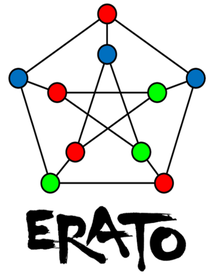

What Can Neural Networks Reason About?

ICLR 2020 Spotlight

Keyulu Xu, Jingling Li, Mozhi Zhang, Simon S. Du

Ken-ichi Kawarabayashi, Stefanie Jegelka

Future of AI - Reasoning, beyond pattern recognition

Learn features

Learn reasoning process

Pattern recognition
for perception

Reason about
movement of blocks

Reasoning tasks

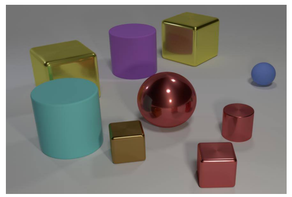

Perception

Universe
(collection of objects)

h_1

…

h_n

?

Neural Network

answer

Learns reasoning process (algorithm)

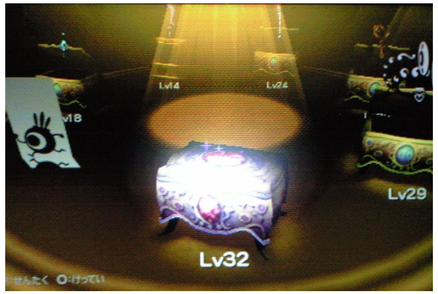

Summary statistics
What is the maximum value difference among treasures?

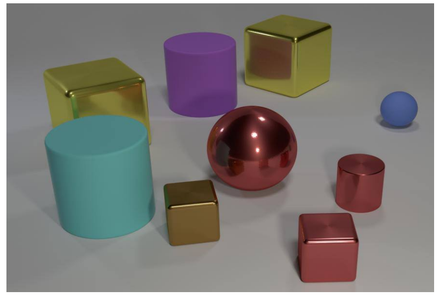

Relational argmax
What are the colors of the furthest pair of objects?

Dynamic programming
What is the cost to defeat monster X by following the optimal path?

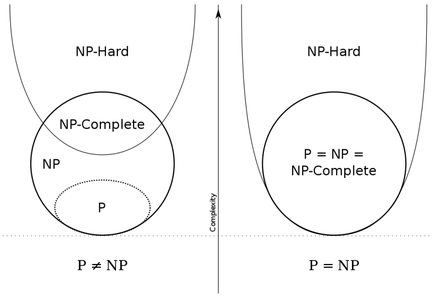

NP-hard problem
Subset sum: Is there a subset that sums to 0?

Challenge for NNs to learn reasoning

MLP

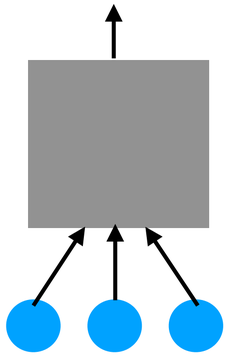

Deep Sets
(Zaheer et al 2017)

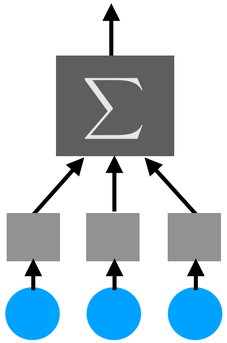

$$y = \mathrm{MLP}_2\left(\sum_{s\in S}\mathrm{MLP}_1(X_s)\right)$$

Graph Neural Network
(GNN)

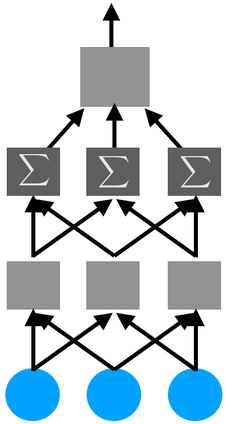

$$h_s^{(k+1)} = \sum_{t\in S}\mathrm{MLP}_1^{(k)}\left(h_s^{(k)}, h_t^{(k)}\right) \quad y = \mathrm{MLP}_2\left(\sum_{s\in S} h_s^{(K)}\right)$$

Lemma MLP, Graph Neural Network (GNN), DeepSet all can **universally approximate reasoning algorithms**.

Equivalent universal expressive power,
BUT big difference in **generalization**

(Johnson et al., 2017a; Weston et al., 2015; Hu et al., 2017; Fleuret et al., 2011; Antol et al., 2015; Battaglia et al., 2016; Watters et al., 2017; Wu et al., 2017, Fragkiadaki et al., 2016; Chang et al., 2017; Saxton et al., 2019; Chang et al., 2019; Santoro et al., 2018; Zhang et al., 2019, …)

Algorithmic alignment

Inductive bias of architectures formally defined.
See paper for mathematical definition.

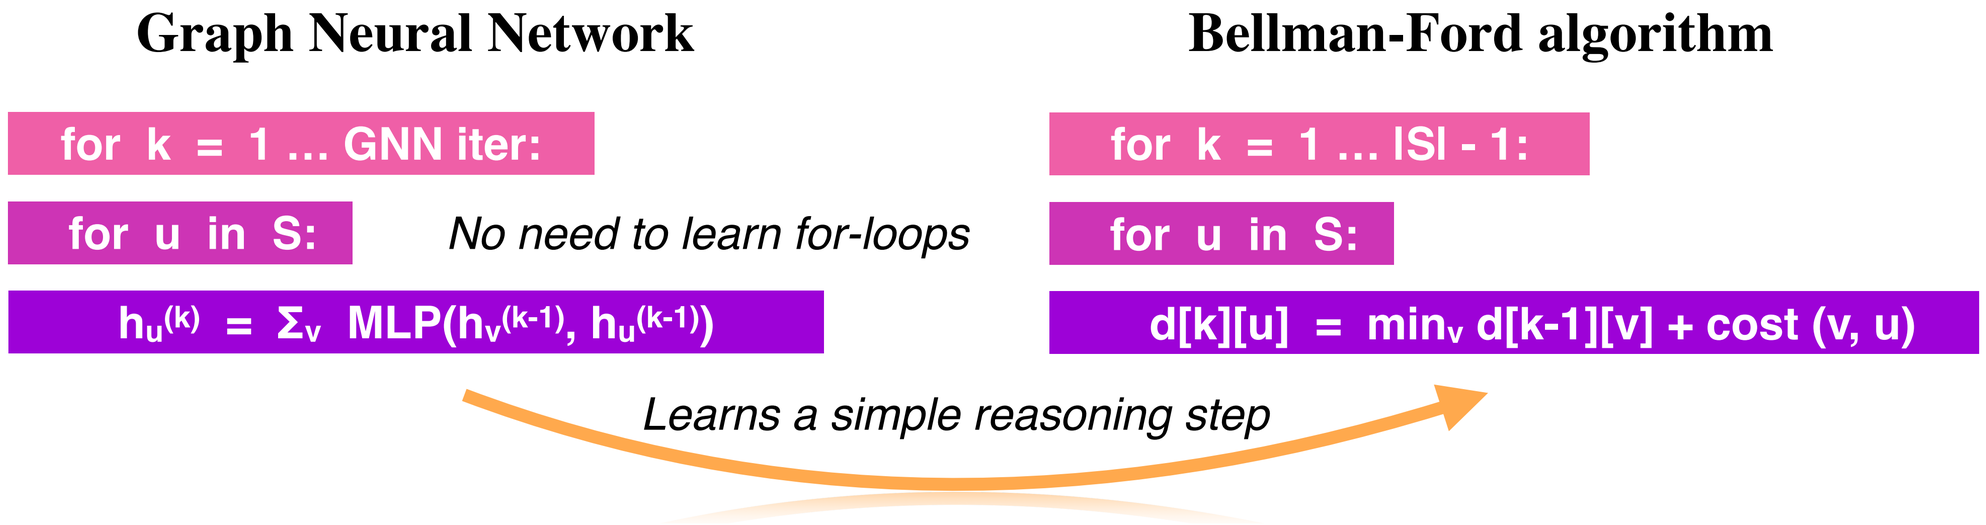

Theorem. Better algorithmic alignment implies better generalization.

DP (GNN) solves many reasoning tasks

More generally: GNNs algorithmically align with Dynamic Programming

$$\text{Answer}[k][i] = \text{DP-Update}(\{\text{Answer}[k-1][j],\ j = 1 \ldots n\})$$

$$h_s^{(k)} = \sum_{t \in S} \text{MLP}_1^{(k)} \left(h_s^{(k-1)}, h_t^{(k-1)} \right)$$

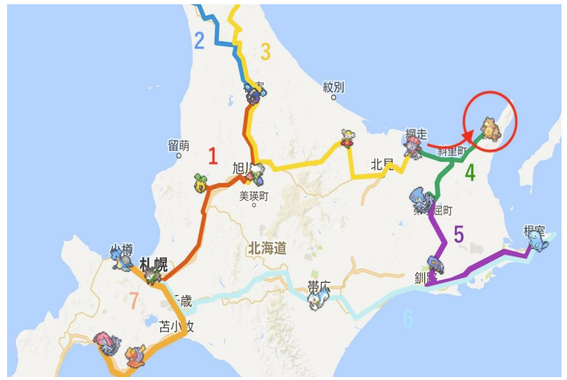

Shortest Path & Graph Algorithms

What is the cost to defeat monster X by following the optimal path?

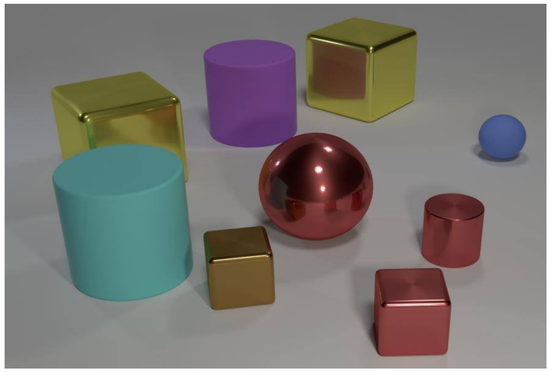

Visual Question Answering

Starting at object X, if each time we jump to the closest object, which object is K jumps away?

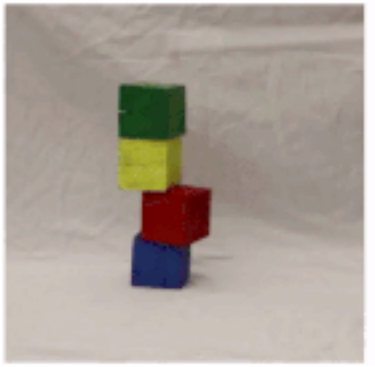

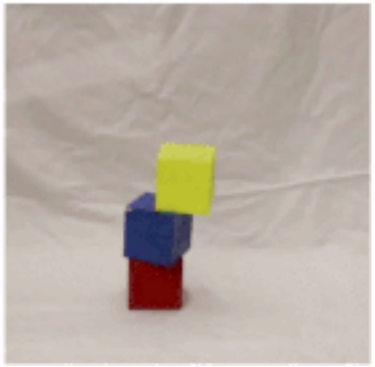

Physical Reasoning

What are the states of the blocks in a second?

Limits of GNN: NP-hard problem (e.g. subset sum)

5 -1 10 -8 7 9 -11 -2 4 3 …

Does any subset of these numbers sum to zero?

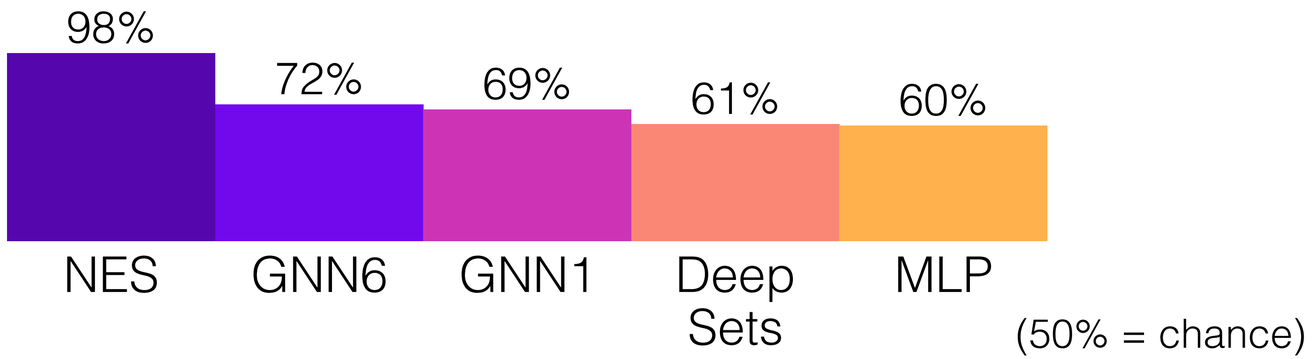

NES (Neural Exhaustive Search) - based on **algo alignment**

$$\mathbf{MLP}_2(\max_{\tau \subseteq S} \mathbf{MLP}_1 \circ \mathbf{LSTM}(X_1, ..., X_{|\tau|} : X_1, ..., X_{|\tau|} \in \tau))$$

y = max$_S$ 1[h(S) = 0], h(S) = $\Sigma_{x \text{ in } S}$ X

What can neural networks reason about?

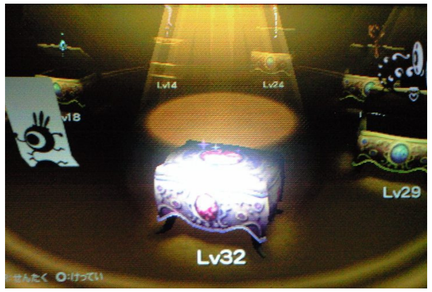

Summary statistics
What is the maximum value difference among treasures?

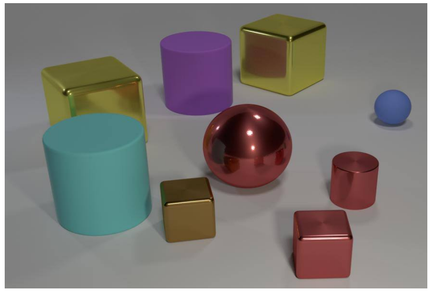

Relational argmax
What are the colors of the furthest pair of objects?

Dynamic programming
What is the cost to defeat monster X by following the optimal path?

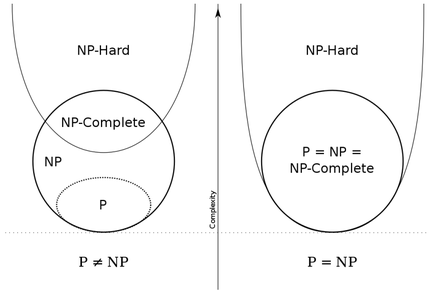

NP-hard problem
Subset sum: Is there a subset that sums to 0?

Graph Neural Network (GNN)	✓	✓	✓	✗
DeepSets	✓	✗	✗	✗
MLP	✗	✗	✗	✗
Neural Exhaustive Search (NES)				✓

Summary of authors

https://openreview.net/forum?id=rJxbJeHFPS

https://github.com/NNReasoning/What-Can-Neural-Networks-Reason-About

Keyulu Xu

超絶弱い主人公

The weakest character

Simon S. Du

深層学習理論の四天王

Bishop of DL theory

Stefanie Jegelka

機械学習の女帝

The Last Jedi of ML

Ken-ichi Kawarabayashi

史上最強のグラフ理論の大魔王

King of Graph Theory

Mozhi Zhang

NLPの魔法使い

Sorcerer of NLP

Jingling Li

AGIを夢見る少女

AGI dreamer