, and *relationships* between them, and uses those to represent the real system."

—Jay Forrester on mental models (1971)



Jay Wright Forrester

Forrester 1971

Origins of Graph Theory



The soluton of the classical problem of the "Bridges of Königsberg" by Euler in 1736 first showed the power of graphs to abstract out the geometry ("geometria situs")

1741

Euler 1741

Origins of Topology

Poincaré's *"analysis situs."* His famous Conjecture appeared in a supplement published in 1904.



Euler 1741; Poincaré 1895

Origin of the name "Graph"

CHEMISTRY AND ALGEBRA

I T may not be wholly without interest to some of the readers of NATURE to be made acquainted with an analogy that has recently forcibly impressed me between branches of human knowledge apparently so dissimilar as modern chemistry and modern algebra.

The weight of an invariant is identical with the number of the bonds in the chemicograph of the analogous chemical substance, and the weight of the leading term (or basic differentiant) of a co-variant is the same as the number of bonds in the chemicograph of the analogous compound radical. Every invariant and covariant thus becomes expressible by a graph precisely identical with a Kekuléan diagram or chemicograph.

Baltimore, January I J. J. SYLVESTER

The term "graph" appeared first in the chemical context



J. Sylvester

1878

Sylvester 1878

GRAPHS: THE BASICS

Types of Learning Tasks on Graphs







Graph regression water solubility?

Node classification who is a spammer

Link prediction WTF (whom to follow)?

Invariant vs Equivariant tasks





who is a spammer?

GNNs = *Parametric graph functions*



Note: we use the term "GNN" to refer to general parametric graph functions. The particular class of GNNs we consider are Message Passing Neural Networks (MPNNs). Petar Veličković argues that "it's all message passing"

Types of Learning Tasks on Graphs



Transductive vs Inductive tasks



Training







Inductive different graph

Graphs: The Basics

- A graph is a pair G = (V, E)
 - *V* = **vertices** (or **nodes**)
 - $E \subseteq V \times V =$ **edges** (or **links**). If $(u, v) \in E$, we will write $u \sim v$
- Edges in general are **directed**, i.e., we might have *u~v* but not *v~u*.



Graphs: The Basics

- A graph is a pair G = (V, E)
 - *V* = **vertices** (or **nodes**)
 - $E \subseteq V \times V =$ edges (or links). If $(u, v) \in E$, we will write $u \sim v$
- Edges in general are **directed**, i.e., we might have *u~v* but not *v~u*.
- Nodes and edges can have **attributes**. We will typically assume
 - *d*-dimensional vector **node features**, denoted **x**_{*u*}
 - scalar **edge weights**, denoted *w*_{*uv*}



Extensions of Graphs







Multigraph

Hypergraph

Cellular complex

Representations of Graphs as matrices



- Nodes indexed as $V = \{1, ..., n\}$ in *arbitrary order*
- Node features stacked row-wise in an $n \times d$ matrix **X**

Representations of Graphs as matrices



• Nodes indexed as $V = \{1, ..., n\}$ in *arbitrary order*

- Node features stacked row-wise in an $n \times d$ matrix **X**
- Graph structure represented by $n \times n$ adjacency matrix **A** with $a_{ij} = 1$ iff $i \sim j$

Representations of Graphs as matrices



The ordering of nodes is arbitrary!

What do we want now with Graphs?



- The function now also depends on edges (adjacency **A**) in addition to node features **X**
- Permutation invariance: $f(\mathbf{P}\mathbf{X}, \mathbf{P}\mathbf{A}\mathbf{P}^{\mathrm{T}}) = f(\mathbf{X}, \mathbf{A})$
- Permutation equivariance:

 $\mathbf{F}(\mathbf{P}\mathbf{X}, \mathbf{P}\mathbf{A}\mathbf{P}^{\mathrm{T}}) = \mathbf{P}\mathbf{F}(\mathbf{X}, \mathbf{A})$

Invariant Graph Functions



Equivariant Graph Functions



DeepSets



On sets, we can only process each node independently (or all together)

Graph Neural Networks



On graphs, we have a notion of a **neighbourhood** of ever node $\mathcal{N}_i = \{j \in V : i \sim j\}$

Neighbour Aggregation



Neighbour Aggregation



GNN Layer

$$\mathbf{F}(\mathbf{X}, \mathbf{A}) = \begin{pmatrix} -\phi(\mathbf{x}_1, \mathbf{X}_{\mathcal{N}_1}) - \\ \vdots \\ -\phi(\mathbf{x}_i, \mathbf{X}_{\mathcal{N}_i}) - \\ \vdots \\ -\phi(\mathbf{x}_n, \mathbf{X}_{\mathcal{N}_n}) - \end{pmatrix}$$

permutation equivariant

GNNs as an Instance of Geometric Deep Learning Blueprint

Graph G = (V, E)

Node features $\mathcal{X}(G)$

Functions $\mathcal{F}(\mathcal{X}(G))$







Permutation group $\pi \in S_n$ Permutation matrix **PX**, **PAP**^{\top}

Message passing $F(PX, PAP^{\top}) = PF(X, A)$

GNN FLAVOURS










Convolutional GNNs

- Simplest GNN
- Highly scalable
- Industrial use cases
- Folklore: works only on homophilic graphs





Rossi, Frasca et B 2020 (SIGN); Ying et al. 2018 (PinSAGE)

Attentional GNNs



weights

Monti et B 2017; Veličković et al. 2018 (GAT)

Message-Passing GNNs



message from node j to node i

Gilmer et al. 2017 (MPNN); Battaglia et al. 2018 (Graph Networks) Wang et B, Solomon 2018 (edgeconv)

Flavours of GNNs



Flavours of GNNs



MAIN INGREDIENTS OF AN MPNN



$$\mathbf{x}_i \leftarrow \phi\left(\mathbf{x}_i, \prod_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j)\right)$$



$$\mathbf{x}_i \leftarrow \phi\left(\mathbf{x}_i, \prod_{j \in \mathcal{N}_i} \boldsymbol{\psi}(\mathbf{x}_i, \mathbf{x}_j)\right)$$

Message passing function



$$\mathbf{x}_i \leftarrow \phi\left(\mathbf{x}_i, \bigcup_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j)\right)$$

Message passing function

Aggregation operator





$$\mathbf{x}_i \leftarrow \phi\left(\mathbf{x}_i, \prod_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j)\right)$$

Message passing function

Aggregation operator

Computational graph



$$\mathbf{x}_i \leftarrow \phi\left(\mathbf{x}_i, \prod_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j)\right)$$

Message passing function

Aggregation operator

Computational graph

EXPRESSIVE POWER OF GNNS

Graph isomorphism



• Two graphs G = (V, E) and G' = (V', E') are **isomorphic** if there exists an *edge-preserving bijection* $\varphi: V \to V'$ s.t. $u \sim v$ in *G* iff $\varphi(u) \sim \varphi(v)$ in *G'*

Graph isomorphism



• Two graphs G = (V, E) and G' = (V', E') are **isomorphic** if there exists an *edge-preserving bijection* $\varphi: V \to V'$ s.t. $u \sim v$ in *G* iff $\varphi(u) \sim \varphi(v)$ in *G'*

Note: φ is not unique where the graph has symmetries (edge-preserving automorphism)

Graph isomorphism



- Two graphs G = (V, E) and G' = (V', E') are **isomorphic** if there exists an *edge-preserving bijection* $\varphi: V \to V'$ s.t. $u \sim v$ in *G* iff $\varphi(u) \sim \varphi(v)$ in *G'*
- A set of graphs isomorphic to each other is an isomorphism class
- Complexity of computing graph isomorphism is an open question ("*NP intermediate*": not NP but also no polynomial time algorithm currently known)

Attributed graph isomorphism



- Two node-attributed graphs $G = (V, E, \mathbf{X})$ and $G' = (V', E', \mathbf{X}')$ are **isomorphic** if there exists a *bijection* $\varphi: V \to V'$ s.t.
- Structure preservation: $u \sim v$ in G iff $\varphi(u) \sim \varphi(v)$ in G'
- Feature preservation: $\mathbf{x}_u = \mathbf{x}'_{\varphi(u)}$

Universal Approximation on Graphs

Theorem: A class of functions is universally approximating permutationinvariant functions on graphs with finite node features iff it can discriminate graph isomorphisms.

Universal approximation on graphs is equivalent to graph isomorphism testing

Chen 2019



What graphs can MPNNs represent?

Adapted from Bevilacqua et al. (LoG Tutorial 2022)





Expressive power of MPNNs

Functions that can be computed by MPNN

All permutationinvariant functions

Adapted from Bevilacqua et al. (LoG Tutorial 2022)

All graph-isomorphism discriminating functions

WEISFEILER-LEHMAN TEST





- A. Le(h)man
- **B.** Weisfeiler

















non-isomorphic graphs that are WL-equivalent

Necessary but insufficient condition for graph isomorphism!

What does WL test see?



What WL cannot test?

Example of non-isomorphic graphs that cannot be distinguished by Weisfeiler-Lehman test (outputs "possibly isomorphic")



r-*r*egular graphs (deg=r at every node) with the same number of nodes

What WL cannot test?

Example of non-isomorphic graphs that cannot be distinguished by Weisfeiler-Lehman test (outputs "possibly isomorphic")



Any induced connected pattern with ≥ 3 nodes (triangles, cycles, etc.)

What WL cannot test?

Example of non-isomorphic graphs that cannot be distinguished by Weisfeiler-Lehman test (outputs "possibly isomorphic")



Important implications e.g. in chemistry!

Expressive power of Weisfeiler-Lehman



Adapted from Bevilacqua e (LoG Tutorial 2022)

MPNNs vs WL

• WL-test: $\mathbf{x}_i \leftarrow \phi(\mathbf{x}_i, \{\!\!\{\mathbf{x}_j : j \in \mathcal{N}_i\}\!\!\})$ where ϕ is *injective* (hash function)

• MPNN:
$$\mathbf{x}_i \leftarrow \phi\left(\mathbf{x}_i, \prod_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j)\right)$$



MPNN expressive power is *upper-bounded* by the Weisfeiler-Lehman test
MPNNs vs WL

• WL-test: $\mathbf{x}_i \leftarrow \phi(\mathbf{x}_i, \{\!\!\{\mathbf{x}_j : j \in \mathcal{N}_i\}\!\!\})$ where ϕ is *injective* (hash function)

• MPNN:
$$\mathbf{x}_i \leftarrow \phi\left(\mathbf{x}_i, \prod_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j)\right)$$



When is MPNN as expressive as the Weisfeiler-Lehman test?

Expressive power of MPNNs

Functions that can be computed by some MPNNs

Functions depending on counts of certain substructures

Functions that can be computed by WL

All permutationinvariant functions

Adapted from Bevilacqua et al. (LoG Tutorial 2022)

Are all Aggregators the same?

Input

max "skeleton" of the multiset

Distribution

mean

sum

of the multiset

Are all Aggregators the same?



mean and **max** fail to distinguish



max fails to distinguish



mean and **max** fail to distinguish

Graph Isomorphism Network (GIN)

Theorem: Assume graph node features are from a countable set. Then, an MPNN with with injective aggregator \Box , update function ϕ , and graph-wise readout function, is as powerful as the Weisfeiler-Lehman test.

- Assumption of discrete countable features (often not the case in practice)
- Proof similar to DeepSets, with the difference that we now deal with *multisets*, where popular injective set functions such as mean are not injective anymore
- GIN: uses an injective multiset function of the form

$$\mathsf{MLP}\left((1+\epsilon)\mathbf{x}_i + \sum_{j\in\mathcal{N}_i}\mathbf{x}_j\right)$$

Xu 2019; (Morris 2019)

Expressive power of GIN ("best MPNN")



(LoG Tutorial 2022)

What to do with continuous features?

- Many practical applications rely on assumption of continuous (uncountable) features
- Most of the results we have seen (DeepSets, GINs, etc.) do not work in this setting!

Theorem: In order to discriminate between real multisets of size *n*, at least *n* aggregators are needed.

Proof: relying on Borsuk-Ulam Theorem (continuous function from \mathbb{S}^n to \mathbb{R}^n maps some pairs of antipodal points to the same point)

Corso et al. 2020

Principal Neighbourhood Aggregation (PNA)

• Use a *combination* of multiple aggregators defined as moments of neighbour features

$$\Box = \begin{bmatrix} I \\ S(D, \alpha = 1) \\ S(D, \alpha = -1) \end{bmatrix} \otimes \begin{bmatrix} \mu \\ \sigma \\ \max \\ \min \end{bmatrix}$$

(*scalers* $S(d, \alpha) = \left(\frac{\log(d+1)}{\delta}\right)^{\alpha}$ emphasize hub nodes and allow for aggregators like sum)

• Empirically good performance

Corso et al. 2020

Principal Neighbourhood Aggregation (PNA)

		Nodes tasks		Graph tasks				
Model	Average score	1	2	3	4	5	6	
PNA	-3.13	-2.89	-2.89	-3.77	-2.61	-3.04	-3.57	Best
PNA (no scalers)	-2.77	-2.54	-2.42	-2.94	-2.61	-2.82	-3.29	
MPNN (max)	-2.53	-2.36	-2.16	-2.59	-2.54	-2.67	-2.87	
MPNN (sum)	-2.50	-2.33	-2.26	-2.37	-1.82	-2.69	-3.52	
GAT	-2.26	-2.34	-2.09	-1.60	-2.44	-2.40	-2.70	
GCN	-2.04	-2.16	-1.89	-1.60	-1.69	-2.14	-2.79	
GIN	-1.99	-2.00	-1.90	-1.60	-1.61	-2.17	-2.66	
Baseline	-1.38	-1.87	-1.50	-1.60	-0.62	-1.30	-1.41	Worst
 Single-source shortest-paths Eccentricity 				4. 5.	Conne Diame			
3. Laplacian features				6.	Spectral radius			

Mean log error of different aggregators on different tasks

Corso et al. 2020

Towards More Expressive GNNs









Higher-order WL tests

Maron et al. 2019 Morris et al. 2019

Positional & Structural encoding

Monti, Otness et B 2018 Sato 2020 Dwivedi et al. 2020 Bouritsas, Frasca et B 2020 ...many more Subgraph GNNs

Bevilacqua, Frasca et B, Maron 2021

Papp et al. 2021

Cotta et al. 2021

Zhao et al. 2021

Frasca et B, Maron 2022

Topological message passing

Bodnar, Frasca et B 2021

HIGHER-ORDER GNNs

WL test



Node colour refinement

k-WL test

• Colour adjacent *k*-tuples $\mathbf{v} \in V^k$ instead of nodes



Subgraph colour refinement

k-WL test



Subgraph colour refinement

- Colour adjacent *k*-tuples $\mathbf{v} \in V^k$ instead of nodes
- *k*-tuples are **adjacent** if they differ in one node; the *j*th neighbourhood of **v** is defined as

$$\mathcal{N}_{\mathbf{v},j} = \{ (v_1, \dots, v_{j-1}, w, v_{j+1}, \dots, v_k) : w \in V \}$$

k-WL test



Subgraph colour refinement

- Colour adjacent *k*-tuples $\mathbf{v} \in V^k$ instead of nodes
- *k*-tuples are **adjacent** if they differ in one node; the *j*th neighbourhood of **v** is defined as

$$\mathcal{N}_{\mathbf{v},j} = \left\{ \left(v_1, \dots, v_{j-1}, w, v_{j+1}, \dots, v_k \right) : w \in V \right\}$$

• For every tuple $\mathbf{v} \in V^k$, update the colour

$$\phi(c_{\mathbf{v}}, c_{\mathcal{N}_{\mathbf{v},1}}, \dots, c_{\mathcal{N}_{\mathbf{v},k}})$$

where $c_{\mathcal{N}_{\mathbf{v},j}} = \{ c_{\mathbf{w}} : \mathbf{w} \in \mathcal{N}_{\mathbf{v},j} \}$



Note: there are two slightly different versions of the high-dimensional WL test referred to as "k-WL" and "folklore k-WL" that differ in the update step.

What graphs can k-WL distinguish?



Cai, Fürer, Immerman 1992 (CFI graphs)

What graphs cannot 3-WL distinguish?



Strongly regular graphs cannot be distinguished by 3-WL

What graphs cannot k-WL distinguish?



Cai-Fürer-Immerman (CFI) graphs

Cai, Fürer, Immerman 1992 (CFI graphs)

Various k-WL-like GNNs

• *k*-GNNs mimics the *k*-WL test by working with subgraphs of size *k*:

$$\mathbf{x}_{\mathbf{v}} \leftarrow \varphi \left(\mathbf{W}_{1} \mathbf{x}_{\mathbf{v}} + \sum_{\mathbf{u} \in \mathcal{N}_{\mathbf{v}}} \mathbf{W}_{2} \mathbf{x}_{\mathbf{u}} \right)$$

where $\mathcal{N}_{\mathbf{v}} = \{\mathbf{u} \in V^k : |\mathbf{u} \cap \mathbf{v}| = k - 1\}$

- Complexity: $O(n^k)$
- More efficient variants
 - Sparse neighbourhoods
 - Sparse sets of *k*-tuples

Morris et al. 2019; Morris et al. 2020; Morris et al. 2022

Towards More Expressive GNNs









Higher-order WL tests

Maron et al. 2019 Morris et al. 2019

Positional & Structural encoding

Monti, Otness et B 2018 Sato 2020 Dwivedi et al. 2020 Bouritsas, Frasca et B 2020 ...many more Subgraph GNNs

Bevilacqua, Frasca et B, Maron 2021

Papp et al. 2021

Cotta et al. 2021

Zhao et al. 2021

Frasca et B, Maron 2022

Topological message passing

Bodnar, Frasca et B 2021

POSITIONAL ENCODING

What does WL test see?



Node encoding



"Colouring" nodes removes (some) ambiguity

How to "colour" nodes?

- Random
- Substructure count
- Laplacian/Adjacency eigenvectors
- Gradients of global encoding ("direction")
- Shortest path distance
- Diffusion kernel
- Random walk kernel
- ...many more

Random node features

- rMPNN: Attach a random feature to every node of the graph, then apply MPNN
- Output of rMPNN is a *random variable*
- Not permutation invariant! (only in expectation)

Probabilistic Universal Approximation: Let *f* be a permutation-invariant graph function. Then, for all ε , $\delta > 0$, there exists an rMPNN \hat{f} that (ε, δ) -approximates *f*, in the sense that

 $\mathbb{P}(\left|f(G) - \hat{f}(G)\right| \le \varepsilon) \ge 1 - \delta$

- Embedding dimension $\mathcal{O}(n^2 \delta)$
- Extensions to equivariant functions, weighted graphs, etc.

Abboud et al. 2021; Sato 2021

- Choose a *bank of substructures* containing graphs *H* of size k = O(1)
- Count the occurrence of each *H* in every node/edge of the input graph
 - *Subgraph* counts
 - *Induced subgraph* counts



- Choose a *bank of substructures* containing graphs *H* of size k = O(1)
- Count the occurrence of each *H* in every node/edge of the input graph
- Use the counts as additional node/edge features



- Choose a *bank of substructures* containing graphs *H* of size k = O(1)
- Count the occurrence of each *H* in every node/edge of the input graph
- Use the counts as additional node/edge features



substructure bank



- Choose a *bank of substructures* containing graphs *H* of size k = O(1)
- Count the occurrence of each *H* in every node/edge of the input graph
- Use the counts as additional node/edge features
- Complexity: *precomputation* (substructure counting, which is worse case is $O(n^k)$ but in practice much more optimistic) + *standard MPNN* (linear complexity $O(|E|) \sim O(n)$)

Theorem: GSN is strictly more expressive than WL if

- *H* is not a star graph, and counting is done using subgraph matching; or
- *H* is of size $k \ge 3$, and counting is using induced subgraph matching.

- Choose a *bank of substructures* containing graphs *H* of size k = O(1)
- Count the occurrence of each *H* in every node/edge of the input graph
- Use the counts as additional node/edge features
- Complexity: *precomputation* (substructure counting, which is worse case is $O(n^k)$ but in practice much more optimistic) + *standard MPNN* (linear complexity $O(|E|) \sim O(n)$)

Theorem: GSN is not less expressive than 3-WL.

Proof: by example

What graphs can GSN distinguish?



Strongly Regular (SR) graphs cannot be distinguished by 3-WL but can be distinguished by GSN with 4-clique count

What graphs can GSN distinguish?



What graphs can GSN distinguish hypothetically?



What graphs can GSN distinguish hypothetically?







Bouritsas, Frasca et B 2020 Slide adapted from Bevilacqua et al. (LoG 2022 tutorial)
Graph Substructure Network in practice



Molecule of caffeine

application-specific inductive bias

Bouritsas, Frasca et B 2020 Slide adapted from Bevilacqua et al. (LoG 2022 tutorial)

Towards More Expressive GNNs









Higher-order WL tests

Maron et al. 2019 Morris et al. 2019

Positional & Structural encoding

Monti, Otness et B 2018 Sato 2020 Dwivedi et al. 2020 Bouritsas, Frasca et B 2020 ...many more Subgraph GNNs

Bevilacqua, Frasca et B, Maron 2021

Papp et al. 2021

Cotta et al. 2021

Zhao et al. 2021

Frasca et B, Maron 2022

Topological message passing

Bodnar, Frasca et B 2021

SUBGRAPH GNNs

Subgraph GNNs



Subgraph GNNs



Graph perturbation allows to distinguish between structures otherwise indistinguishable by Weisfeiler-Lehman

Papp et al. 2021; Cotta et al. 2021; Zhao et al. 2021 Bevilacqua, Frasca et B, Maron 2022; Frasca et B, Maron 2022



A multiset of subgraphs obtained by edge deletion

Bevilacqua, Frasca et B, Maron 2022; Frasca et B, Maron 2022



A multiset of subgraphs obtained by node deletion

Bevilacqua, Frasca et B, Maron 2022; Frasca et B, Maron 2022

Graph Reconstruction Conjectures

Graph Reconstruction Conjecture: A graph *H* is said to be a *reconstruction* of *G* (denoted $H \sim G$) if they have the same multiset (*deck*) of node-removed subgraphs (*cards*). If *G* and *H* are two finite, undirected, simple graphs with at least three vertices and *H* is the reconstruction of *G*, then $H \simeq G$.



P. Kelly

S. Ulam

Collection of subgraphs determines the graph isomorphism class

Kelly 1942 (PhD thesis where the Conjecture appeared); Kelly 1957 (reconstruction from *k*-subgraphs); McKay 1997 (proof for small graphs); Nýdl 2001 (k-reconstructability for certain graph families)

Graph Reconstruction Conjectures

Graph Reconstruction Conjecture: A graph *H* is said to be a *reconstruction* of *G* (denoted $H \sim G$) if they have the same multiset (*deck*) of node-removed subgraphs (*cards*). If *G* and *H* are two finite, undirected, simple graphs with at least three vertices and *H* is the reconstruction of *G*, then $H \simeq G$.



- Proven for small graphs with $n \leq 11$ nodes
- Open question in general
- Generalisations for subgraphs of size n k

Kelly 1942 (PhD thesis where the Conjecture appeared); McKay 1997 (proof for small graphs); Kelly 1957 (reconstruction from *k*-subgraphs); Nýdl 2001 (k-reconstructability for certain graph families)

Equivariant Subgraph Aggregation Networks



Symmetry group of subgraph colection $G = S_n \times S_m$

Bevilacqua, Frasca et B, Maron 2022; Frasca et B, Maron 2022



Towards More Expressive GNNs









Higher-order WL tests

Maron et al. 2019 Morris et al. 2019

Positional & Structural encoding

Monti, Otness et B 2018 Sato 2020 Dwivedi et al. 2020 Bouritsas, Frasca et B 2020 ...many more Subgraph GNNs

Bevilacqua, Frasca et B, Maron 2021

Papp et al. 2021

Cotta et al. 2021

Zhao et al. 2021

Frasca et B, Maron 2022

Topological message passing

Bodnar, Frasca et B 2021

TOPOLOGICAL MESSAGE PASSING

Cell complexes



Topological Message Passing



- Lift the graph into a cell complex
- Hierarchical message passing
- Strictly more expressive than Weisfeiler-Lehman



Takeaways

- Message-Passing GNNs are upper-bounded in their expressive power by the Weisfeiler-Lehman graph isomorphism test
- Several ways to enhance expressive power:
 - Higher-order WL-tests
 - Structural and positional encoding
 - Subgraph aggregation methods
 - Graph lifting to simplicial/cellular complexes & topological message passing
- Generalisation power of GNNs is an open question
- Next: is Weisfeiler-Lehman formalism the end of story? (No!)

Key Concepts

- Graph isomorphism
- Weisfeiler-Lehman tests
- Positional and structural coding
- Reconstruction Conjectures

Main References

- M. Bronstein et al., <u>Geometric deep learning</u>, *arXiv:2104.13478*, 2021. Section 4.1 "Graphs and sets" and Section 5.3 "Graph neural networks"
- C. Morris et al., <u>Weisfeiler and Leman go Machine Learning: The Story so far</u>, *arXiv:*2112.099992, 2021. Introductory text on *k*-WL tests and various equivalent GNNs
- B. Bevilacqua, F. Frasca, H. Maron, <u>Exploring the practical and theoretical landscape of</u> <u>expressive GNNs</u>, LoG tutorial 2022. Expressive power and advanced GNN architectures

Additional References

- Z. Chen et al., <u>On the equivalence between graph isomorphism testing and function</u> <u>approximation with GNNs</u>, *NeurIPS* 2019. Universal approximation in GNNs
- K. Xu et al., <u>How powerful are graph neural networks</u>?, *ICLR* 2019. Proof of equivalence between WL-test and MPNN + GIN architecture
- M. Bronstein, L. Cotta, F. Frasca, H. Maron, <u>Using subgraphs for more expressive</u> <u>GNNs</u>, Towards Data Science blog post 2021. Diverse subgraph methods + Reconstruction conjectures
- C. Bodnar, F. Frasca, et al., <u>Weisfeiler and Lehman go cellular: CW networks</u>, *NeurIPS* 2021. Topological message passing

Main Ingredients of an MPNN



Weisfeiler-Lehman hierarchy



Weisfeiler, Lehman 1968 (2-WL); Babai, Mathon 1979 (k-WL) Cai, Furer, Immerman 1992 (CFI graphs)

Weisfeiler-Lehman hierarchy



GNN expressive power

Weisfeiler, Lehman 1968 (2-WL); Babai, Mathon 1979 (*k*-WL); Cai, Furer, Immerman 1992 (CFI graphs)

Structural encoding

Bouritsas, Frasca et B 2020

Weisfeiler-Lehman hierarchy



Graphs may be <mark>unfriendly</mark> for message passing resulting in "bottlenecks"



GNN expressive power

Weisfeiler, Lehman 1968 (2-WL); Babai, Mathon 1979 (*k*-WL); Cai, Fürer, Immerman 1992 (CFI graphs)

"Graph rewiring"

Alon, Yahav 2020 (bottlenecks); Hamilton et al. 2017 (neighbour sampling); Klicpera et al. 2019 (diffusion); Topping, Di Giovanni et B 2022 (Ricci flow); Deac et al. 2022 (expanders)

MORE STRUCTURE



MORE STRUCTURE



TRANSFORMERS



What graph should we use in GNNs?

"Too much flexibility"

- Assume the graph is *complete* (every node is connected to every node)
- Apply a Convolutional GNN:

$$\phi\left(\mathbf{x}_{i}, \sum_{j\in\mathcal{N}_{i}}a_{ij}\psi(\mathbf{x}_{j})\right)$$

- Assume the graph is *complete* (every node is connected to every node)
- Apply a Convolutional GNN:

$$\phi\left(\mathbf{x}_{i}, \sum_{j=1}^{n} a_{ij} \Psi(\mathbf{x}_{j})\right)$$

equal for every node!

- Assume the graph is *complete* (every node is connected to every node)
- Apply an Attentional GNN:

$$\phi\left(\mathbf{x}_{i}, \sum_{j=1}^{n} \boldsymbol{\alpha}(\mathbf{x}_{i}, \mathbf{x}_{j}) \boldsymbol{\psi}(\mathbf{x}_{j})\right)$$

- Attention weights = *learned graph adjacency* on the downstream task
- **Transformer** architecture, which has recently become predominant in NLP task, is a special type of GNN!

Vaswani et al. 2017; Joshi 2020

- Assume the graph is *complete* (every node is connected to every node)
- Apply an Attentional GNN:

$$\phi\left(\mathbf{x}_{i}, \sum_{j=1}^{n} \alpha(\mathbf{x}_{i}, \mathbf{x}_{j}, \mathbf{p}_{i}, \mathbf{p}_{j})\psi(\mathbf{x}_{j})\right)$$

- Attention weights = *learned graph adjacency* on the downstream task
- **Transformer** architecture, which has recently become predominant in NLP task, is a special type of GNN!
- Tasks in NLP often require *position-dependent* functions (not permutation-invariant), which is achieved through **positional encoding**

Vaswani et al. 2017; Joshi 2020



Vaswani et al. 2017; Joshi 2020

Positional encoding



Typical positional encoding used in Transformers on a 1D grid

Positional encoding

	Unique?	Complexity
Substructures of size <i>k</i>	Yes	$\mathcal{O}ig(n^k)^{\star}$
k Laplacian eigenvectors	No 2 ^k sign flips	$\mathcal{O}(n^2) \div \mathcal{O}(n^3)^{**}$
Random Walk kernels	Yes	$\mathcal{O}(n^2) \div \mathcal{O}(n^3) * *$

*In practice better for many substructures such as triangles **Depending on the sparsity of the Laplacian



Wang et B 2018
Differentiable Graph Module



Differentiable Graph Module (DGM) allowing to construct the graph from the data and use it for feature learning

Method	TADPOLE		UK Biobank	
	Transductive	Inductive	Transductive	Inductive
DGCNN	$84.59 {\pm} 4.33$	$\textbf{82.99}{\pm}\textbf{4.91}$	$58.35 {\pm} 0.91$	$51.84{\pm}8.16$
LDS	$87.06{\pm}3.67$	Ť	OOM	†
\mathbf{cDGM}	$92.91{\pm}2.50$	$91.85{\pm}2.62$	$61.32{\pm}1.51$	$55.91{\pm}3.49$
dDGM	$94.10{\pm}2.12$	$\textbf{92.17}{\pm}\textbf{3.65}$	$\textbf{63.22}{\pm}\textbf{1.12}$	$\textbf{57.34}{\pm}\textbf{5.32}$

Kazi, Cosmo et B 2020

Disease classification accuracy

Latent Simplicial Complex Learning



(a) Input, $\%_p = 100, h = 0.11$ (b) Epoch 40, $\%_p = 90, h = 0.4$ (c) Epoch 180, $\%_p = 20, h = 0.99$



Differentiable Cell Complex Module

Battiloro et B, Scardapane, Di Lorenzo 2024

OVERSQUASHING & GRAPH REWIRING

What graph should we use in GNNs?



What graph should we use in GNNs?



What graph should we use in GNNs?



- Preserve locality ("input graph inductive bias")
- Preserve sparsity (computational efficiency)
- Improve connectivity (to reduce over-squashing)

Barbero, Velingker, Saberi, B, Di Giovanni 2023

"Failure of Message Passing to propagate information on the graph"

Over-squashing





In small-world graphs metric ball volume grows exponentially with ball radius



Alon, Yahav 2020

Over-squashing

• Consider an MPNN of the form

$$\mathbf{x}_{i}^{(k+1)} = \sigma \left(\mathbf{W}_{1} \mathbf{x}_{i}^{(k)} + \sum_{j} a_{ij} \mathbf{W}_{2} \mathbf{x}_{j}^{(k)} \right)$$

- *L* = *depth* (number of layers)
- *p* =*width* (hidden dimension)
- Nonlinearity σ is c_{σ} -Lipschitz-continuous
- $w = \text{maximum element of weight matrices } \mathbf{W}_1, \mathbf{W}_2$







Over-squashing: small Jacobian $\left\| \partial \mathbf{x}_{i}^{(L)} / \partial \mathbf{x}_{j}^{(0)} \right\|$ indicates poor information propagation from input node

Preventing over-squashing

$$\left\| \frac{\partial \mathbf{x}_{i}^{(L)}}{\partial \mathbf{x}_{j}^{(0)}} \right\|_{1} \leq (\mathbf{C}_{\mathbf{\sigma}} \mathbf{w} \mathbf{p})^{L} (\mathbf{I} \blacksquare \mathbf{A})_{ij}^{L}$$

model topology

- Width *p* helps mitigate over-squashing (potentially at the risk of worse generalization)
- **Depth** *L* does not help
 - If *L*~diam(*G*), over-squashing occurs between distant nodes
 - If $L \gg 1$, we transition from over-squashing to vanishing gradients
- **Topology** of *G* has the largest effect on over-squashing, which occurs
 - Between nodes with high *effective resistance* $\text{Res}(i, j) \propto \tau(i, j) = commute time$
 - On graphs with small *Cheeger constant* h(G) = energy required to disconnect the graph into two communities)
 - Edges with strongly negative discrete *Ricci curvature* κ(*i*, *j*)

Di Giovanni et B 2023

Effective Resistance & Commute Time

- Commute time τ(i, j) = expected number of steps a random walk on a graph starting from node *i* will take to reach node *j* and come back
- **Effective resistance** Res(*i*, *j*) = voltage difference between nodes *i* and *j* if a unit current flows through the graph where every edge has unit resistance

 $\tau(i,j) = 2|E|\operatorname{Res}(i,j)$

Chandra et al. 1996

Cheeger constant

• The **Cheeger constant** of a graph G = (V, E) is defined as $h(G) = \min_{U \subset V} \frac{|\{(i, j) \in E \text{ s.t. } i \in U \text{ and } j \in V \setminus U\}|}{\min\{\operatorname{vol}(U), \operatorname{vol}(V)\}}$ where $\operatorname{vol}(U) = \sum_{j \in U} d_i$.



(Cheeger 1970)

Cheeger constant

• The **Cheeger constant** of a graph G = (V, E) is defined as

 $h(G) = \min_{U \subset V} \frac{|\{(i, j) \in E \text{ s.t. } i \in U \text{ and } j \in V \setminus U\}|}{\min\{\operatorname{vol}(U), \operatorname{vol}(V)\}}$

where $\operatorname{vol}(U) = \sum_{j \in U} d_i$.



(Cheeger 1970)

Cheeger constant

• The **Cheeger constant** of a graph G = (V, E) is defined as

$$h(G) = \min_{U \subset V} \frac{|\{(i,j) \in E \text{ s.t. } i \in U \text{ and } j \in V \setminus U\}|}{\min\{\operatorname{vol}(U), \operatorname{vol}(V)\}}$$

where $\operatorname{vol}(U) = \sum_{j \in U} d_i$.

- Energy required to disconnect *G* into separate communities ("bottleneck" if *h* is small)
- **Cheeger inequality:** $2h(G) \ge \lambda_1 > \frac{h^2(G)}{2}$, where λ_1 is the first nonzero eigenvalue ("spectral gap") of the normalised graph Laplacian



(Cheeger 1970)

Preventing over-squashing

$$\left\| \frac{\partial \mathbf{x}_{i}^{(L)}}{\partial \mathbf{x}_{j}^{(0)}} \right\|_{1} \leq (\mathbf{C}_{\sigma} \boldsymbol{w} \boldsymbol{p})^{L} (\mathbf{I} + \mathbf{A})_{ij}^{L}$$
model topology





Di Giovanni et B 2023

When oversquashing might occur?



Ricci Curvature on Manifolds



Spherical (>0)



Euclidean (=0)

"geodesic dispersion"



Hyperbolic (<0)

Ricci 1903

Graph analogues of Ricci curvature



Graph Ricci curvature

- Multiple definitions of curvature on graphs, two main
 - Forman (combinatorial)
 - Ollivier (optimal transport)
- **Balanced Forman Curvature** of edge *i*~*j* in simple unweighted graph is

•
$$\kappa(i,j) = 0$$
 if $\min\{d_i, d_j\} = 1$, otherwise
Triangles based at $i - j$
• $\kappa(i,j) = \frac{2}{d_i} + \frac{2}{d_j} + 2\frac{|\#_{\Delta}(i,j)|}{\max\{d_i,d_j\}} + \frac{|\#_{\Delta}(i,j)|}{\min\{d_i,d_j\}} + \frac{\gamma_{\max}^{-1}}{\max\{d_i,d_j\}} \left(\left| \#_{\Box}^i(i,j) \right| + \left| \#_{\Box}^j(i,j) \right| \right) - 2$
Pegree of i
Neighbours of i forming a 4-cycle based at $i - j$
Neighbours of i forming a 4-cycle based at $i - j$
(w/o diagonals)

Forman 2003; Ollivier 2007; Topping, di Giovanni, et B. 2021

Graph Ricci curvature

- Multiple definitions of curvature on graphs, two main
 - Forman (combinatorial)
 - Ollivier (optimal transport)
- Balanced Forman Curvature of edge *i~j* in simple unweighted graph behaves similarly to the continuous Ricci curvature



Forman 2003; Ollivier 2007; Topping, di Giovanni, et B. 2021

What contributes to over-squashing?

Theorem: (informal) *strong negatively-curved edges* contribute to over-squashing.



Topping, Di Giovanni et B 2021; *No contradiction to using expander graphs (which have negative curvature)

Curvature-based rewiring

Input: graph G = (V, E), temperature $\tau > 0$, (optional *C*)

• For edge $i \sim j$ with smallest $\kappa(i, j)$

- Calculate the improvement $\delta_{kl} = \kappa_{G'}(i, j) \kappa(i, j)$ from adding edge $k \sim l$ with $k \in B_1(i)$ and $l \in B_1(j)$
- Sample index k, l with probability Softmax(τδ_{kl}) and add edge k~l to E'

• (optional) Remove edge $i \sim j$ with largest Ric(i, j) > C

Output: new graph G' = (V, E')

Topping, di Giovanni, et B. 2021

Ricci flow

• Ricci flow: "diffusion of the Riemannian metric"



Evolution of a manifold under Ricci flow



G. Ricci-Curbastro R. Hamilton

Ricci 1903; Hamilton 1988;







G. Perelman

G. Ricci-Curbastro

R. Hamilton

Ricci 1903; Hamilton 1988; Perelman 2003

Curvature- vs Diffusion-based Rewiring



Topping, di Giovanni, et B. 2021; Klicpera et al. 2019 (DIGL)

What contributes to over-squashing?

Theorem: (informal) large *commute-time distances* contribute to over-squashing.

- **Spectral rewiring:**^{1,2} increase the Cheeger constant of the graph ("clusteredness"), which leads to lower commute time
- **Spatial rewiring:**³⁻⁵ inserting edges reduces the total effective resistance of the graph (=commute-time distance up to scale)
- Bounded-degree expanders⁶ have commute time O(|E|); over-squashing effect is independent on the size of the graph



Di Giovanni, Giusti, Barbero, Luise, Lio', B 2023

¹Arnaiz-Rodriguez et al. 2022 (DiffWire); ²Karhadkar et al. 2022; ³Abboud et al. 2022; ⁴Mialon et al. 2022; ⁵Abu-El-Haija et al. 2019; ³Deac et al. 2022 (EGP) & many more

Graph Rewiring Approaches



Connectivity Diffusion (DIGL) Gasteiger et al. 2019

- KNN graph on PPR embedding
- + Good for homophilic graphs
- Bad for heterophilic graphs
- Drastically different graph



Discrete Ricci flow (SDRF) Topping, Di Giovanni et B. 2022

- Remove negatively-curved edges
- + Good for both homophilic and heterophilic graphs
- + "Surgical" rewiring
- Curvature is computationally expensive

No relation to the task!



Expander Propagation (EGP) Deac et al. 2022

- Fixed expander graph
- Alternate MP on original and new graph
- + Optimal graph for MP
- No relation to input graph
- No permutation equivariance

Why it is important to consider the task?



Same graph+features, different task Whether the graph is good depends on the task!

LONG-RANGE INTERACTIONS & EXPRESSIVE POWER

Long-range interactions in graph tasks

- **Task** = a function $f(\mathbf{X})$ on the node features of a graph G
- The interaction between features in nodes *i* and *j* required for the task is given by

Mixing of
$$f$$
: $\min_{f}(i,j) = \max_{\mathbf{X}} \max_{1 \le \alpha, \beta \le d} \left| \frac{\partial^2 f(\mathbf{X})}{\partial x_i^{\alpha} \partial x_j^{\beta}} \right|$

- $f(\mathbf{X}) = \phi(\mathbf{x}_i) + \phi(\mathbf{x}_j)$ is fully separable, thus $\min_f(i, j) = 0$
- $f(\mathbf{X}) = \phi(\langle \mathbf{x}_i, \mathbf{x}_j \rangle)$ mixing depends on how non-linear ϕ is

Di Giovanni, Rusch, B, Deac, Lackenby, Mishra, Veličković 2023

Capacity bounds

$$\begin{split} \min_{\boldsymbol{f}}(i,j) &\leq \sum_{k=1}^{L-1} (\boldsymbol{c}_{\boldsymbol{\sigma}} \boldsymbol{w})^{2L-k-1} \left(\boldsymbol{w} \left(\boldsymbol{S}^{L-k} \right)^{\mathrm{T}} \mathrm{diag} \left(\boldsymbol{1}^{\mathrm{T}} \boldsymbol{S}^{k} \right) \boldsymbol{S}^{L-k} + \boldsymbol{C} \boldsymbol{Q}_{k} \right)_{ij} \\ \max_{\boldsymbol{k} \in \mathcal{K}} \operatorname{topology} \end{split}$$

What is the capacity of MPNN required for a given task? model + topology mixing

Di Giovanni, Rusch, B, Deac, Lackenby, Mishra, Veličković 2023

Capacity bounds

$$\operatorname{mix}_{f}(i,j) \leq \sum_{k=1}^{L-1} (C_{\sigma} W)^{2L-k-1} \left(W (\mathbf{S}^{L-k})^{\mathrm{T}} \operatorname{diag} (\mathbf{1}^{\mathrm{T}} \mathbf{S}^{k}) \mathbf{S}^{L-k} + C \mathbf{Q}_{k} \right)_{ij}$$

Bound on weights *w*



- d_{\min} =min node degree
- Fixed depth L = [d(i, j)/2]
- q =number of paths of length d(i, j) between i and j

Di Giovanni, Rusch, B, Deac, Lackenby, Mishra, Veličković 2023

Bound on depth *L*

$$L \ge \frac{\mathbf{a}(\mathbf{i}, \mathbf{j})}{4c_2} + \frac{|\mathbf{E}|}{\sqrt{\mathbf{d}_1 \mathbf{d}_1}} \left(\alpha \min_{f} (\mathbf{i}, \mathbf{j}) - \mathbf{\beta} \right)$$

- d_i =degree of node i
- α , β =model-related constants
- |E| =number of edges
- Bounded weights

Capacity bounds

$$\operatorname{mix}_{f}(i,j) \leq \sum_{k=1}^{L-1} (C_{\sigma} W)^{2L-k-1} \left(W (\mathbf{S}^{L-k})^{\mathrm{T}} \operatorname{diag} (\mathbf{1}^{\mathrm{T}} \mathbf{S}^{k}) \mathbf{S}^{L-k} + C \mathbf{Q}_{k} \right)_{ij}$$

Bound on weights *w*

$$w \ge \frac{d_{\min}}{c_2} \left(\frac{\min_f(i,j)}{q}\right)^{1/d(i,j)}$$

"weights need to be large enough to allow mixing"

Di Giovanni, Rusch, B, Deac, Lackenby, Mishra, Veličković 2023

Bound on depth *L*

$$L \ge \frac{\tau(i,j)}{4c_2} + \frac{|E|}{\sqrt{d_i d_j}} \left(\alpha \operatorname{mix}_f(i,j) - \beta \right)$$

- Depth must be \sim commute time $\tau(i, j)$
- Rewiring tries to improve τ
- τ can be as large as $O(n^3)$, which implies impossibility statements

New way of characterising expressive power

Expressive power (informal): MPNN with $L \le n$ layers **cannot learn tasks** that require high mixing among features at nodes with large commute time.





A. Lehman

B. Weisfeiler

Di Giovanni, Rusch, B, Deac, Lackenby, Mishra, Veličković 2023

EXOTIC MESSAGE PASSING














Gutteridge, Di Giovanni et B 2023





Experimental validation

Model	Peptides-func	Peptides-struct	PCQM-Contact	PascalVOC-SP
	$\mathrm{AP}\uparrow$	$\mathrm{MAE}\downarrow$	$\mathrm{MRR}\uparrow$	$F1\uparrow$
Classical MPNN	$0.6069 {\pm} 0.0035$	$0.3357{\pm}0.0006$	$0.3242{\pm}0.0008$	0.2860 ± 0.0085
Rewiring+MPNN (DIGL)	$0.6830{\pm}0.0026$	$0.2616{\pm}0.0018$	$0.1707{\pm}0.0021$	$0.2921{\pm}0.0038$
Multi-hop (MixHop-GCN)	$0.6843{\pm}0.0049$	$0.2614{\pm}0.0023$	$0.3250{\pm}0.0010$	$0.2506{\pm}0.0133$
Graph Transformer (GT)	$0.6326{\pm}0.0126$	$0.2529{\pm}0.0016$	$0.3174{\pm}0.0020$	$0.2694{\pm}0.0098$
GraphGPS (GT+MPNN)	$0.6535 {\pm} 0.0041$	$0.2500{\pm}0.0005$	$0.3337 {\pm} 0.0006$	$0.3748 {\pm} 0.0109$
ν DRew-MPNN	$0.7150 {\pm} 0.0044$	$0.2536{\pm}0.0015$	$0.3442{\pm}0.0006$	$0.3314 {\pm} 0.0024$

Evaluation on Long-Range Graph Benchmark

Gutteridge et B 2023

Cooperative Message Passing



Standard Message Passing each node Broadcasts & Listens

Finkelshtein, Huang, B, Ceylan 2023



Cooperative Message Passing each node individually decides

Cooperative Message Passing



Broadcast & Listen Listen

Broadcast

Isolate

Finkelshtein, Huang, B, Ceylan 2023

Cooperative Message Passing



Takeaways

- Input graph is not necessarily the best choice for message passing
- *Graph rewiring* separates input & computational graphs in attempt to avoid *over-squashing*
- Various rewiring techniques based on
 - Discrete Ricci curvature
 - Commute time / resistance distances
 - Spectral properties such as "connectedness"
- The graph structure impacts the expressive power of MPNN, which can be captured through the *mixing ability* of the GNN
- Advanced rewiring: *dynamic* (DReW) or *implicit* (cooperative GNNs)

Key Concepts

- Transformers and latent graph learning
- Graph rewiring
- Graph Ricci curvature
- Commute time
- Dynamic rewiring
- Co-operative GNNs

Main References

- C. Joshi, <u>Transformers are Graph Neural Networks</u>, 2021
- M. Bronstein, <u>Beyond Message Passing: a Physics-Inspired Paradigm for Graph Neural</u> <u>Networks</u>, *The Gradient* 2022. blog post on various exotic versions of message passing
- M. Bronstein, <u>Graph Neural Networks through the lens of Differential Geometry and</u> <u>Algebraic Topology</u>, *Towards Data Science* 2022.
- M. Bronstein, <u>Over-squashing</u>, <u>Bottlenecks</u>, and <u>Graph Ricci curvature</u>, *Towards Data Science* 2022.
- M. Bronstein, <u>Dynamically Rewired Delayed Message Passing GNNs</u>, Towards Data *Science* 2023.
- M. Bronstein, <u>Co-operative Graph Neural Networks</u>, *Towards Data Science* 2023.