

Motivation

Logic-based Expert Systems

- No training data
- Interpretable
- No generalization beyond what is manually defined in rules

Representation Learning


- Behavior is learned from input-output examples
- Achieves strong generalization
- Needs a lot of training data
- Generally not interpretable

Can we get the best of both worlds?

Differentiable Backward Chaining

- Neural network for proving queries to a knowledge base
- Proof success is differentiable with respect to vector representations of symbols
- Learn vector representations of symbols using SGD
- Make use of provided rules in soft proofs
- Induce interpretable first-order logic rules using SGD

Proof States and Modules



- Proof state $S = (\varphi, \rho)$ is a tuple consisting of
 - S_φ : Substitution set (variable bindings)
 - S_ρ : Neural network calculating real-valued proof success
- Modules map upstream proof state to a list of new proof states
 - Extending the substitution set (adding variable bindings)
 - Extending the neural network (adding nodes to comp. graph)

Unification

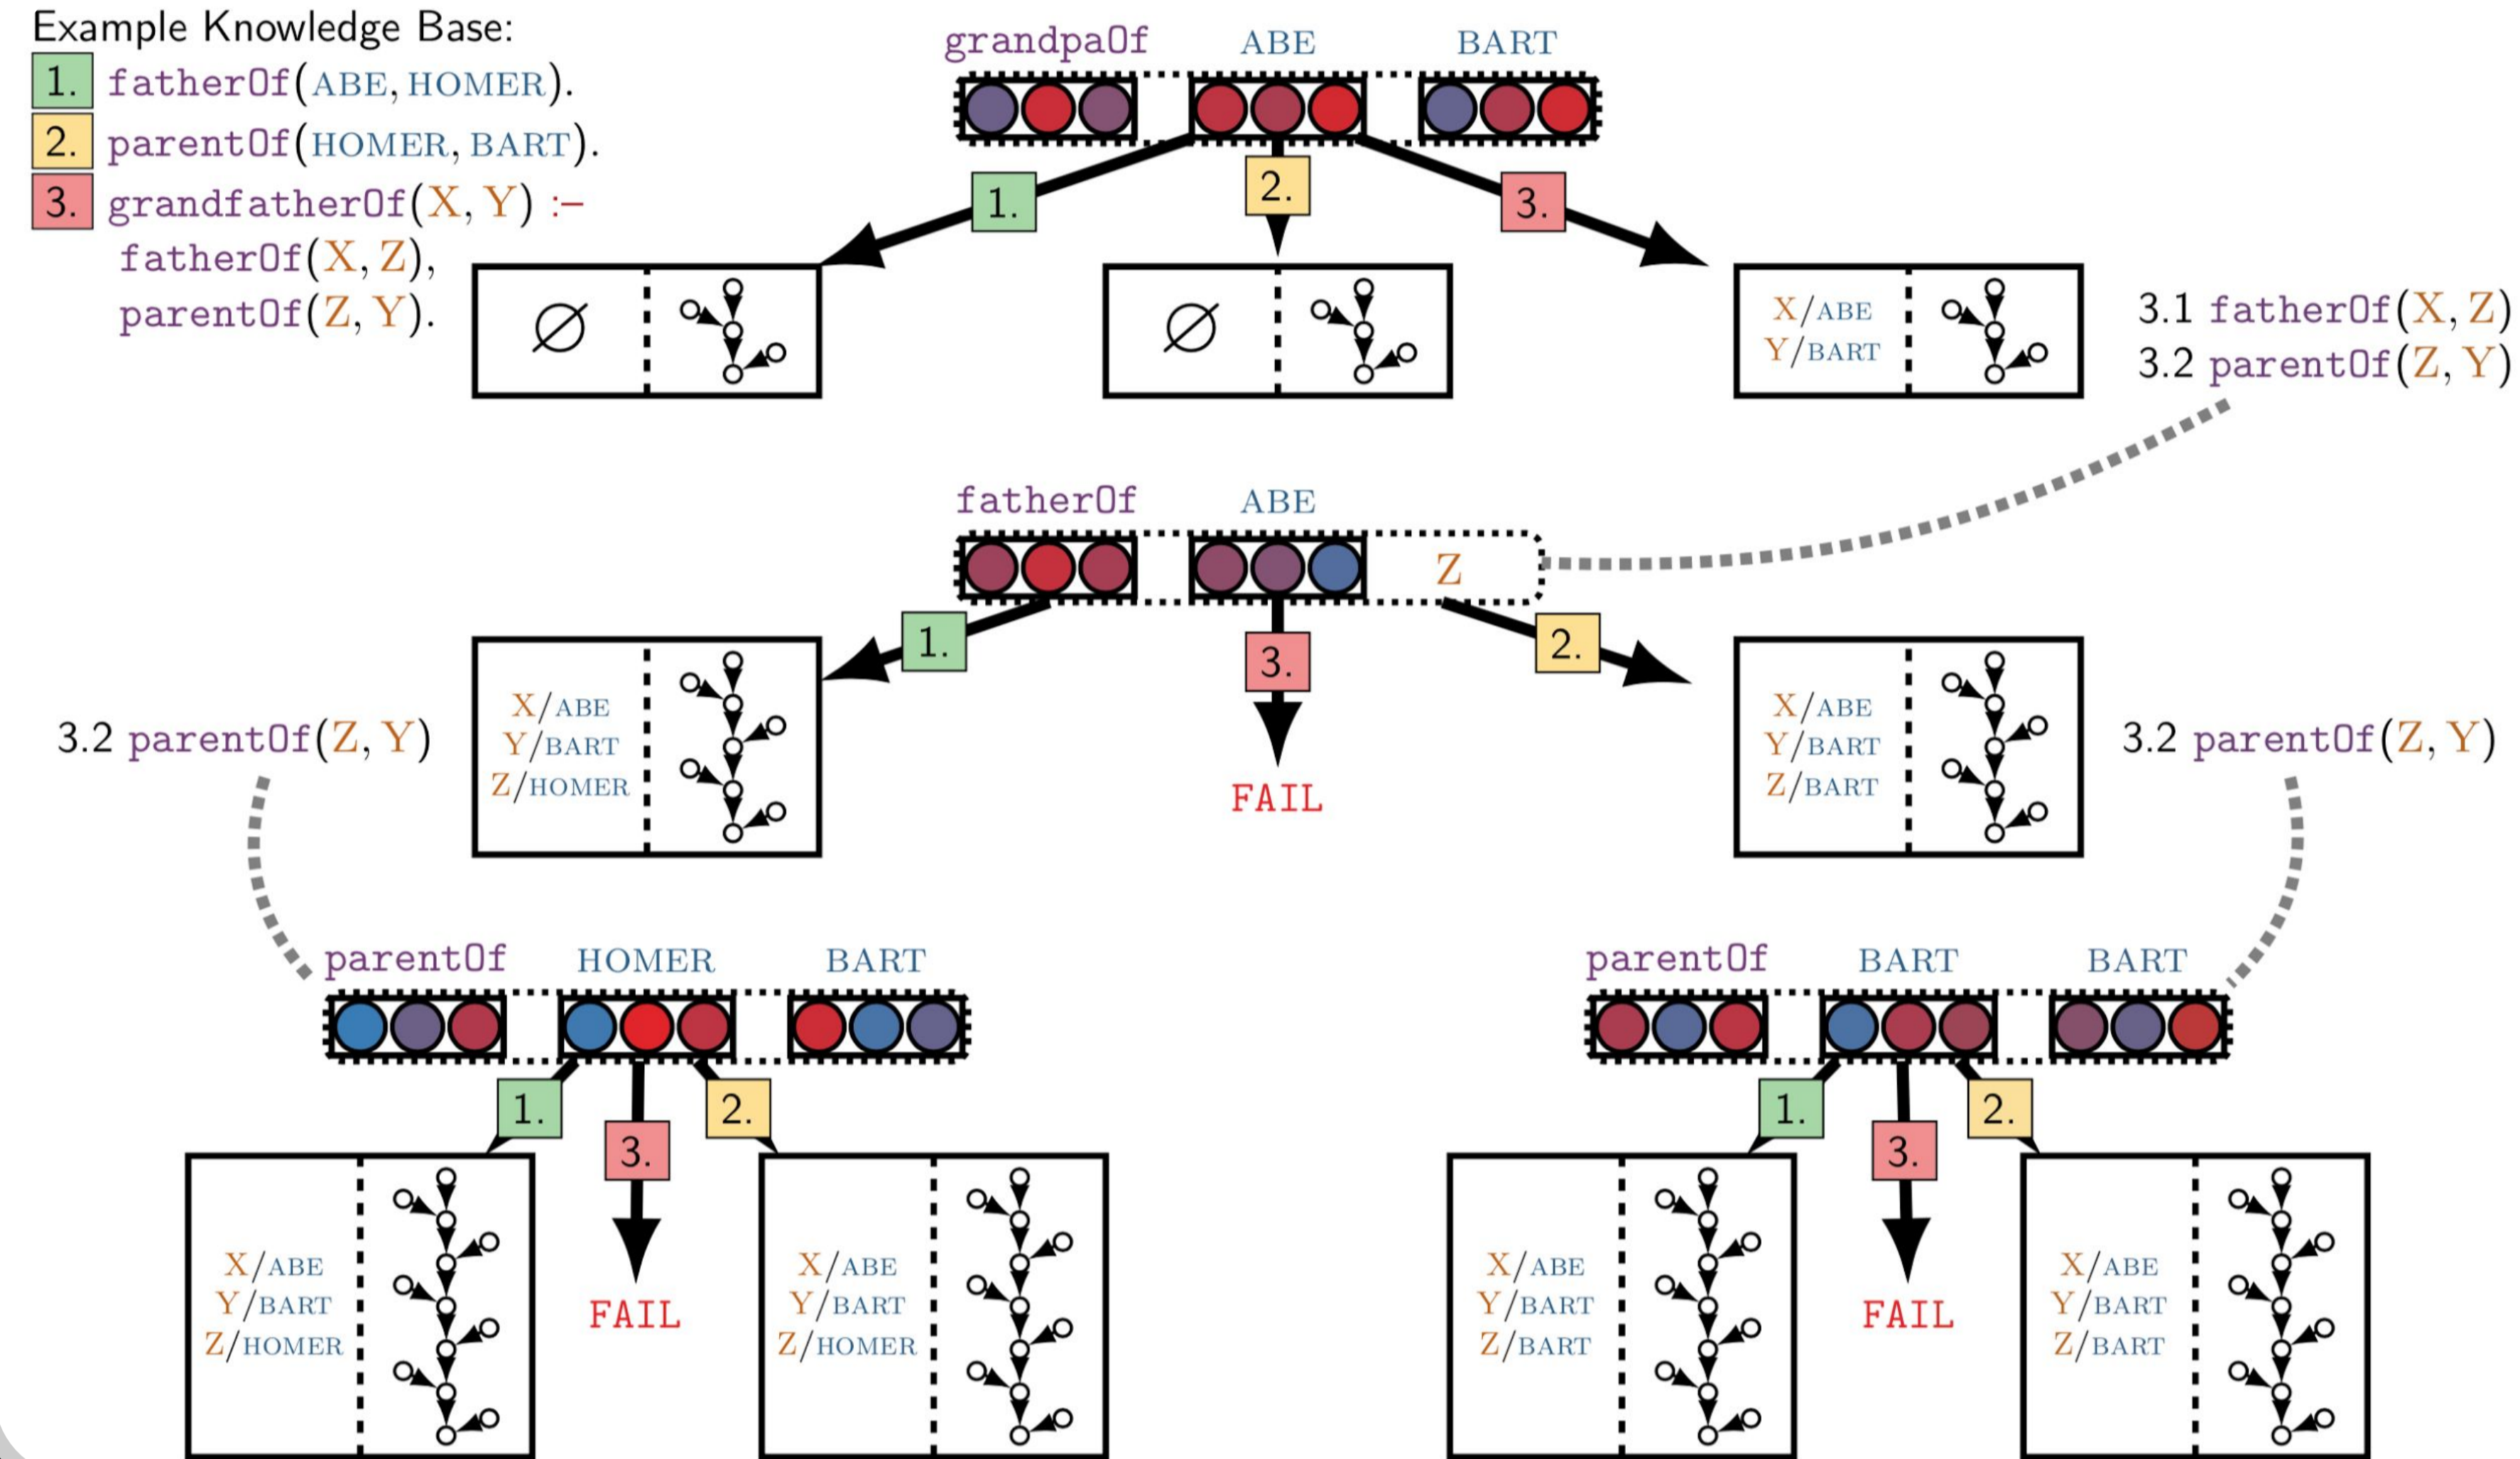
- Update substitution set S_φ by creating new variable bindings
- Compare vector representations of non-variable symbols using a Radial Basis Function kernel (extending neural net S_ρ)

1. $\text{unify}_\theta([], [], S) = S$
2. $\text{unify}_\theta([], _, _) = \text{FAIL}$
3. $\text{unify}_\theta(_, [], _) = \text{FAIL}$
4. $\text{unify}_\theta(h : H, g : G, S) = \text{unify}_\theta(H, G, S') = (S'_\varphi, S'_\rho)$ where
$$S'_\varphi = \begin{cases} S_\varphi \cup \{h/g\} & \text{if } h \in \mathcal{V} \\ S_\varphi \cup \{g/h\} & \text{if } g \in \mathcal{V}, h \notin \mathcal{V} \\ S_\varphi & \text{otherwise} \end{cases}, \quad S'_\rho = \min \left(S_\rho, \begin{cases} \exp\left(\frac{-\|\theta_h - \theta_g\|_2}{2r^2}\right) & \text{if } h, g \notin \mathcal{V} \\ 1 & \text{otherwise} \end{cases} \right)$$

Example

Example Knowledge Base:

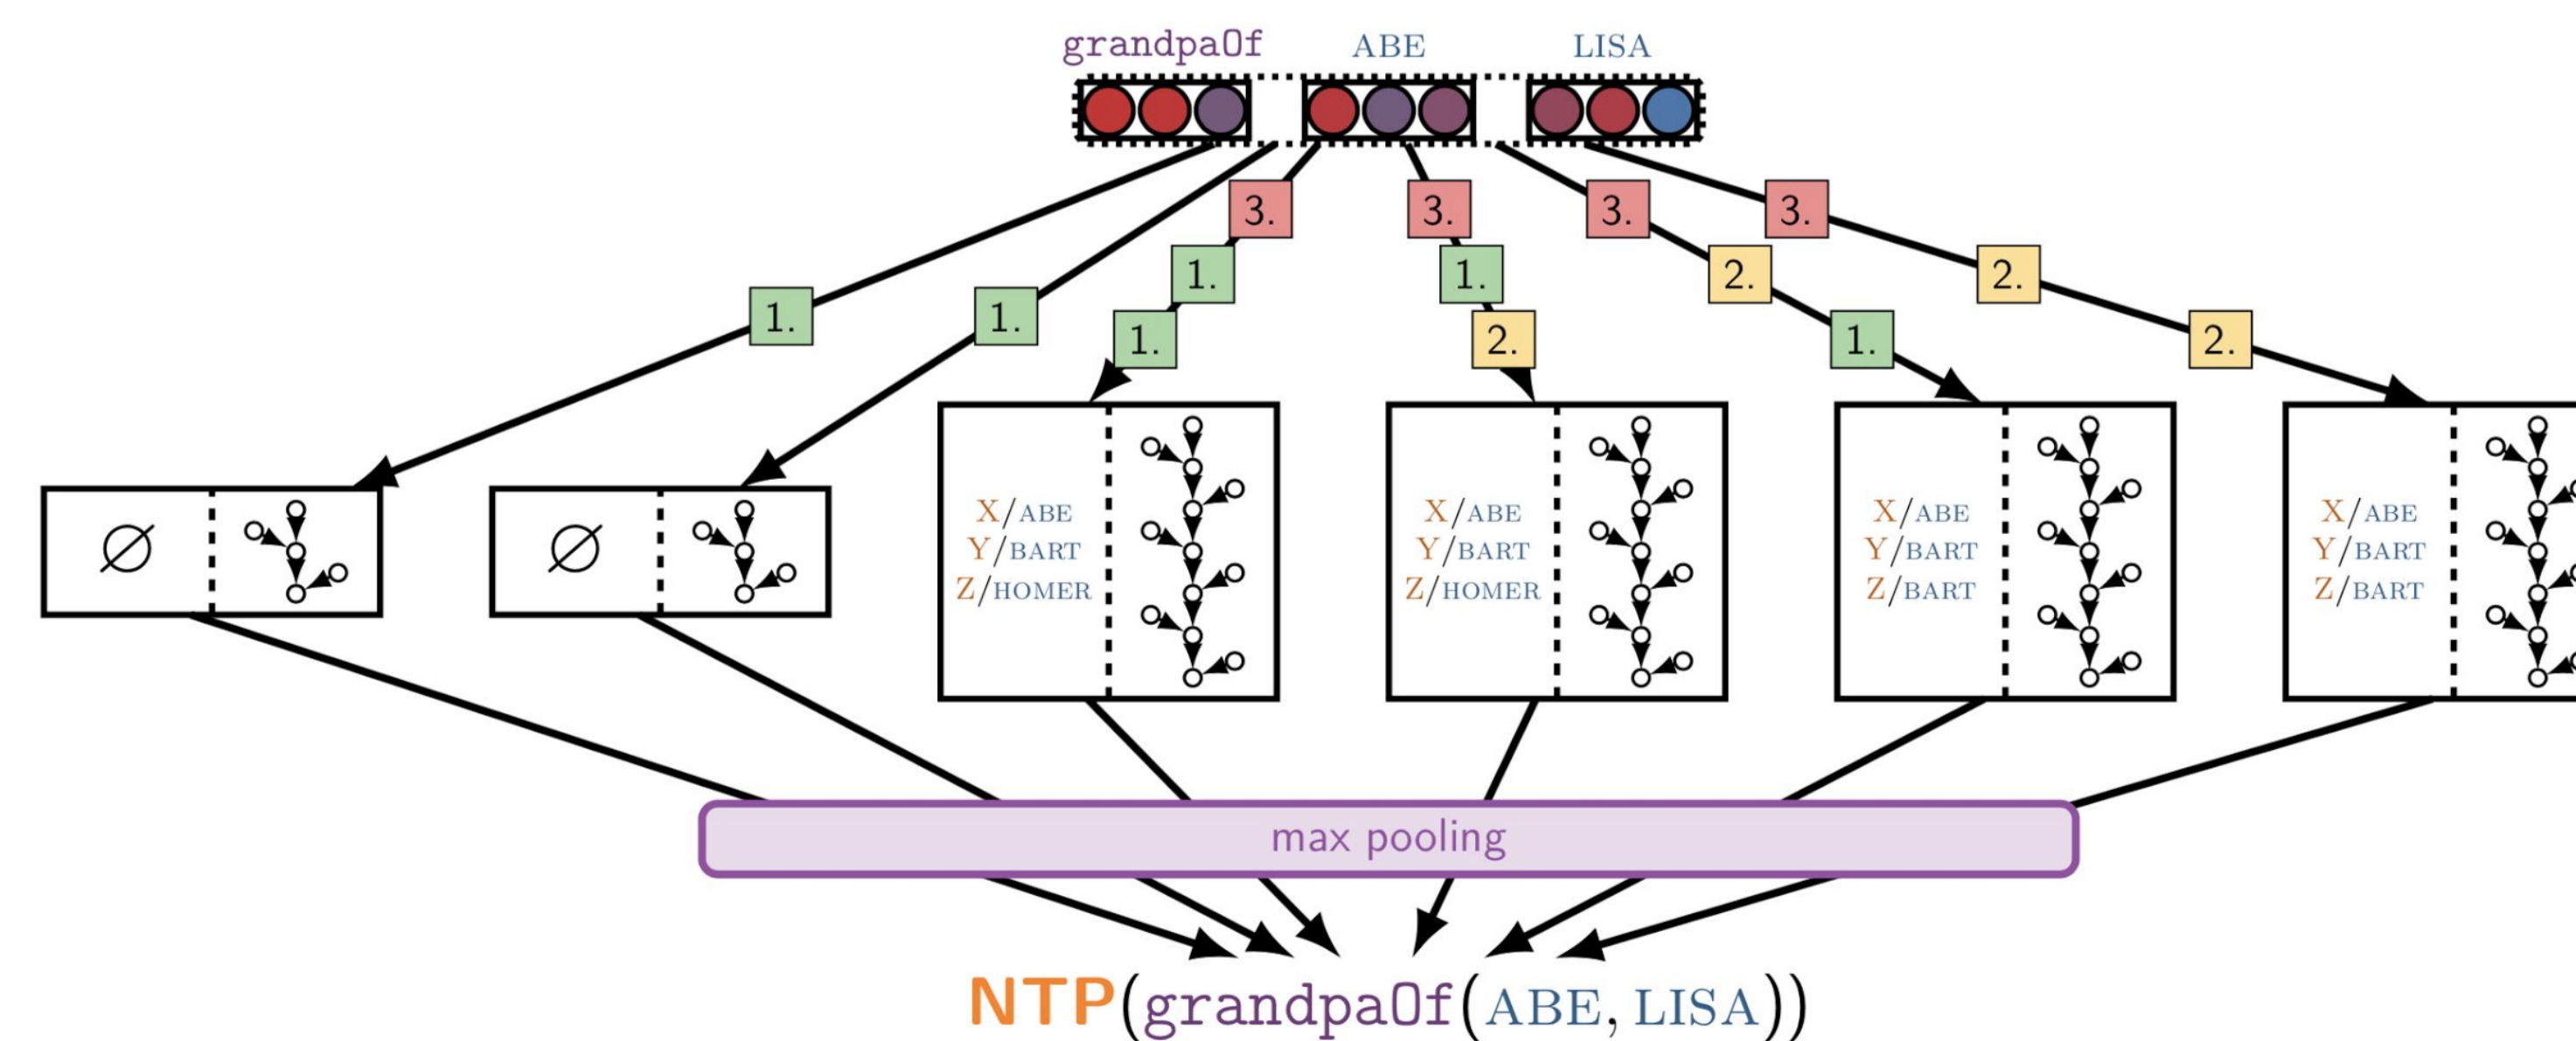
1. $\text{fatherOf}(\text{ABE}, \text{HOMER}).$
2. $\text{parentOf}(\text{HOMER}, \text{BART}).$
3. $\text{grandfatherOf}(X, Y) :- \text{fatherOf}(X, Z), \text{parentOf}(Z, Y).$



Recursion

- Iterate through all rules in the knowledge base and unify goal with rule heads
 1. $\text{or}_\theta^{\mathbb{R}}(G, d, S) = [S' \mid S' \in \text{and}_\theta^{\mathbb{R}}(\mathbb{B}, d, \text{unify}_\theta(H, G, S')) \text{ for } H :- \mathbb{B} \in \mathbb{R}]$
- Recursively prove subgoals in rule body
 1. $\text{and}_\theta^{\mathbb{R}}(_, _, \text{FAIL}) = \text{FAIL}$
 2. $\text{and}_\theta^{\mathbb{R}}(_, 0, _) = \text{FAIL}$
 3. $\text{and}_\theta^{\mathbb{R}}([], _, S) = S$
 4. $\text{and}_\theta^{\mathbb{R}}(G : \mathbb{G}, d, S) = [S'' \mid S'' \in \text{and}_\theta^{\mathbb{R}}(\mathbb{G}, d, S') \text{ for } S' \in \text{or}_\theta^{\mathbb{R}}(\text{substitute}(G, S_\psi), d - 1, S)]$

Training Objective



- Loss: $L(r_s(e_i, e_j), y) = -y \log(\text{NTP}(r_s(e_i, e_j))) - (1-y) \log(1 - \text{NTP}(r_s(e_i, e_j)))$
- **NTP** variant: SotA neural link prediction model (**Complex**) as auxiliary task

Neural Inductive Logic Programming

- Architecture allows us to induce rules of predefined structure
- We can, for instance, incorporate the inductive bias of a transitivity relationship in the knowledge base
$$\theta_1(X, Y) :- \theta_2(X, Z), \theta_3(Z, Y).$$
- θ_i are vector representations for unknown predicates
- They can be learned like all other vector representations
- They can be decoded at test time by finding the closest known relation using the RBF kernel
- Rule confidence is minimum RBF similarity over all decodings
- Confidence is an upper bound on the proof success that can be achieved when applying the induced rule

Results

Knowledge Base	Metric	Model		
		Complex	NTP	NTP λ
Countries	S1 AUC-PR	99.37 \pm 0.4	90.83 \pm 15.4	100.00 \pm 0.0
	S2 AUC-PR	87.95 \pm 2.8	87.40 \pm 11.7	93.04 \pm 0.4
	S3 AUC-PR	48.44 \pm 6.3	56.68 \pm 17.6	77.26 \pm 17.0
Kinship	MRR	0.81	0.60	0.80
	HITS@1	0.70	0.48	0.76
Nations	MRR	0.75	0.75	0.74
	HITS@1	0.62	0.62	0.59
UMLS	MRR	0.89	0.88	0.93
	HITS@1	0.82	0.82	0.87

Induced Rules

Knowledge Base	Examples of induced rules and their confidence	
Countries	S1	0.90 $\text{locatedIn}(X, Y) :- \text{locatedIn}(X, Z), \text{locatedIn}(Z, Y).$
	S2	0.63 $\text{locatedIn}(X, Y) :- \text{neighborOf}(X, Z), \text{locatedIn}(Z, Y).$
	S3	0.32 $\text{locatedIn}(X, Y) :- \text{neighborOf}(X, Z), \text{neighborOf}(Z, W), \text{locatedIn}(W, Y).$
Nations		0.68 $\text{blockpositionindex}(X, Y) :- \text{blockpositionindex}(Y, X).$
		0.46 $\text{expeldiplomats}(X, Y) :- \text{negativebehavior}(X, Y).$
		0.38 $\text{negativecomm}(X, Y) :- \text{commonbloc}(X, Y).$
UMLS		0.38 $\text{intergovorgs3}(X, Y) :- \text{intergovorgs}(Y, X).$
		0.88 $\text{interacts_with}(X, Y) :- \text{interacts_with}(X, Z), \text{interacts_with}(Z, Y).$
		0.71 $\text{isa}(X, Y) :- \text{isa}(X, Z), \text{isa}(Z, Y).$

Limitations and Future Work

- Scale to larger knowledge bases (beyond 10k facts)
 - Hierarchical attention for unification with facts
 - Reinforcement learning for pruning proof tree
- Train jointly with RNNs that encode natural language statements which can then be used in proofs
- Learn to prove mathematical theorems
- Incorporate commonsense knowledge for Visual Q&A