

Learning and Reasoning with Graph Data: Integrating SRL and GNN

Manfred Jaeger




Aalborg University

- ▶ Learning and reasoning with graphs: from logic to graph neural networks
- ▶ A few notes on GNNs
- ▶ A few notes on SRL
- ▶ Relational Bayesian Networks
- ▶ GNN-RBN integration


Real-world networks

- ▶ Social networks
- ▶ Knowledge graphs
- ▶ Sensor networks
- ▶ Traffic networks
- ▶ ...





Real-world networks

- ▶ Social networks
- ▶ Knowledge graphs
- ▶ Sensor networks
- ▶ Traffic networks
- ▶ ...

**Abstract: graphs**

Real-world networks

- ▶ Social networks
- ▶ Knowledge graphs
- ▶ Sensor networks
- ▶ Traffic networks
- ▶ ...

**Abstract: graphs****Predicate logic (relational)**


$$\forall x(r(x) \rightarrow \exists y(e(x, y) \wedge b(y)))$$

$$\exists z, x, y \neg(e(x, y) \wedge e(x, z) \wedge (e(y, z)))$$

...

Real-world networks

- ▶ Social networks
- ▶ Knowledge graphs
- ▶ Sensor networks
- ▶ Traffic networks
- ▶ ...

**Abstract: graphs****Predicate logic (relational)**

$$\forall x(r(x) \rightarrow \exists y(e(x, y) \wedge b(y)))$$

$$\exists z, x, y \neg(e(x, y) \wedge e(x, z) \wedge (e(y, z)))$$

...


+*uncertainty*

➔

- ▶ Probabilistic Logic
- ▶ Statistical (Random) Graph Theory
- ▶ Statistical Relational Learning (Probabilistic Logic Learning)
- ▶ Graph Learning and Mining
- ▶ Graph Neural Networks

Real-world networks

- ▶ Social networks
- ▶ Knowledge graphs
- ▶ Sensor networks
- ▶ Traffic networks
- ▶ ...

**Abstract: graphs****Predicate logic (relational)**

$$\forall x(r(x) \rightarrow \exists y(e(x, y) \wedge b(y)))$$


$$\exists z, x, y \neg(e(x, y) \wedge e(x, z) \wedge (e(y, z)))$$

...


+uncertainty

- ▶ Probabilistic Logic
- ▶ Statistical (Random) Graph Theory
- ▶ **Statistical Relational Learning (Probabilistic Logic Learning)**
- ▶ Graph Learning and Mining
- ▶ **Graph Neural Networks**

Graph Representations




Graph: (V, E)



Graph: (V, E)


Attributed graph: (V, E, \mathbf{A}) . Node attributes \mathbf{A} : *Boolean, categorical, or numeric*



Graph: (V, E)

Attributed graph: (V, E, \mathbf{A}) . Node attributes \mathbf{A} : *Boolean, categorical, or numeric*

Attributed multirelational graph: $(V, \mathbf{E}, \mathbf{A})$. \mathbf{E} : set of different edge relations




Graph: (V, E)

Attributed graph: (V, E, \mathbf{A}) . Node attributes \mathbf{A} : *Boolean, categorical, or numeric*

Attributed multirelational graph: $(V, \mathbf{E}, \mathbf{A})$. \mathbf{E} : set of different edge relations

Attributed multirelational hyper-graph: (V, \mathbf{R}) . \mathbf{R} : set of 1,2,3,...-ary relations (subsumes \mathbf{A}, \mathbf{E})



Graph: (V, E)

Attributed graph: (V, E, \mathbf{A}) . Node attributes \mathbf{A} : *Boolean, categorical, or numeric*

Attributed multirelational graph: $(V, \mathbf{E}, \mathbf{A})$. \mathbf{E} : set of different edge relations


Attributed multirelational hyper-graph: (V, \mathbf{R}) . \mathbf{R} : set of 1,2,3,...-ary relations (subsumes \mathbf{A}, \mathbf{E})

Examples for higher arity relations (logic, relational databases):

3-ary traffic network relation: *on_shortest_path*(location,location,location)


3-ary movie data: *made_contract*(agent,actor,movie)

Node	color
1	blue
2	blue
3	green
4	red




(a)

Node	c_blue	c_green	c_red
1	1	0	0
2	1	0	0
3	0	1	0
4	0	0	1




(b)




(c)

- (a): unary, categorical values
- (b): unary, Boolean/binary values (one-hot encoding)
- (c): binary relation between objects and attribute values materialized as nodes (example: knowledge graphs)




(a)

(a): as tuples of nodes




(a)




(b)

(a): as tuples of nodes

(b): materialize tuples of nodes;




(a)




(b)

- (a): as tuples of nodes
- (b): materialize tuples of nodes;
connect tuple-nodes with entity-nodes by binary relations



(a)



(b)

(a): as tuples of nodes

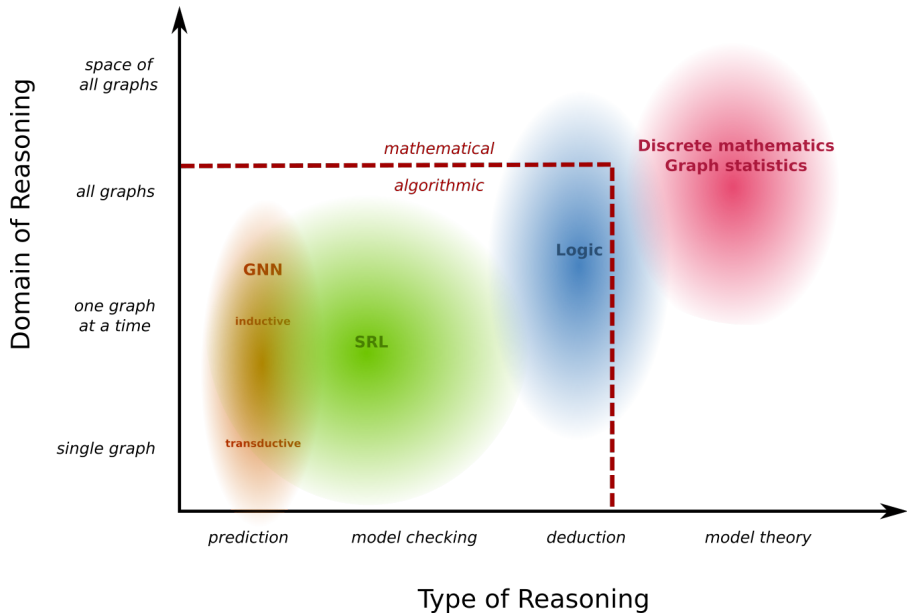
(b): materialize tuples of nodes;
connect tuple-nodes with entity-nodes by binary relations


➡ Categorical attributes and relations of higher arities can be reduced to Boolean attributes (one-hot-encodings) and binary relations (but this can be user-unfriendly).

N	number of nodes/vertices
\mathcal{R}	a <i>signature</i> of 1,2,3,...-ary relation symbols
\mathbf{R}	specific values of the relations in \mathcal{R} in a graph $G = (V, \mathbf{R})$.
$G = (V, \mathbf{R})$	a graph with node set V , and relations \mathbf{R}
$\mathcal{G}(V, \mathcal{R})$	set of all graphs with node set V , and relations in the signature \mathcal{R}
$\Delta\mathcal{G}(V, \mathcal{R})$	set of all probability distributions over $\mathcal{G}(V, \mathcal{R})$

Generally assume that $V = \{1, \dots, N\}$, and $i, j, \dots \in \mathbb{N}$ denote nodes.

Reasoning






Given: a probabilistic model for the random generation/evolution of graphs.

Question: what is the probability that the graph becomes (stays) connected, as the number of nodes goes to infinity?

➡ Or many other questions about the global properties of a random graph model.

➡ Mostly (human powered) mathematics, not algorithmic reasoning




Given: a knowledge base

$$\forall x \exists y \text{ follows}(x, y)$$
$$\exists y \neg \exists x \text{ follows}(x, y)$$

Question: Does the knowledge base imply a given query statement?


$$(\exists y \exists x \geq 2 \text{ follows}(x, y)) \vee \exists x \geq 10,000 x ?$$

- ➡ Reasoning about all possible graphs
- ➡ Algorithmic reasoning implemented by *theorem provers*.



Given: a generative probabilistic model for graphs.


Question: for a single partially observed graph, what are the probabilities of unobserved features?






Given: a generative probabilistic model for graphs.


Question: for a single partially observed graph, what are the probabilities of unobserved features?






Given: a generative probabilistic model for graphs.


Question: for a single partially observed graph, what are the probabilities of unobserved features?






Given: a generative probabilistic model for graphs.


Question: for a single partially observed graph, what are the probabilities of unobserved features?





Given: a discriminative model for specific node label.

Question: for an input graph (edges, node attributes), what are predicted node labels?




➡ Similarly: link prediction, graph classification.


Learning

Learning and Reasoning about a single graph:

Training data




Reasoning domain




Learning and Reasoning about different graphs:

Training data



Reasoning domains (a.k.a. test cases)



Graph Neural Networks: Basics

$\mathbf{h}^k(i)$: d^k -dimensional vector representation of node i at k th iteration (layer).

A basic form of message passing updates:

$$\begin{aligned}\mathbf{h}^0(i) &= \text{initial node feature vector of node } i \\ \mathbf{h}^{k+1}(i) &= f\left(\mathbf{W}^k \mathbf{h}^k(i) + \mathbf{U}^k \sum_{j \in N_i} \mathbf{h}^k(j)\right)\end{aligned}$$

with ingredients:

- ▶ $\mathbf{W}^k, \mathbf{U}^k$: weight matrices (dimensions: $d^{k+1} \times d^k$)
- ▶ f : (nonlinear) activation function (component-wise)

$\mathbf{h}^k(i)$: d^k -dimensional vector representation of node i at k th iteration (layer).

A basic form of message passing updates:

$$\begin{aligned} \mathbf{h}^0(i) &= \text{initial node feature vector of node } i \\ \mathbf{h}^{k+1}(i) &= f\left(\mathbf{W}^k \mathbf{h}^k(i) + \mathbf{U}^k \sum_{j \in N_i} \mathbf{h}^k(j)\right) \end{aligned}$$

with ingredients:

- ▶ $\mathbf{W}^k, \mathbf{U}^k$: weight matrices (dimensions: $d^{k+1} \times d^k$)
- ▶ f : (nonlinear) activation function (component-wise)


In full matrix notation:

$$\mathbf{H}^{k+1} = f\left(\mathbf{H}^k (\mathbf{W}^k)^T + \mathbf{E} \mathbf{H}^k (\mathbf{U}^k)^T\right)$$

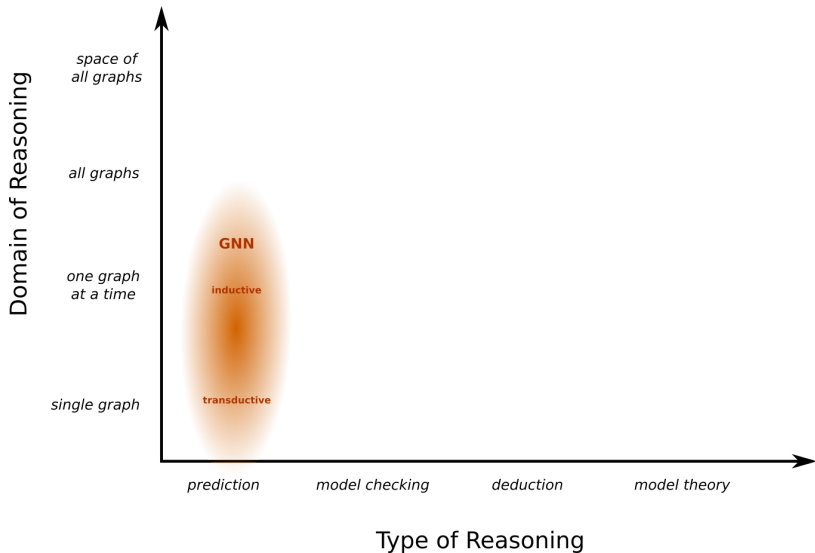
with ingredients:

- ▶ $\mathbf{H}^k, \mathbf{H}^{k+1}$: $n \times d^k$ and $n \times d^{k+1}$ matrices
- ▶ \mathbf{E} : $n \times n$ adjacency matrix


Representation as NN architecture/computation graph:



- ▶ At each layer: one vector for each node (picture: $N = 3$)
- ▶ At top: task-specific (node or graph classification) transformations of final node representations
- ▶ self, neighbors: dependence of vectors in following layer on previous layer




Initial features: node identifiers (typically: one-hot encoded).



Can represent/learn classification rule: node is *red*, if it has distance ≤ 3 to node 26.

➡ this only works in transductive settings.


Initial features: node attributes (e.g. $color \in \{yellow, blue\}$)



Can represent/learn classification rule: node is *red*, if it has distance ≤ 2 to a blue node.

→ this works in inductive settings: rule can be applied to new graphs with yellow/blue nodes.


Initial features: none (then can say e.g.: $\mathbf{h}^0(i) = 1$ for all i).



Can represent/learn classification rule: node is *red*, if it has distance ≤ 2 to a node with degree ≥ 5 .

➡ this works in inductive settings: rule can be applied to new graphs.

Discriminative power: when can two nodes be distinguished by a GNN?



a, b indistinguishable by any GNN.

a, c indistinguishable by 2-layer GNNs, distinguishable by 3-layer GNNs.

➡ l -layer GNNs can only access information in the $l - 1$ hop node neighborhood.

- ▶ GNNs cannot access “global” graph properties. Examples: cannot recognize whether a graph is connected/disconnected
- ▶ GNNs cannot reason about “identity” of nodes (unless node identifiers provided as initial features). Example: cannot recognize whether a node is a member of a clique of size ≥ 3 .

Main theorem of [Barceló et al.]:

*Every node property that can be expressed in the **two-variable fragment of first-order logic with counting quantifiers** (FOC_2) can be captured by an ACR-GNN.*

Example

In FOC_2 :

$$\alpha_1(X) \equiv \exists^{[8,10]} Y (blue(Y) \wedge \neg edge(X, Y))$$

("there exist 8-10 blue nodes that are not neighbors of X ")

Not in FOC_2 :

$$\delta(X) \equiv \exists Y, Z (X \neq Y \wedge X \neq Z \wedge Y \neq Z \wedge edge(X, Y) \wedge edge(X, Z) \wedge edge(Y, Z))$$

(" X is part of a triangle ")

[Barceló, Pablo, et al. "The logical expressiveness of graph neural networks." 2020]

Result from [Jaeger, Relational Bayesian Networks, 1997]:

Let $\phi(\mathbf{x})$ be a first-order formula over signature \mathcal{R} . Then there exists a probability formula $F_\phi(\mathbf{x})$ over \mathcal{R} , s.t. for every multi-relational graph $G = (V, \mathbf{R})$ and every $|\mathbf{x}|$ -tuple \mathbf{v} of nodes: $F_\phi(\mathbf{v}) = 1$ iff $\phi(\mathbf{v})$ holds in G (and $F_\phi(\mathbf{d}) = 0$ otherwise).

Statistical Relational Learning

An SRL *framework* consists of

- ▶ **Syntax:** a formal representation language over relational signatures \mathcal{R}
- ▶ **Semantics:** defines for any domain V , a probability distribution over the space $\mathcal{G}(V, \mathcal{R})$; formally: a mapping

$$V \mapsto P_V \in \Delta\mathcal{G}(V, \mathcal{R})$$

- ▶ **Inference (reasoning):** algorithms for the computation of *conditional probabilities*

$$P_V(A|B) \text{ for some } A, B \subseteq \mathcal{G}(V, \mathcal{R})$$

Also: computing *most probable explanation (MPE)*:

$$\max_{G \in \mathcal{G}(V, \mathcal{R})} P_V(G|B)$$

- ▶ **Learning:** methods for learning models from graph (relational) data. Typically divided into:
 - ▶ Structure learning: determines (logical) structure of the model (here also: knowledge-driven design)
 - ▶ Parameter learning: fitting numerical parameters

Representatives for main paradigms:

RBN	Directed probabilistic graphical models
MLN	Undirected probabilistic graphical models
ProbLog	(Inductive) logic programming

Relational Bayesian Networks

Chain Rule

For fixed V , P_V is a distribution over values $\mathbf{R} = (R_1, \dots, R_r)$. Let $R_{1:h} := (R_1, \dots, R_h)$. This distribution can be factored as

$$P_V(\mathbf{R}) = P_V(R_1) \cdot P_V(R_2|R_1) \cdot \dots \cdot P_V(R_h|R_{1:h-1}) \cdot \dots \cdot P_V(R_r|R_{1:r-1}).$$

Conditional independence of relations

Conditional independencies lead to simplifications:


$$P_V(R_h|R_{1:h-1}) = P_V(R_h|Pa(R_h)) \text{ for some } Pa(R_h) \subset R_{1:h-1}$$


↳ directed acyclic graph over relations (*relation DAG*).

$$P_V(\text{gender}, \text{republican}, \text{bloodtype}, \text{friends}) =$$


$$P_V(\text{gender})P_V(\text{republican}|\text{gender})P_V(\text{bloodtype}|\text{republican}, \text{gender})P_V(\text{friends}|\text{bloodtype}, \text{republican}, \text{gender}) \stackrel{\text{assume}}{=}$$

$$P_V(\text{gender})P_V(\text{republican}|\text{gender})P_V(\text{bloodtype}|\text{gender})P_V(\text{friends}|\text{republican}, \text{gender})$$





- Defines full generative probabilistic model for graphs in signature \mathcal{R}




- ▶ Defines full generative probabilistic model for graphs in signature \mathcal{R}
- ▶ Sometimes: assume some relations $R \in \mathcal{R}$ are **predefined input relations**:

$$\mathcal{R} = \mathcal{R}_{prob} \cup \mathcal{R}_{in}$$

- ▶ make these relations roots in the relation DAG
- ▶ do not define a distribution $P_V(R_h)$ for these relations
- ▶ defines a *conditional* distribution

$$P_V(\mathbf{R}_{prob} | \mathbf{R}_{in})$$



- ▶ Defines full generative probabilistic model for graphs in signature \mathcal{R}
- ▶ Sometimes: assume some relations $R \in \mathcal{R}$ are **predefined input relations**:

$$\mathcal{R} = \mathcal{R}_{prob} \cup \mathcal{R}_{in}$$

- ▶ make these relations roots in the relation DAG
- ▶ do not define a distribution $P_V(R_h)$ for these relations
- ▶ defines a *conditional* distribution

$$P_V(\mathbf{R}_{prob} | \mathbf{R}_{in})$$

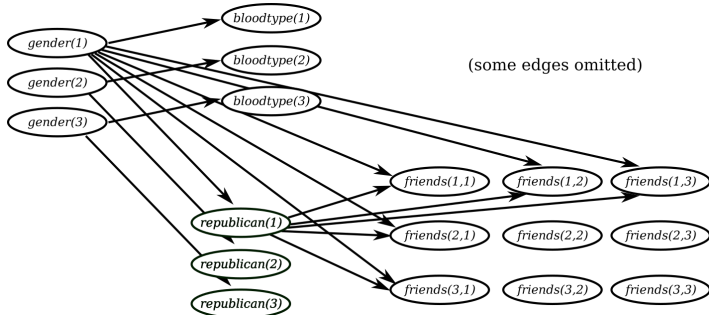
➡ All SRL frameworks support divisions $\mathcal{R} = \mathcal{R}_{prob} \cup \mathcal{R}_{in}$

Atom independence

Assume atoms of one relation are mutually independent, given the parent relations:

$$P_V(R_h | Pa(R_h)) := \prod_{i \in \text{arity}(R_h)} P_V(R_h(i) | Pa(R_h))$$

As a Bayesian network:



↳ Leads to limitations for modeling e.g. symmetry constraints $friends(1,2) \Leftrightarrow friends(2,1)$, or homophily (exist modeling tricks to circumvent this!).

A relational Bayesian network for signature \mathcal{R} consists of

- ▶ a directed acyclic graph whose nodes are the relations $R \in \mathcal{R}$,
- ▶ for each $R \in \mathcal{R}$ a *probability formula* F_R in the signature $Pa(R)$ that defines the conditional probabilities

$$P_V(R(\mathbf{i})|Pa(R))$$

Probability formulas: semantics

A probability formula F maps tuples of entities \mathbf{i} in a graph $G = (V, \mathbf{R})$ to a probability value

$$eval(F, \mathbf{i}, G) \in [0, 1]$$

[M. Jaeger: Relational Bayesian Networks. UAI 1997]

Constants

For any $q \in [0, 1]$,

$$F \equiv q$$

is a probability formula with

$$\text{eval}(F, i, G) = q$$

for all i, G .

Example

Let $\mathcal{R} = \{edge\}$. Then

$$F_{edge(x,y)} \equiv 0.5$$

defines the classic Erdős-Rényi random graph model.

Atoms

For any $R \in \mathcal{R}$, and variables $Y_1, \dots, Y_{arity(R)}$

$$F \equiv R(Y_1, \dots, Y_{arity(R)})$$

is a probability formula with

$$eval(F, \mathbf{i}, G) = \begin{cases} 1 & \text{if } R(\mathbf{i}) \text{ is true in } G \\ 0 & \text{if } R(\mathbf{i}) \text{ is false in } G \end{cases}$$

WIF-THEN-ELSE


If F_1, F_2, F_3 are probability formulas, then

$$F \equiv \text{WIF } F_1 \text{ THEN } F_2 \text{ ELSE } F_3$$

is a probability formula with

$$eval(F, \mathbf{i}, G) = eval(F_1, \mathbf{i}, G)eval(F_2, \mathbf{i}, G) + (1 - eval(F_1, \mathbf{i}, G))eval(F_3, \mathbf{i}, G)$$

↳ Generalization of Boolean operations ($F_i \in \{0, 1\}$)



- ▶ Nodes partitioned into *blocks*
- ▶ Probability of edges depends on block memberships

With the constructs introduced so far:

A. partitioning into red, green, blue nodes:


$$F_{red(x)} \equiv 0.5$$

$$F_{blue(x)} \equiv \text{WIF } red(X) \text{ THEN } 0 \text{ ELSE } 0.7$$

$$F_{green(x)} \equiv \text{WIF } red(X) \vee blue(X) \text{ THEN } 0 \text{ ELSE } 1.0$$

B. generating edges:

$$F_{edge(x,y)} \equiv \text{WIF } red(X) \wedge red(Y) \text{ THEN } 0.6 \text{ ELSEIF } \dots$$




Combination Function

(related to first-order quantifiers \forall, \exists , GNN message passing aggregation, ...)

If F_1, \dots, F_t are probability formulas, then

$$F \equiv \begin{array}{l} \text{COMBINE } F_1, \dots, F_t \\ \text{WITH } \langle \textit{combination function} \rangle \\ \text{FORALL } \langle \textit{variables} \rangle \\ \text{WHERE } \langle \textit{logical constraint} \rangle \end{array}$$


is a probability formula.




$\mathcal{R} = \{red, blue, edge\}$. In figure: yellow \sim not blue; black \sim not red.

$P_V(red(i))$ higher if

- ▶ i is blue
- ▶ i is part of many triangles






$\mathcal{R} = \{red, blue, edge\}$. In figure: yellow \sim not blue; black \sim not red.


$P_V(red(i))$ higher if

- ▶ i is blue
- ▶ i is part of many triangles



Defining triangles:


$$F_{triangle(x,y,z)} \equiv edge(x, y) \wedge edge(x, z) \wedge edge(y, z)$$



$\mathcal{R} = \{\text{red}, \text{blue}, \text{edge}\}$. In figure: yellow \sim not blue; black \sim not red.

$P_V(\text{red}(i))$ higher if

- ▶ i is blue
- ▶ i is part of many triangles




Defining triangles:

$$F_{\text{triangle}(X,Y,Z)} \equiv \text{edge}(X, Y) \wedge \text{edge}(X, Z) \wedge \text{edge}(Y, Z)$$

Counting triangles:


$$F_{\text{triangle_count}(X)} \equiv \text{COMBINE 1.0} \\ \text{WITH } \text{sum} \\ \text{FORALL } Y, Z \\ \text{WHERE } F_{\text{triangle}(X,Y,Z)}(X, Y, Z)$$



$\mathcal{R} = \{red, blue, edge\}$. In figure: yellow \sim not blue; black \sim not red.

$P_V(red(i))$ higher if

- ▶ i is blue
- ▶ i is part of many triangles



Defining triangles:

$$F_{triangle(x,y,z)} \equiv edge(x,y) \wedge edge(x,z) \wedge edge(y,z)$$

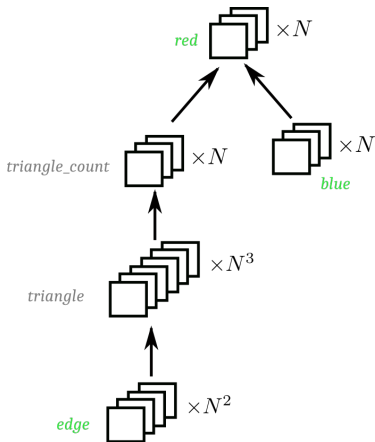
Counting triangles:

$$F_{triangle_count(x)} \equiv \begin{array}{l} \text{COMBINE } 1.0 \\ \text{WITH } sum \\ \text{FORALL } Y, Z \\ \text{WHERE } F_{triangle(x,y,z)}(X, Y, Z) \end{array}$$

Logistic regression of *triangle_count* and *blue* feature:

$$F_{red(x)} \equiv \begin{array}{l} \text{COMBINE } 0.6 \cdot F_{triangle_count(x)}(X), \\ \quad 0.3 \cdot blue(X), \\ \quad -3.0 \\ \text{WITH } logistic\ regression \end{array}$$

The computation graph of the probability formula for *red*. In **green**: relations from \mathcal{R} . In **gray**: synthetic names for intermediate formulas (“layers”).



- ▶ Each probability (sub-)formula defines a feature of 0, 1, 2, ...-tuples of entities
- ▶ Nested formulas give “deep” models
- ▶ Aggregation (message passing) along “channels” defined by the constraints in combination functions
- ▶ Scalar features

GNN-2-RBN

From [Barceló et al.,2020]:

Input graphs defined by signature:


$$\mathcal{R}_{in} = \{blue, green, red, yellow, purple, edge\}$$

Target concept to represent/learn:

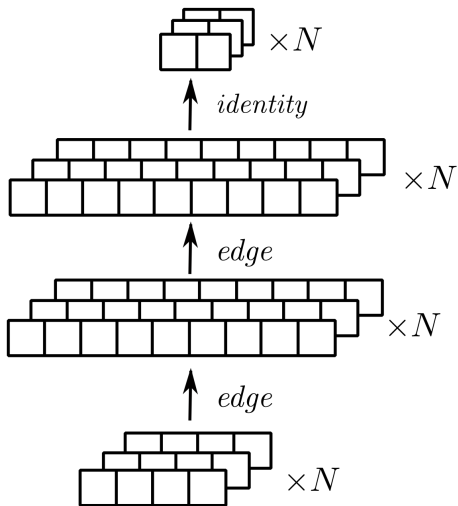
$$\alpha_1(X) \equiv \exists^{[8,10]} Y (blue(Y) \wedge \neg edge(X, Y))$$

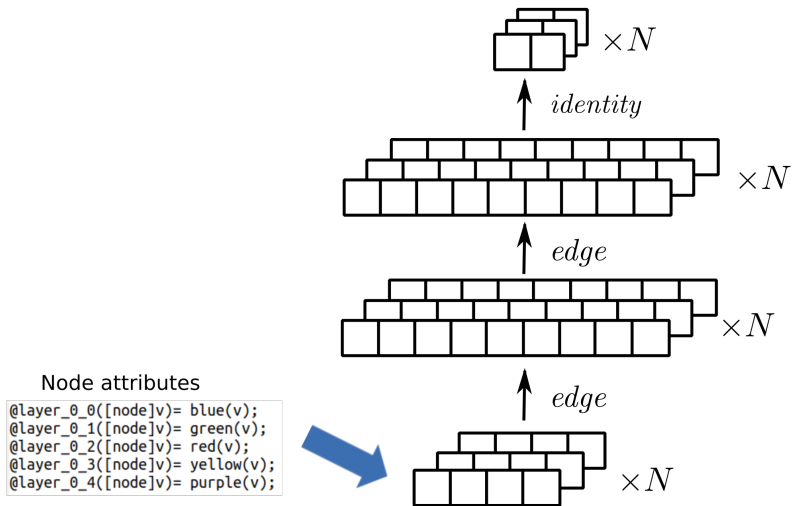
(cf. slide 22)

A GNN α_1 classifier defines a conditional distribution $P(\alpha_1 | blue, green, red, yellow, purple, edge)$ satisfying the *Atom Independence* property (slide 28).



➡ This distribution can be encoded by a probability formula.






Linear combination and activation

```
@layer_1_0([node]v)= COMBINE
  ($c_1_0_0*@layer_0_0(v)),
  ($c_1_0_1*@layer_0_1(v)),
  ($c_1_0_2*@layer_0_2(v)),
  ($c_1_0_3*@layer_0_3(v)),
  ($c_1_0_4*@layer_0_4(v)),
  ($A_1_0_0*@agg_0_0(v)),
  ($A_1_0_1*@agg_0_1(v)),
  ($A_1_0_2*@agg_0_2(v)),
  ($A_1_0_3*@agg_0_3(v)),
  ($A_1_0_4*@agg_0_4(v)),
  ($R_1_0_0*@read_0_0()),
  ($R_1_0_1*@read_0_1()),
  ($R_1_0_2*@read_0_2()),
  ($R_1_0_3*@read_0_3()),
  ($R_1_0_4*@read_0_4()),
  $b_1_0
  WITH l-reg
```

Aggregation

```
@agg_0_1([node]v) = COMBINE @layer_0_1(w)
  WITH sum
  FORALL w
  WHERE (edge(v,w)|edge(w,v));
```



Linear combination and activation


```

@layer_1_0([node]v)= COMBINE
  ($c_1_0_0*@layer_0_0(v)),
  ($c_1_0_1*@layer_0_1(v)),
  ($c_1_0_2*@layer_0_2(v)),
  ($c_1_0_3*@layer_0_3(v)),
  ($c_1_0_4*@layer_0_4(v)),
  ($A_1_0_0*@agg_0_0(v)),
  ($A_1_0_1*@agg_0_1(v)),
  ($A_1_0_2*@agg_0_2(v)),
  ($A_1_0_3*@agg_0_3(v)),
  ($A_1_0_4*@agg_0_4(v)),
  ($R_1_0_0*@read_0_0()),
  ($R_1_0_1*@read_0_1()),
  ($R_1_0_2*@read_0_2()),
  ($R_1_0_3*@read_0_3()),
  ($R_1_0_4*@read_0_4()),
  $b_1_0
WITH l-reg
  
```

Aggregation

```

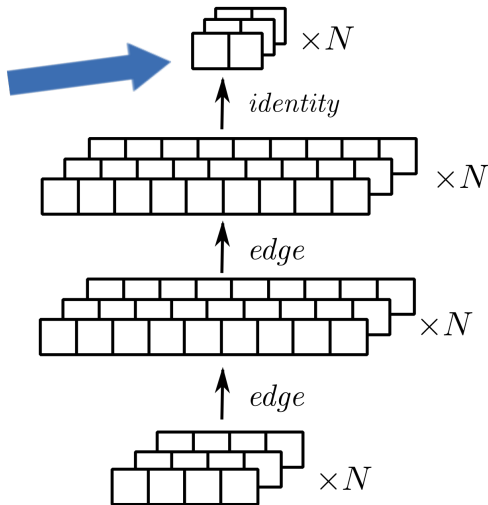
@agg_0_1([node]v) = COMBINE @layer_0_1(w)
WITH sum
FORALL w
WHERE (edge(v,w)|edge(w,v));
  
```



Yellow highlight: trainable parameters \sim entries of GNN weight matrices


Linear classifier

```
alpha1([node]v)= COMBINE
  ($w_0*@layer_1_0(v)),
  ($w_1*@layer_1_1(v)),
  ($w_2*@layer_1_2(v)),
  ($w_3*@layer_1_3(v)),
  $w_5
  WITH l-reg
```



Yellow highlight: trainable parameters \sim entries of GNN weight matrices

- ▶ One-to-one mapping of representation and parameterization
- ▶ Matrix-vector level GNN specifications broken down to the “scalar” level
- ▶ GNN training \sim RBN learning (same objective, same gradients, ...)




Yellow highlight: trainable parameters \sim entries of GNN weight matrices


Learning the α_1 target.

Training data: 5000 random graphs of size $N \in 40..50$ (data from [Barceló et al.]).

Pytorch geometric implementation
of ACR-GNN:



Primula implementation of
RBN encoding:




(blue: loss, red: accuracy (on training data); 20 epochs, 10 restarts with random parameter initializations)


Learning the α_1 target.

Training data: 5000 random graphs of size $N \in 40..50$ (data from [Barceló et al.]).

Pytorch geometric implementation
of ACR-GNN:




Primula implementation of
RBN encoding:



(blue: loss, red: accuracy (on training data); 20 epochs, 10 restarts with random parameter initializations)


But: Primula takes **much** longer ...

Building a conditional generative model:



Determine relational dependencies and input relations


Building a conditional generative model:



Determine relational dependencies and input relations

Rich data, little knowledge: use GNN modules (many parameters, little structure) to define conditional distributions

Building a conditional generative model:




Determine relational dependencies and input relations

Rich data, little knowledge: use GNN modules (many parameters, little structure) to define conditional distributions

Sparse data, expert knowledge, constraints: use customized probability formula (few parameters, highly structured) to define conditional distributions

Building a conditional generative model:




Determine relational dependencies and input relations

Rich data, little knowledge: use GNN modules (many parameters, little structure) to define conditional distributions

Sparse data, expert knowledge, constraints: use customized probability formula (few parameters, highly structured) to define conditional distributions

➡ No a-priori distinction of low-level perceptual vs. high-level cognitive reasoning (cf. DeepProbLog)

Building a conditional generative model:



Determine relational dependencies and input relations


Rich data, little knowledge: use GNN modules (many parameters, little structure) to define conditional distributions

Sparse data, expert knowledge, constraints: use customized probability formula (few parameters, highly structured) to define conditional distributions

➡ No a-priori distinction of low-level perceptual vs. high-level cognitive reasoning (cf. DeepProbLog)


➡ The resulting neuro-symbolic model supports all types of model checking reasoning.

Making the α_1 model generative:




Conditional (prediction) model for α_1 given all other relations as input.

Making the α_1 model generative:



Conditional (prediction) model for α_1 given all other relations as input.



Generative model for node attributes and label, given *edge* as input

$$F_{yellow(x)} = 0.18;$$


$$F_{blue(x)} = 0.26;$$

$$F_{red(x)} = 0.18;$$


$$F_{green(x)} = 0.18;$$

$$F_{purple(x)} = 0.18;$$

Making the α_1 model generative:



Conditional (prediction) model for α_1 given all other relations as input.



Generative model for node attributes and label, given *edge* as input

$$F_{yellow(x)} = 0.18;$$


$$F_{blue(x)} = 0.26;$$

$$F_{red(x)} = 0.18;$$

$$F_{green(x)} = 0.18;$$


$$F_{purple(x)} = 0.18;$$

MPE task: given observed α_1 labels, what is the most probable configuration of the *blue* attribute?




- ▶ Graph with observed α_1 relation (21 nodes)

[R. Pojer, A. Passerini, M. Jaeger: Generalized Reasoning with Graph Neural Networks by Relational Bayesian Network Encodings. In Learning on Graphs Conference (2023)].




- ▶ Graph with observed α_1 relation (21 nodes)
- ▶ MAP for *blue* with RBN-GNN manually set parameters (exactly implementing logical definition of α_1 ; test accuracy: 1.0).

[R. Pojer, A. Passerini, M. Jaeger: Generalized Reasoning with Graph Neural Networks by Relational Bayesian Network Encodings. In Learning on Graphs Conference (2023)].



- ▶ Graph with observed α_1 relation (21 nodes)
- ▶ MAP for *blue* with RBN-GNN manually set parameters (exactly implementing logical definition of α_1 ; test accuracy: 1.0).
- ▶ MAP for *blue* with RBN-GNN learned parameters (approximately implementing logical definition of α_1 ; test accuracy: 1.0).

[R. Pojer, A. Passerini, M. Jaeger: Generalized Reasoning with Graph Neural Networks by Relational Bayesian Network Encodings. In Learning on Graphs Conference (2023)].



- ▶ Graph with observed α_1 relation (21 nodes)
- ▶ MAP for *blue* with RBN-GNN manually set parameters (exactly implementing logical definition of α_1 ; test accuracy: 1.0).
- ▶ MAP for *blue* with RBN-GNN learned parameters (approximately implementing logical definition of α_1 ; test accuracy: 1.0).

▶ perfect accuracy on primary prediction task does not guarantee perfect accuracy for other reasoning tasks.

[R. Pojer, A. Passerini, M. Jaeger: Generalized Reasoning with Graph Neural Networks by Relational Bayesian Network Encodings. In Learning on Graphs Conference (2023)].

- ▶ There is more to reasoning than prediction!
- ▶ GNNs: good at learning accurate predictors from data
- ▶ SRL: good at flexible reasoning
- ▶ RBNs: the SRL framework most closely related to GNNs
- ▶ RBN+GNN: seamless integration of GNN prediction models into SRL model