

Logic and the Power of Recurrent Graph Neural Networks

Martin Grohe

 Graph neural networks (GNNs) are deep learning architectures for machine learning problems on graphs.

- Graph neural networks (GNNs) are deep learning architectures for machine learning problems on graphs.
- A GNN may be viewed as distributed message passing algorithm operating on the vertices of its input graph. The algorithm is controlled by parameters typically learned from data.



- Graph neural networks (GNNs) are deep learning architectures for machine learning problems on graphs.
- A GNN may be viewed as distributed message passing algorithm operating on the vertices of its input graph. The algorithm is controlled by parameters typically learned from data.



In this talk, we study the expressiveness of GNNs:

Which functions on graphs can be computed by a GNN?

GNN N with d layers maps graph G to sequence of signals

$$s^{(t)}: V(G) \to \mathbb{R}^{p_t}$$
 for $t = 0, \dots, d$



GNN N with d layers maps graph G to sequence of signals

$$s^{(t)}: V(G) \to \mathbb{R}^{p_t}$$
 for $t = 0, \dots, d$



GNN N with d layers maps graph G to sequence of signals

$$s^{(t)}: V(G) \to \mathbb{R}^{p_t}$$
 for $t = 0, \dots, d$



GNN N with d layers maps graph G to sequence of signals

$$\mathfrak{s}^{(t)}: V(G) \to \mathbb{R}^{p_t} \quad \text{for } t = 0, \dots, d$$



Initialisation: $\mathfrak{z}^{(0)}(v) \in \mathbb{R}^{p_0}$ encodes node labels or initial signal on the graph.

GNN N with d layers maps graph G to sequence of signals

$$\mathfrak{z}^{(t)}: V(G) \to \mathbb{R}^{p_t}$$
 for $t = 0, \ldots, d$



Initialisation: $\mathfrak{z}^{(0)}(v) \in \mathbb{R}^{p_0}$ encodes node labels or initial signal on the graph.

Aggregation:
$$a^{(t)}(v) \coloneqq \operatorname{agg}_t \left(\left\{ s^{(t-1)}(w) \mid w \in N_G(v) \right\} \right)$$

 \blacktriangleright agg_t aggregation function: coordinatewise sum, mean or max

GNN N with d layers maps graph G to sequence of signals

$$s^{(t)}: V(G) \to \mathbb{R}^{p_t}$$
 for $t = 0, \ldots, d$



Initialisation: $\mathfrak{z}^{(0)}(v) \in \mathbb{R}^{p_0}$ encodes node labels or initial signal on the graph.

Aggregation:
$$a^{(t)}(v) \coloneqq \operatorname{agg}_t\left(\left\{\!\!\left\{\mathfrak{s}^{(t-1)}(w) \mid w \in N_G(v)\right\}\!\!\right\}\!\right)$$

agg_t aggregation function: coordinatewise sum, mean or max

Combination: $s^{(t)}(v) \coloneqq \operatorname{comb}_t(s^{(t-1)}(v), a^{(t)}(v))$

• $\operatorname{comb}_t : \mathbb{R}^{2p_{t-1}} \to \mathbb{R}^{p_t}$ combination function computed by a feedforward neural network (a.k.a. multilayer perceptron)





► Nodes and edges are weighted.



- ► Nodes and edges are weighted.
- Node with weight b ∈ ℝ and incoming edges with weights w₁,..., w_k ∈ ℝ computes function

$$x_1,\ldots,x_k\mapsto\sigma\left(b+\sum_{i=1}^kw_ix_i\right)$$
,

where $\sigma : \mathbb{R} \to \mathbb{R}$ is the activation function.



- ► Nodes and edges are weighted.
- Node with weight b ∈ ℝ and incoming edges with weights w₁,..., w_k ∈ ℝ computes function

$$x_1,\ldots,x_k\mapsto\sigma\left(b+\sum_{i=1}^kw_ix_i
ight)$$
 ,

where $\sigma : \mathbb{R} \to \mathbb{R}$ is the activation function.





rectified linear unit (ReLU)

Typical activation functions are:

5



- ► Nodes and edges are weighted.
- Node with weight b ∈ ℝ and incoming edges with weights w₁,..., w_k ∈ ℝ computes function

$$x_1,\ldots,x_k\mapsto\sigma\left(b+\sum_{i=1}^kw_ix_i
ight)$$
 ,

where $\sigma : \mathbb{R} \to \mathbb{R}$ is the activation function.

In this talk, we always use

 $\operatorname{ReLu}(x) \coloneqq \max\{0, x\}$

as activation function.



Universal Approximation Theorem

Theorem (Cybenko 1989, Hornik 1991)

Let $f : K \to \mathbb{R}^n$ be a continuous function defined on a compact set $K \subseteq \mathbb{R}^m$. Then for every $\epsilon > 0$ there is a feedforward neural network N computing a function f_N such that

$$\sup_{\mathbf{x}\in K} \|f(\mathbf{x}) - f_N(\mathbf{x})\| < \epsilon.$$

Computation of GNNs (revisited)

GNN N with d layers maps graph G to sequence of signals

$$s^{(t)}: V(G) \to \mathbb{R}^{p_t}$$
 for $t = 0, \dots, d$

Initialisation: $\mathfrak{z}^{(0)}(v) \in \mathbb{R}^{p_0}$ encodes node labels or initial signal on the graph.

Aggregation:
$$a^{(t)}(v) \coloneqq \operatorname{agg}_t\left(\left\{\!\!\left\{\mathfrak{s}^{(t-1)}(w) \mid w \in N_G(v)\right\}\!\!\right\}\!\right)$$

agg_t aggregation function: coordinatewise sum, mean or max

Combination: $s^{(t)}(v) \coloneqq \operatorname{comb}_t(s^{(t-1)}(v), a^{(t)}(v))$

• $\operatorname{comb}_t : \mathbb{R}^{2p_{t-1}} \to \mathbb{R}^{p_t}$ combination function computed by a feedforward neural network (a.k.a. multilayer perceptron)

Global Aggregation

$$\mathscr{g}^{(t)} \coloneqq \operatorname{agg}_t^{\prime}(\{\!\!\{\mathfrak{s}^{(t-1)}(w) \mid w \in V(G)\}\!\!\}),$$

where the aggregation function agg'_t is coordinatewise sum, mean or max.

Global Aggregation

$$\mathscr{Q}^{(t)} \coloneqq \operatorname{agg}_t^{\prime}(\{\!\!\{\mathfrak{s}^{(t-1)}(w) \mid w \in V(G)\}\!\!\}),$$

where the aggregation function agg'_t is coordinatewise sum, mean or max. Then combination becomes

$$\mathfrak{z}^{(t)}(v)\coloneqq \mathsf{comb}_t\bigl(\mathfrak{z}^{(t-1)}(v), \, a^{(t)}(v), \, \mathfrak{g}^{(t)}\bigr)$$

with $\operatorname{comb}_t : \mathbb{R}^{3p_{t-1}} \to \mathbb{R}^{p_t}$ computed by feedforward neural network.

Global Aggregation

$$\mathscr{Q}^{(t)} \coloneqq \operatorname{agg}_t^{\prime}(\{\!\!\{\mathfrak{s}^{(t-1)}(w) \mid w \in V(G)\}\!\!\}),$$

where the aggregation function agg'_t is coordinatewise sum, mean or max. Then combination becomes

$$\mathfrak{s}^{(t)}(v)\coloneqq \mathsf{comb}_t\bigl(\mathfrak{s}^{(t-1)}(v), \, a^{(t)}(v), \, \mathcal{g}^{(t)}\bigr)$$

with $\operatorname{comb}_t : \mathbb{R}^{3p_{t-1}} \to \mathbb{R}^{p_t}$ computed by feedforward neural network.

Random Initialisation

 $s^{(0)}(v) = (s_1, \ldots, s_{p_0}, \frac{s_{p_0+1}}{s_{p_0+1}}) \in \mathbb{R}^{p_0+1}$ where (s_1, \ldots, s_{p_0}) endodes node labels or initial signal and s_{p_0+1} is chosen uniformly at random from [0, 1].

Global Aggregation

$$\mathscr{g}^{(t)} \coloneqq \operatorname{agg}_t^{\prime}(\{\!\!\{\mathfrak{z}^{(t-1)}(w) \mid w \in V(G)\}\!\!\}),$$

where the aggregation function agg'_t is coordinatewise sum, mean or max. Then combination becomes

$$\mathfrak{s}^{(t)}(v)\coloneqq \mathsf{comb}_t\bigl(\mathfrak{s}^{(t-1)}(v), \, a^{(t)}(v), \, \mathscr{g}^{(t)}\bigr)$$

with $\operatorname{comb}_t : \mathbb{R}^{3p_{t-1}} \to \mathbb{R}^{p_t}$ computed by feedforward neural network.

Random Initialisation

 $s^{(0)}(v) = (s_1, \ldots, s_{p_0}, \frac{s_{p_0+1}}{s_{p_0+1}}) \in \mathbb{R}^{p_0+1}$ where (s_1, \ldots, s_{p_0}) endodes node labels or initial signal and $\frac{s_{p_0+1}}{s_{p_0+1}}$ is chosen uniformly at random from [0, 1].

In this talk, we always consider GNNs with global aggregation. We sometimes consider random initialisation, but mention it explicitly when we do.

Node-Level Functions

GNNs compute a function $\not\in$ that maps each graph G (possibly with node labels or an initial signal) to a signal $\not\in (G, \cdot) : V(G) \to \mathbb{R}^q$ defined by

$$f(G, v) = s^{(d)}(v).$$

Node-Level Functions

GNNs compute a function $\not \in$ that maps each graph G (possibly with node labels or an initial signal) to a signal $\not \in (G, \cdot) : V(G) \to \mathbb{R}^q$ defined by

$$f(G, v) = s^{(d)}(v).$$

Graph-Level Functions

To compute a function $\mathscr{F}: \mathscr{G} \to \mathbb{R}^q$ mapping graphs to vectors of real numbers,

Node-Level Functions

GNNs compute a function $\not \in$ that maps each graph G (possibly with node labels or an initial signal) to a signal $\not \in (G, \cdot) : V(G) \to \mathbb{R}^q$ defined by

$$f(G, v) = s^{(d)}(v).$$

class of all graphs

Graph-Level Functions

To compute a function $\mathscr{F}: \mathscr{G} \xrightarrow{\sim} \mathbb{R}^q$ mapping graphs to vectors of real numbers,

Node-Level Functions

GNNs compute a function $\not \in$ that maps each graph G (possibly with node labels or an initial signal) to a signal $\not \in (G, \cdot) : V(G) \to \mathbb{R}^q$ defined by

$$f(G, v) = s^{(d)}(v).$$

Graph-Level Functions

To compute a function $\mathscr{F}: \mathscr{G} \to \mathbb{R}^q$ mapping graphs to vectors of real numbers, we aggregate the values and then apply a readout function:

$$\mathscr{F}(G) = \operatorname{ro}\left(\operatorname{agg}\left(\left\{\!\!\left\{\mathfrak{s}^{(d)}(v) \mid v \in V(G)\right\}\!\!\right\}\right).$$

agg aggregation function: sum, mean or max

▶ ro : $\mathbb{R}^{p} \to \mathbb{R}^{q}$ readout function computed by a feedforward neural network

In this talk, we focus on graph-level functions.

In this talk, we focus on graph-level functions.

Invariance

Let $\mathscr{F}: \mathscr{G} \to \mathbb{R}^q$ be computed by a GNN. Then for all isomorphic graphs G, H,

 $\mathcal{F}(G) = \mathcal{F}(H).$

In this talk, we focus on graph-level functions.

Invariance

Let $\mathscr{F}: \mathscr{G} \to \mathbb{R}^q$ be computed by a GNN. Then for all isomorphic graphs G, H,

$$\mathscr{F}(G) = \mathscr{F}(H).$$

Graph Properties

When comparing GNNs with logic and circuit complexity, we often look at (isomorphism invariant) Boolean functions $\mathscr{P} : \mathscr{G} \to \{0, 1\}$, that is, properties of graphs.

In this talk, we focus on graph-level functions.

Invariance

Let $\mathscr{F}: \mathscr{G} \to \mathbb{R}^q$ be computed by a GNN. Then for all isomorphic graphs G, H,

$$\mathcal{F}(G) = \mathcal{F}(H).$$

Graph Properties

When comparing GNNs with logic and circuit complexity, we often look at (isomorphism invariant) Boolean functions $\mathscr{P} : \mathscr{G} \to \{0, 1\}$, that is, properties of graphs.

Remark

GNNs with random initialisation compute an isomorphism invariant random variable.

In this talk, we focus on graph-level functions.

Invariance

Let $\mathscr{F}: \mathscr{G} \to \mathbb{R}^q$ be computed by a GNN. Then for all isomorphic graphs G, H,

$$\mathcal{F}(G) = \mathcal{F}(H).$$

Graph Properties

When comparing GNNs with logic and circuit complexity, we often look at (isomorphism invariant) Boolean functions $\mathscr{P} : \mathscr{G} \to \{0, 1\}$, that is, properties of graphs.

Remark

GNNs with random initialisation compute an isomorphism invariant random variable. To compute a function, we can run the GNN several times with independent random initialisations and then take the average value, or for Boolean functions, the majority.

The Logic of GNNs



 $\left\{ \begin{array}{c} \mathsf{C}^2\\ \mathsf{FO}^2+\mathsf{C} \end{array} \right\}$ 2-variable fragments of first-order logic with counting quantifiers $\exists^{\geq n} x \varphi$



 $\begin{bmatrix} C^2 \\ FO^2 + C \end{bmatrix}$ 2-variable fragments of first-order logic with counting quantifiers $\exists^{\geq n} x \varphi$

Difference:

▶ in C², the number *n* in a counting quantifier $\exists^{\geq n} x$ is a constant;



 $\begin{bmatrix} C^2 \\ FO^2 + C \end{bmatrix}$ 2-variable fragments of first-order logic with counting quantifiers $\exists^{\geq n} x \varphi$

Difference:

- ▶ in C², the number *n* in a counting quantifier $\exists^{\geq n} x$ is a constant;
- \blacktriangleright in FO²+C it is a variable or term, and we allow arithmetic on such number terms.



2-variable fragments of first-order logic with counting quantifiers $\exists^{\geq n} x \varphi$

Difference:

- ▶ in C², the number *n* in a counting quantifier $\exists^{\geq n} x$ is a constant;
- in FO^2+C it is a variable or term, and we allow arithmetic on such number terms.

Example

• A C²-formula expressing that vertex x has degree 8:

$$\exists^{\geq 8} y E(x, y) \land \neg \exists^{\geq 9} y E(x, y).$$
2-Variable Counting Logic



 $\begin{bmatrix} C^2 \\ EO^2 + C \end{bmatrix}$ 2-variable fragments of first-order logic with counting quantifiers $\exists^{\geq n} x \varphi$

Difference:

- ▶ in C², the number *n* in a counting quantifier $\exists^{\geq n} x$ is a constant;
- \blacktriangleright in FO²+C it is a variable or term, and we allow arithmetic on such number terms.

Example

 \triangleright A C²-formula expressing that vertex x has degree 8:

$$\exists^{\geq 8} y E(x, y) \land \neg \exists^{\geq 9} y E(x, y).$$

An FO²+C-formula expressing that vertex x has even degree:

$$\exists n \Big(\exists^{\geq 2n} y \, E(x, y) \land \neg \exists^{\geq 2n+1} y \, E(x, y) \Big).$$





- 1. G, H are C^2 -equivalent.
- 2. *G*, *H* are FO^2+C -equivalent.



- 1. G, H are C^2 -equivalent.
- 2. G, H are FO²+C-equivalent.
- 3. The Colour Refinement algorithm does not distinguish G, H.



- 1. G, H are C^2 -equivalent.
- 2. G, H are FO^2+C -equivalent.
- 3. The Colour Refinement algorithm does not distinguish G, H.
- 4. *G* and *H* are fractionally isomorphic, that is, there is a doubly stochastic matrix X such that $A_G X = XA_H$.



- 1. G, H are C^2 -equivalent.
- 2. G, H are FO²+C-equivalent.
- 3. The Colour Refinement algorithm does not distinguish G, H.
- 4. *G* and *H* are fractionally isomorphic, that is, there is a doubly stochastic matrix X such that $A_G X = XA_H$.
- 5. For all trees T, the number of homomorphisms from T to G equals the number of homomorphisms from T to H.

The Colour Refinement algorithm iteratively computes a colouring of the vertices of graph G.

The Colour Refinement algorithm iteratively computes a colouring of the vertices of graph G.

Initialisation All vertices get the same colour.

The Colour Refinement algorithm iteratively computes a colouring of the vertices of graph G.

Initialisation All vertices get the same colour.

Refinement Step Two nodes v, w get different colours if there is some colour c such that v and w have different numbers of neighbours of colour c.

The Colour Refinement algorithm iteratively computes a colouring of the vertices of graph G.

Initialisation All vertices get the same colour.

Refinement Step Two nodes v, w get different colours if there is some colour c such that v and w have different numbers of neighbours of colour c. Refinement is repeated until colouring stays stable.

The Colour Refinement algorithm iteratively computes a colouring of the vertices of graph G.

Initialisation All vertices get the same colour.

Refinement Step Two nodes v, w get different colours if there is some colour c such that v and w have different numbers of neighbours of colour c. Refinement is repeated until colouring stays stable.

Remark

The algorithm is also refered to as naive vertex classification and as 1-dimensional Weisfeiler-Leman algorithm. One needs to be careful though, because there are two slightly different versions of the algorithm, and this difference does play a role in our work (though not in this talk).

























Colour Refinement distinguishes two graphs G, H if their colour histograms differ, that is, some colour appears a different number of times in G and H.

Colour Refinement distinguishes two graphs G, H if their colour histograms differ, that is, some colour appears a different number of times in G and H. Thus Colour Refinement can be used as an incomplete isomorphism test.

Colour Refinement distinguishes two graphs G, H if their colour histograms differ, that is, some colour appears a different number of times in G and H.

Thus Colour Refinement can be used as an incomplete isomorphism test.

works on almost all graphs (Babai, Erdös, Selkow 1980)

Colour Refinement distinguishes two graphs G, H if their colour histograms differ, that is, some colour appears a different number of times in G and H.

Thus Colour Refinement can be used as an incomplete isomorphism test.

works on almost all graphs (Babai, Erdös, Selkow 1980)

fails on some very simple graphs:



The Distinguishing Power of GNNs

Theorem (Morris, Ritzert, Fey, Hamilton, Lenssen, Rattan, G. 2019, Xu, Hu, Leskovec, Jegelka 2019)

For all graphs G, H, the following are equivalent:

- 1. G and H are C^2 -equivalent.
- 2. G and H are not distinguishable by a GNN, that is, for all GNNs N we have $\mathscr{F}_N(G) = \mathscr{F}_N(H)$;

The Distinguishing Power of GNNs

Theorem (Morris, Ritzert, Fey, Hamilton, Lenssen, Rattan, G. 2019, Xu, Hu, Leskovec, Jegelka 2019)

For all graphs G, H, the following are equivalent:

- 1. G and H are C^2 -equivalent.
- 2. *G* and *H* are not distinguishable by a GNN, that is, for all GNNs N we have $\mathscr{F}_N(G) = \mathscr{F}_N(H)$;

Corollary

Every function $\mathscr{F} : \mathscr{G} \to \mathbb{R}^q$ computable by a GNN is \mathbb{C}^2 -invariant, that is, if G, H are \mathbb{C}^2 -equivalent then $\mathscr{F}(G) = \mathscr{F}(H)$.

The Distinguishing Power of GNNs

Theorem (Morris, Ritzert, Fey, Hamilton, Lenssen, Rattan, G. 2019, Xu, Hu, Leskovec, Jegelka 2019)

For all graphs G, H, the following are equivalent:

- 1. G and H are C^2 -equivalent.
- 2. *G* and *H* are not distinguishable by a GNN, that is, for all GNNs N we have $\mathscr{F}_N(G) = \mathscr{F}_N(H)$;

Corollary

Every function $\mathscr{F} : \mathscr{G} \to \mathbb{R}^q$ computable by a GNN is \mathbb{C}^2 -invariant, that is, if G, H are \mathbb{C}^2 -equivalent then $\mathscr{F}(G) = \mathscr{F}(H)$.

Remark

(Morris et al. 2019) ensure that graphs G, H that are not C²-equivalent can be distinguished by a GNN N of size polynomial in |G|, |H|.

Theorem (Barceló, Kostylev, Monet, Pérez, Reutter, Silva 2019) Let \mathscr{P} be a graph property expressible in the logic C². Then there is a GNN (with rational weights) computing \mathscr{P} .

Theorem (Barceló, Kostylev, Monet, Pérez, Reutter, Silva 2019) Let \mathscr{P} be a graph property expressible in the logic C². Then there is a GNN (with rational weights) computing \mathscr{P} .

Theorem (G. 2023)

Let \mathscr{P} be a graph property computable by a GNN with rational weights. Then \mathscr{P} is expressible in in FO²+C.

Theorem (Barceló, Kostylev, Monet, Pérez, Reutter, Silva 2019) Let \mathscr{P} be a graph property expressible in the logic C². Then there is a GNN (with rational weights) computing \mathscr{P} .

Theorem (G. 2023)

Let \mathscr{P} be a graph property computable by a GNN with rational weights. Then \mathscr{P} is expressible in in FO²+C.

Corollary

$$C^2 \subset GNN \subset FO^2+C.$$

Theorem (Barceló, Kostylev, Monet, Pérez, Reutter, Silva 2019) Let \mathscr{P} be a graph property expressible in the logic C². Then there is a GNN (with rational weights) computing \mathscr{P} .

Theorem (G. 2023)

Let \mathscr{P} be a graph property computable by a GNN with rational weights. Then \mathscr{P} is expressible in in FO²+C.

Corollary

$$C^2 \subset GNN \subset FO^2+C.$$

It can be shown that both inclusions are strict.

Theorem (G. 2023)

For all graph properties \mathcal{P} , the following are equivalent.

1. *P* is computable by a polynomial-size bounded-depth family of GNNs with random initialisation and with arbitrary real weights (that may use activations like sigmoid or tanh besides ReLU).

Theorem (G. 2023)

For all graph properties \mathcal{P} , the following are equivalent.

- 1. *P* is computable by a polynomial-size bounded-depth family of GNNs with random initialisation and with arbitrary real weights (that may use activations like sigmoid or tanh besides ReLU).
- 2. \mathscr{P} is computable by a polynomial-size bounded-depth family of GNNs with random initialisation and with rational weights, using only sum aggregation.

Theorem (G. 2023)

For all graph properties \mathcal{P} , the following are equivalent.

- 1. *P* is computable by a polynomial-size bounded-depth family of GNNs with random initialisation and with arbitrary real weights (that may use activations like sigmoid or tanh besides ReLU).
- 2. \mathscr{P} is computable by a polynomial-size bounded-depth family of GNNs with random initialisation and with rational weights, using only sum aggregation.
- 3. \mathcal{P} is expressible in FO²+C with built-in relations.

Theorem (G. 2023)

For all graph properties \mathcal{P} , the following are equivalent.

- 1. *P* is computable by a polynomial-size bounded-depth family of GNNs with random initialisation and with arbitrary real weights (that may use activations like sigmoid or tanh besides ReLU).
- 2. \mathscr{P} is computable by a polynomial-size bounded-depth family of GNNs with random initialisation and with rational weights, using only sum aggregation.
- 3. \mathcal{P} is expressible in FO²+C with built-in relations.
- 4. \mathscr{P} is in the complexity class TC^0 .

So far, GNNs have a fixed number d of layers, and each layer t has its own functions $agg_t, agg'_t, comb_t$.

So far, GNNs have a fixed number d of layers, and each layer t has its own functions agg_t , agg'_t , $comb_t$.

Recurrent GNN Computation

A recurrent GNN has a single layer and applies it repeatedly. That is, we have aggregation functions agg, agg' and a combination function comb and for $t \ge 1$ let:

$$s^{(t)}(v) = \text{comb}\Big(s^{(t-1)}(v), \text{agg}\big(\{\!\!\{s^{(t-1)}(w) \mid w \in N_G(v)\}\!\!\}\big), \text{agg}'\big(\{\!\!\{s^{(t-1)}(w) \mid w \in V(G)\}\!\!\}\big)\Big)$$

So far, GNNs have a fixed number d of layers, and each layer t has its own functions agg_t , agg'_t , $comb_t$.

Recurrent GNN Computation

A recurrent GNN has a single layer and applies it repeatedly. That is, we have aggregation functions agg, agg' and a combination function comb and for $t \ge 1$ let:

$$s^{(t)}(v) = \operatorname{comb}\left(s^{(t-1)}(v), \operatorname{agg}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in N_G(v)\}\!\!\}\right), \operatorname{agg'}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in V(G)\}\!\!\}\right)\right)$$

Termination and Readout

Each node has flag (say, the last coordinate in s^(t)(v)) indicating whether the local computation at the node is complete.
Recurrent GNNs

So far, GNNs have a fixed number d of layers, and each layer t has its own functions agg_t , agg'_t , $comb_t$.

Recurrent GNN Computation

A recurrent GNN has a single layer and applies it repeatedly. That is, we have aggregation functions agg, agg' and a combination function comb and for $t \ge 1$ let:

$$s^{(t)}(v) = \operatorname{comb}\left(s^{(t-1)}(v), \operatorname{agg}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in N_G(v)\}\!\!\}\right), \operatorname{agg'}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in V(G)\}\!\!\}\right)\right)$$

Termination and Readout

- Each node has flag (say, the last coordinate in s^(t)(v)) indicating whether the local computation at the node is complete.
- ► The global computation halts once every node has set its flag.

Recurrent GNNs

So far, GNNs have a fixed number d of layers, and each layer t has its own functions agg_t , agg'_t , $comb_t$.

Recurrent GNN Computation

A recurrent GNN has a single layer and applies it repeatedly. That is, we have aggregation functions agg, agg' and a combination function comb and for $t \ge 1$ let:

$$s^{(t)}(v) = \operatorname{comb}\left(s^{(t-1)}(v), \operatorname{agg}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in N_G(v)\}\!\!\}\right), \operatorname{agg'}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in V(G)\}\!\!\}\right)\right)$$

Termination and Readout

- Each node has flag (say, the last coordinate in s^(t)(v)) indicating whether the local computation at the node is complete.
- ► The global computation halts once every node has set its flag.
- ► At that point, the signal is aggregated and a readout functions is applied.

Recurrent GNNs

So far, GNNs have a fixed number d of layers, and each layer t has its own functions agg_t , agg'_t , $comb_t$.

Recurrent GNN Computation

A recurrent GNN has a single layer and applies it repeatedly. That is, we have aggregation functions agg, agg' and a combination function comb and for $t \ge 1$ let:

$$s^{(t)}(v) = \operatorname{comb}\left(s^{(t-1)}(v), \operatorname{agg}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in N_G(v)\}\!\!\}\right), \operatorname{agg'}\left(\{\!\!\{s^{(t-1)}(w) \mid w \in V(G)\}\!\!\}\right)\right)$$

Termination and Readout

- Each node has flag (say, the last coordinate in s^(t)(v)) indicating whether the local computation at the node is complete.
- ► The global computation halts once every node has set its flag.
- At that point, the signal is aggregated and a readout functions is applied.
 We do not require convergence of the signals.

Universality of Recurrent GNNs

Theorem (Rosenbluth and G. 2025+)

Let $\mathscr{F} : \mathscr{G} \to \mathbb{Q}^n$ be computable. Then the following are equivalent.

(i) \mathcal{F} is C²-invariant.

(ii) \mathcal{F} is computable by a recurrent GNN.



Eran Rosenbluth

Universality of Recurrent GNNs

Theorem (Rosenbluth and G. 2025+)

Let $\mathscr{F} : \mathscr{G} \to \mathbb{Q}^n$ be computable. Then the following are equivalent.

(i) \mathcal{F} is C²-invariant.

(ii) \mathcal{F} is computable by a recurrent GNN.



Eran Rosenbluth

Furthermore,

- we may choose the GNN in (ii) to have rational weights and to only use SUM aggregation;
- ▶ if ℱ is computable in polynomial time then the GNN stops after polynomially many iterations and only uses internal states of polynomial bitlength.

Corollary

Every computable function $\mathscr{F} : \mathscr{G} \to \mathbb{Q}^n$ is computable by a recurrent GNN with random initialisation only using rational weight and SUM aggregation.

Corollary

Every computable function $\mathscr{F} : \mathscr{G} \to \mathbb{Q}^n$ is computable by a recurrent GNN with random initialisation only using rational weight and SUM aggregation.

Furthermore, if \mathcal{F} is computable in polynomial time then the GNN runs in polynomial time and only uses internal states of polynomial bitlength.

Corollary

Every computable function $\mathscr{F} : \mathscr{G} \to \mathbb{Q}^n$ is computable by a recurrent GNN with random initialisation only using rational weight and SUM aggregation.

Furthermore, if \mathcal{F} is computable in polynomial time then the GNN runs in polynomial time and only uses internal states of polynomial bitlength.

Remarks

Global aggregation simplifies the theorem and its proof. We also have a version without global aggregation, but it requires to pass the size of the graph as part of the initial signal, and it only applies to connected graphs.

Corollary

Every computable function $\mathscr{F} : \mathscr{G} \to \mathbb{Q}^n$ is computable by a recurrent GNN with random initialisation only using rational weight and SUM aggregation.

Furthermore, if \mathcal{F} is computable in polynomial time then the GNN runs in polynomial time and only uses internal states of polynomial bitlength.

Remarks

- Global aggregation simplifies the theorem and its proof. We also have a version without global aggregation, but it requires to pass the size of the graph as part of the initial signal, and it only applies to connected graphs.
- There are also a versions of the results for node-level functions (actually, we prove these first and derive the graph-level results as a consequence).

We want to compute $\mathcal{F}(G)$ for some graph G by a recurrent GNN.

We want to compute $\mathcal{F}(G)$ for some graph G by a recurrent GNN.

Phase 1: Computing the Colours

We simulate colour refinement. After Phase 1, the signal at every node v holds a representation D_v of v's its colour.

We want to compute $\mathcal{F}(G)$ for some graph G by a recurrent GNN.

Phase 1: Computing the Colours

We simulate colour refinement. After Phase 1, the signal at every node v holds a representation D_v of v's its colour.

Phase 2: Inversion

Locally at every node v, from the colour D_v we compute a graph G_v that is C²-equivalent to the input graph G.

We want to compute $\mathcal{F}(G)$ for some graph G by a recurrent GNN.

Phase 1: Computing the Colours

We simulate colour refinement. After Phase 1, the signal at every node v holds a representation D_v of v's its colour.

Phase 2: Inversion

Locally at every node v, from the colour D_v we compute a graph G_v that is C²-equivalent to the input graph G.

Phase 3: Computation of \mathcal{F}

Locally at every node v, we compute $\mathscr{F}(G_v)$. Since \mathscr{F} is C²-invariant, we have $\mathscr{F}(G_v) = \mathscr{F}(G)$.

We want to compute $\mathcal{F}(G)$ for some graph G by a recurrent GNN.

Phase 1: Computing the Colours

We simulate colour refinement. After Phase 1, the signal at every node v holds a representation D_v of v's its colour.

Phase 2: Inversion

Locally at every node v, from the colour D_v we compute a graph G_v that is C²-equivalent to the input graph G.

Phase 3: Computation of \mathcal{F}

Locally at every node v, we compute $\mathscr{F}(G_v)$. Since \mathscr{F} is C²-invariant, we have $\mathscr{F}(G_v) = \mathscr{F}(G)$.

Thus after Phase 3, every node holds the correct function value, and readout becomes trivial.

Phase 3: Computation of \mathcal{F}

Locally at every node v, we need to compute $\mathscr{F}(G_v)$ from the graph G_v that is encoded in the state of node v after Phase 2.

Phase 3: Computation of \mathcal{F}

Locally at every node v, we need to compute $\mathscr{F}(G_v)$ from the graph G_v that is encoded in the state of node v after Phase 2.

Local computation of a GNN at a node that ignores the message passing is essentially the computation of a recurrent feedforward neural network.

Phase 3: Computation of \mathcal{F}

Locally at every node v, we need to compute $\mathscr{F}(G_v)$ from the graph G_v that is encoded in the state of node v after Phase 2.

Local computation of a GNN at a node that ignores the message passing is essentially the computation of a recurrent feedforward neural network.

Thus we can use the following result

Theorem (Siegelmann and Sontag 1992)

Every computable function $f : \mathbb{Q}^m \to \mathbb{Q}^n$ is computable by a recurrent feedforward neural network.

What are the "colours"?

What are the "colours"?



Graph G

What are the "colours"?



Graph *G* coloured

What are the "colours"?



What are the "colours"?



What are the "colours"?



To simulate colour refinement by a recurrent GNN, we encode the compact dag representation of colours by rational numbers and in several rounds by decreasing c for each node v count the number of neighbours of v of colour c.

Phase 2: Inversion

The following lemma is an easy adaptation of a result on "inverting C²-invariants" due to Martin Otto.

Phase 2: Inversion

The following lemma is an easy adaptation of a result on "inverting C²-invariants" due to Martin Otto.

Lemma (Otto 1997)

There is a polynomial time algorithm that, given the dag representation D_v of the colour of a node v in a graph G, computes a graph G_v that is C²-equivalent to G.

Our results on recurrent GNNs necessarily require rationals of unbounded precision in the (internal) signals computed by the GNN, even if we are only interested in computing a Boolean function.

- Our results on recurrent GNNs necessarily require rationals of unbounded precision in the (internal) signals computed by the GNN, even if we are only interested in computing a Boolean function.
 - It would be interesting, to understand the precise bitlength required. For example, which functions can we compute with logarithmic bitlength (in the size of the input graph)?

Our results on recurrent GNNs necessarily require rationals of unbounded precision in the (internal) signals computed by the GNN, even if we are only interested in computing a Boolean function.

It would be interesting, to understand the precise bitlength required. For example, which functions can we compute with logarithmic bitlength (in the size of the input graph)?

Is there a version of our main result on recurrent GNNs for higher-order recurrent GNNs and functions invariant under the corresponding higher-order Weisfeiler-Leman algorithm?

Our results on recurrent GNNs necessarily require rationals of unbounded precision in the (internal) signals computed by the GNN, even if we are only interested in computing a Boolean function.

It would be interesting, to understand the precise bitlength required. For example, which functions can we compute with logarithmic bitlength (in the size of the input graph)?

- Is there a version of our main result on recurrent GNNs for higher-order recurrent GNNs and functions invariant under the corresponding higher-order Weisfeiler-Leman algorithm?
- We have a good understanding of the expressiveness of GNNs. But expressiveness results only tell half the story, because they ignore learning. However, many of the results presented here have good experimental support.

Our results on recurrent GNNs necessarily require rationals of unbounded precision in the (internal) signals computed by the GNN, even if we are only interested in computing a Boolean function.

It would be interesting, to understand the precise bitlength required. For example, which functions can we compute with logarithmic bitlength (in the size of the input graph)?

- Is there a version of our main result on recurrent GNNs for higher-order recurrent GNNs and functions invariant under the corresponding higher-order Weisfeiler-Leman algorithm?
- We have a good understanding of the expressiveness of GNNs. But expressiveness results only tell half the story, because they ignore learning. However, many of the results presented here have good experimental support.

If you are looking for a postdoc position, please contact me.

A Few References

```
Grohe, M. (2021). The Logic of Graph Neural Networks.
In: Proc. LICS 2021.
arXiv:2104.14624
```

Grohe, M. (2024). The Descriptive Complexity of Graph Neural Networks. In: TheoretiCS 3(25), 2024. arXiv:2303.04613

Rosenbluth, E. and Grohe, M. (2025) Repetition Makes Perfect: Recurrent Sum-GNNs Match Message Passing Limit. arXiv:2505.00291