



Closed Loop Neural-Symbolic Learning via Integrating Neural Perception, Grammar Parsing, and Symbolic Reasoning



Qing Li¹



Siyuan Huang¹



Yining Hong²



Yixin Chen¹



Ying Nian Wu¹




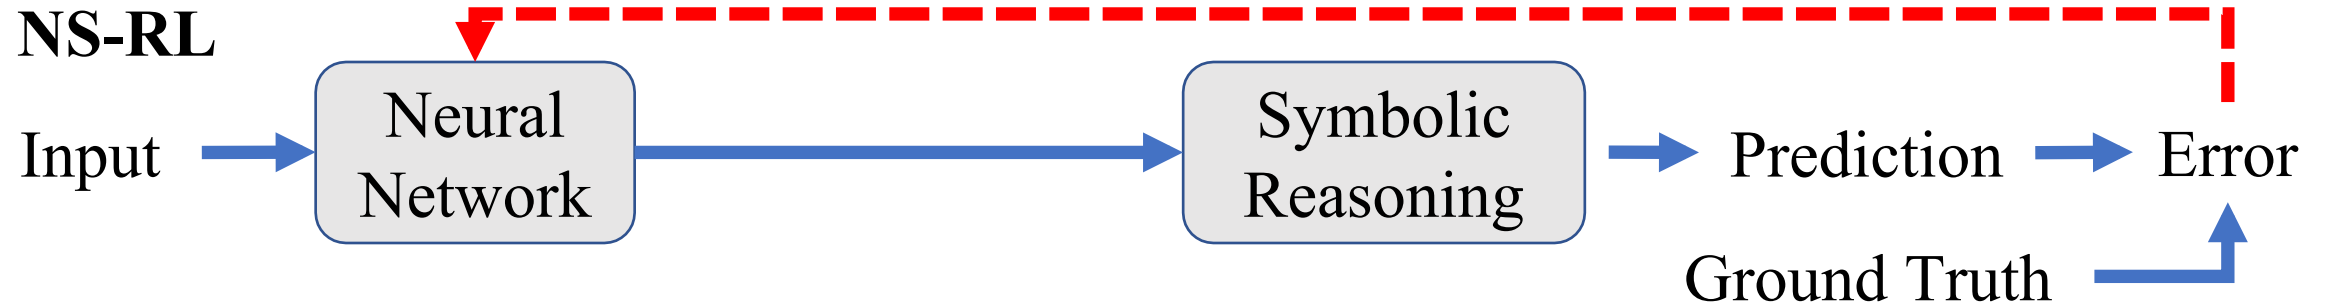
Song-Chun Zhu^{1,2}

¹ Department of Statistics, UCLA



² Department of Computer Science, UCLA

Motivation

NS-RL



How does human do this task?



1. Always generate a valid formula
2. Back-trace the error in the reasoning tree
3. Find the error source and propose a fix
4. Update the perception

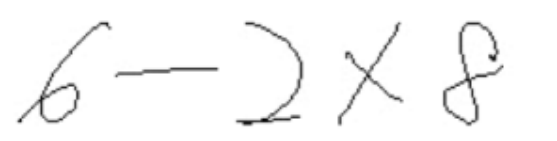
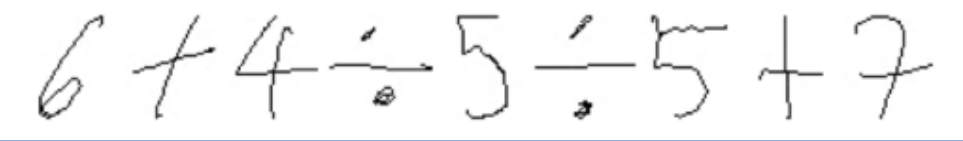
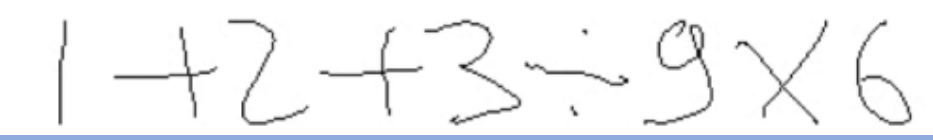
Prior knowledge

Abductive reasoning

Contributions

- **Grammar** to bridge neural network and symbolic reasoning
 - NGS: Neural perception + Grammar parsing + Symbolic reasoning
- **Back-search**
 - Mimic human's ability to learn from failures via abductive reasoning
- A new benchmark **HWF** for neural-symbolic learning
 - Hand-written Formula Recognition with Weak Supervision

Hand-written Formula Recognition (HWF)

Input	Latent	Output
	$6 - 2 \times 8$	-10
	$6 + 4 \div 5 \div 5 + 7$	13.16
	$1 + 2 + 3 \div 9 \times 6$	5

Weakly-supervised!

NGS

6-2x8

Neural Network





Grammar Parsing



Symbolic Reasoning

	'0'	'1'	...	'9'	'+'	'-'	'*'	'/'
6								
-								
...								



Forward Pass (Inference)

- x : the input $6 - 2 \times 8$
- z : the formal representation for x , “ $6 - 2 \times 8$ ”
 - Latent, unobserved
- y : the desired output, e.g., -10

Forward Pass (Inference)

Neural
Perception

$$p_{\theta}(z|x) = \text{softmax}(\phi_{\theta}(z, x)), \\ = \frac{\exp(\phi_{\theta}(z, x))}{\sum_{z'} \exp(\phi_{\theta}(z', x))}$$

Grammar
Parsing

$$\hat{z} = \arg \max_{z \in L(G)} p_{\theta}(z|x)$$

Symbolic
Reasoning

$$\hat{y} = f(\hat{z}; \Delta)$$

Backward Pass (Learning)

- Assumptions
 - Grammar and Symbolic reasoning are perfectly designed by hand, based on our domain knowledge.
 - Only the parameters of neural network need to be learned.
- Gradient descent cannot be applied directly
 - Grammar parsing and symbolic reasoning are non-differentiable.



1-step Back-search (1-BS)

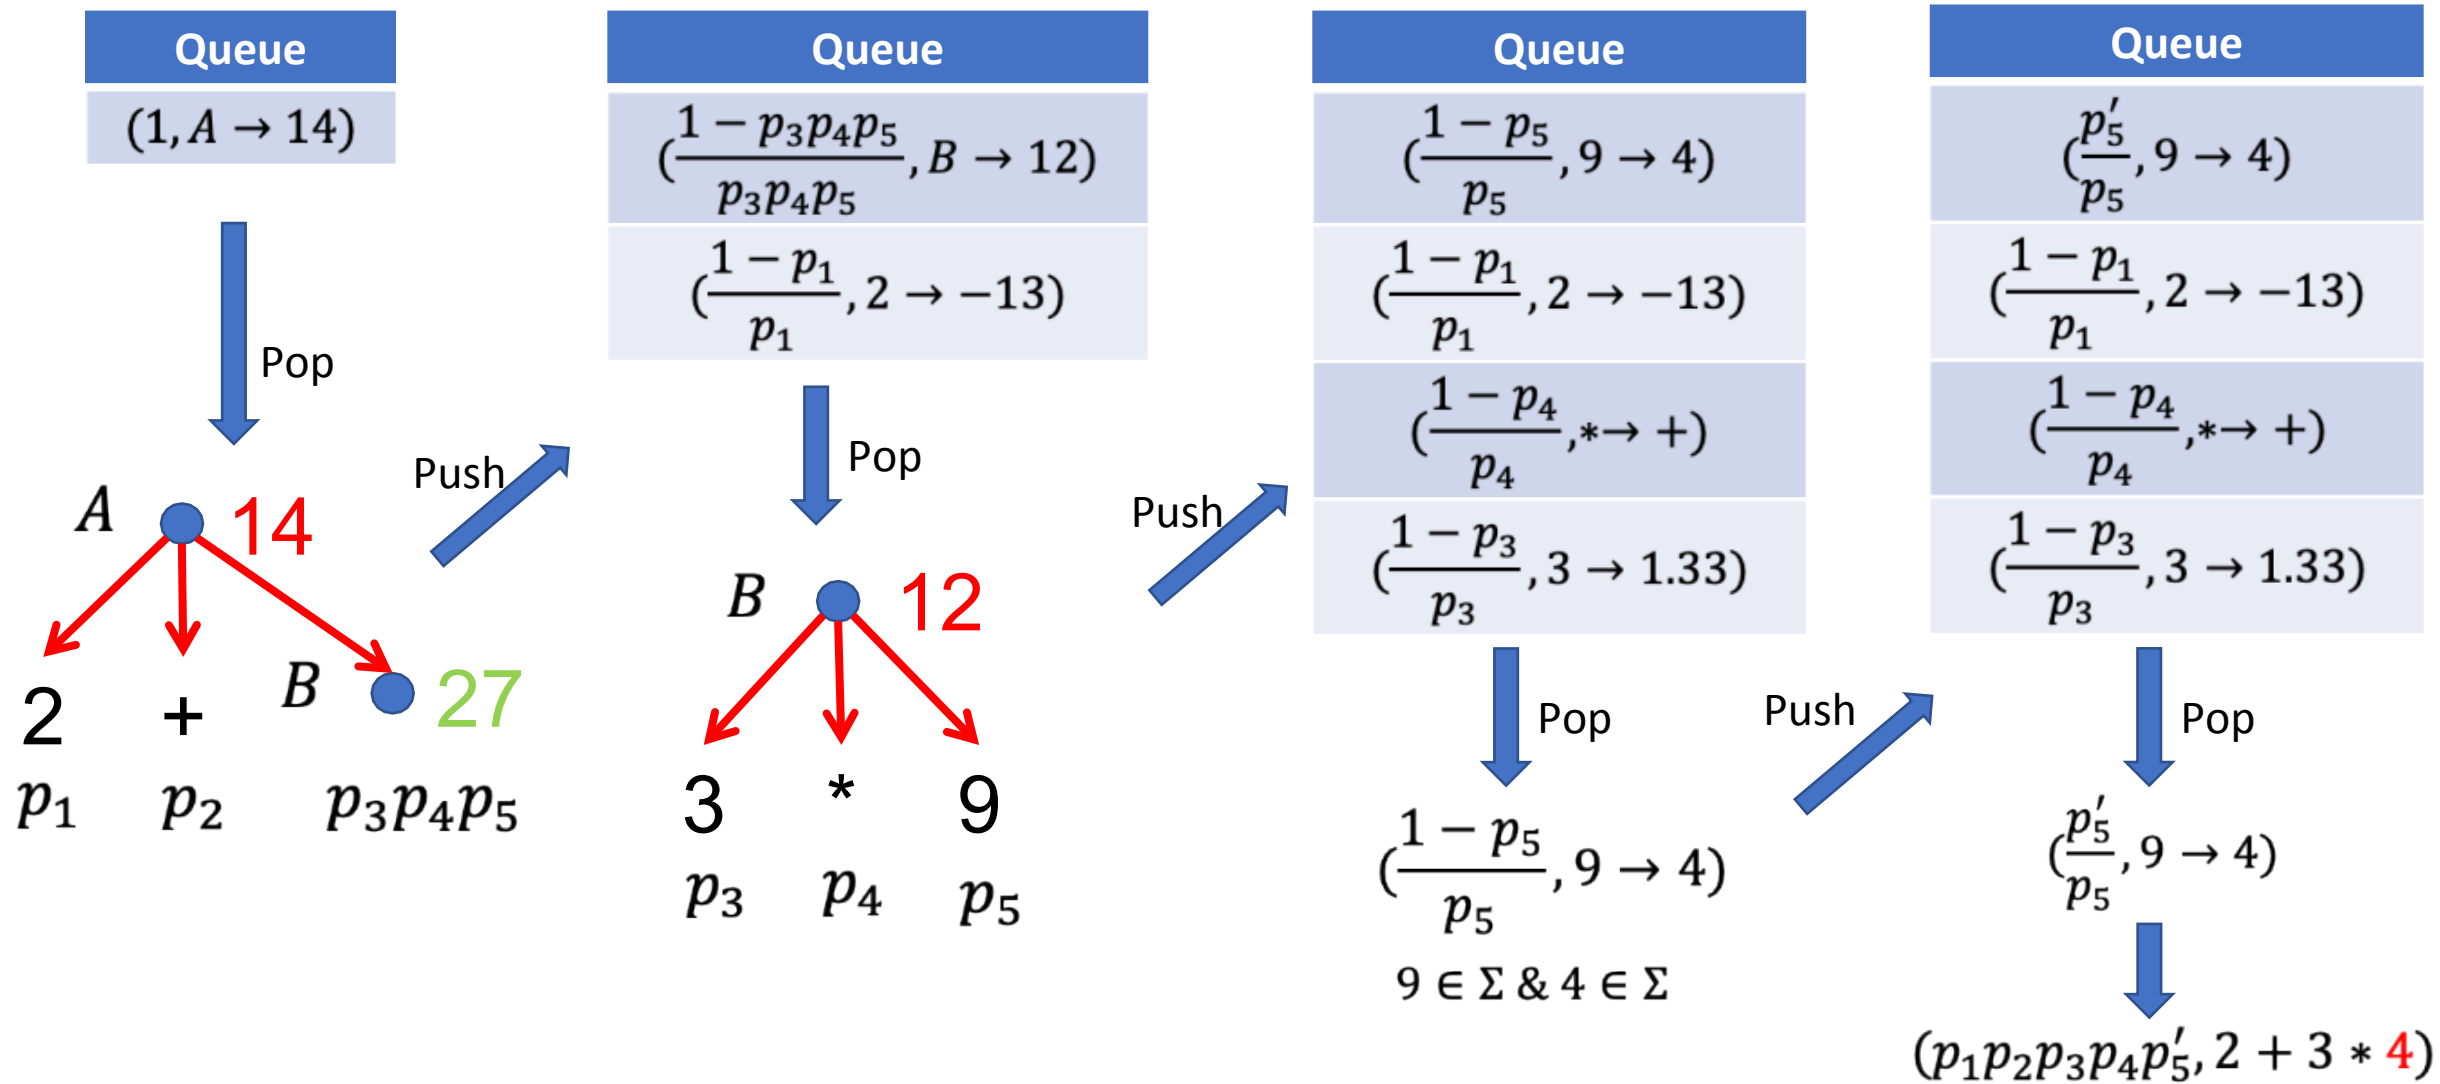
- Top-down search is guided by the bottom-up perception probability
- Dynamic Programming + Priority Queue

Algorithm 1 1-step back-search (1-BS)

```
1: Input:  $\hat{z}, S, y$ 
2:  $q = \text{PriorityQueue}()$ 
3:  $q.\text{push}(S, y, 1)$ 
4: while  $A, \alpha_A, p = q.\text{pop}()$  do
5:   if  $A \in \Sigma$  then
6:      $z^* = \hat{z}(A \rightarrow \alpha_A)$ 
7:     return  $z^*$ 
8:   for  $B \in \text{child}(A)$  do
9:      $\alpha_B = \text{solve}(B, A, \alpha_A | \Delta, G)$ 
10:     $q.\text{push}(B, \alpha_B, p(B \rightarrow \alpha_B))$ 
11: return  $\emptyset$ 
```

A running example for 1-BS





Why can BS be better than RL?

- Learning as Maximum Marginal Likelihood
- REINFORCE as Rejection Sampling
- m-BS as MCMC sampling
 - Metropolis-Hastings sampler

Learning as Maximum Marginal Likelihood

Marginal
likelihood

$$p(y|x) = \sum_z p(y, z|x) = \sum_z p(y|z)p_\theta(z|x)$$

$$\nabla_\theta L(x, y) = \nabla_\theta \log p(y|x)$$

$$= \frac{1}{p(y|x)} \nabla_\theta p(y|x)$$

$$= \sum_z \frac{p(y|z)p_\theta(z|x)}{\sum_{z'} p(y|z')p_\theta(z'|x)} \nabla_\theta \log p_\theta(z|x)$$

$$= \mathbb{E}_{z \sim p(z|x, y)} [\nabla_\theta \log p_\theta(z|x)]$$

$$\approx \nabla_\theta \log p_\theta(\hat{z}|x), \hat{z} \sim p(z|x, y)$$

Monte Carlo
sampling

Posterior distribution

$$\begin{aligned} p(z|x, y) &= \frac{p(y|z)p_\theta(z|x)}{\sum_{z'} p(y|z')p_\theta(z'|x)} \\ &= \begin{cases} 0, & \text{for } z \notin Q \\ \frac{p_\theta(z|x)}{\sum_{z' \in Q} p_\theta(z'|x)}, & \text{for } z \in Q \end{cases} \end{aligned}$$

$$Q = \{z : p(y|z) = 1\} = \{z : f(z; \Delta) = y\}$$

REINFORCE as Rejection Sampling

- Target distribution: $p(z|x, y)$
- Proposal distribution: $p_\theta(z|x)$
- Rejection sampling:
 1. Sample \hat{z} from $p_\theta(z|x)$
 2. If $f(\hat{z}; \Delta) \neq y$ $\dashrightarrow p(\hat{z}|x, y) = 0$
 reject \hat{z}
 else $\dashrightarrow \frac{p(\hat{z}|x, y)}{M p_\theta(\hat{z}|x)} = 1$, where $M = \frac{1}{\sum_{z' \in Q} p_\theta(z'|x)}$
 accept \hat{z}

m-BS as MCMC Sampling

Algorithm 2 *m*-step back-search (*m*-BS)

```
1: Hyperparameters:  $T, \lambda$ 
2: Input:  $\hat{z}, y$ 
3:  $z^{(0)} = \hat{z}$ 
4: for  $t \leftarrow 0$  to  $T - 1$  do
5:    $z^* = \text{1-BS}(z^t, y)$ 
6:   draw  $u \sim \mathcal{U}(0, 1)$ 
7:   if  $u \leq \lambda$  and  $z^* \neq \emptyset$  then
8:      $z^{t+1} = z^*$ 
9:   else
10:     $z^{t+1} = \text{RANDOMWALK}(z^t)$ 
11: return  $z^T$ 
12:
13: function  $\text{RANDOMWALK}(z^t)$ 
14:   sample  $z^* \sim g(\cdot | z^t)$ 
15:   compute acceptance ratio  $a = \min(1, \frac{p_\theta(z^* | x)}{p_\theta(z^t | x)})$ 
16:   draw  $u \sim \mathcal{U}(0, 1)$ 
17:    $z^{t+1} = \begin{cases} z^*, & \text{if } u \leq a \\ z^t, & \text{otherwise.} \end{cases}$ 
```

- m-BS is a Metropolis-Hastings sampler for $p(z|x, y)$

$$\hat{z} \sim p_\theta(z|x) \rightsquigarrow z^* \in Q$$

[Proof in Sec. 3.2.3]

$$g(z_1|z_2) = \text{Poisson}(d(z_1, z_2); \beta)$$

where $d(z_1, z_2)$ denotes the edit distance between z_1, z_2

Experiments

- Hand-written Formula Recognition
- Neural-symbolic VQA

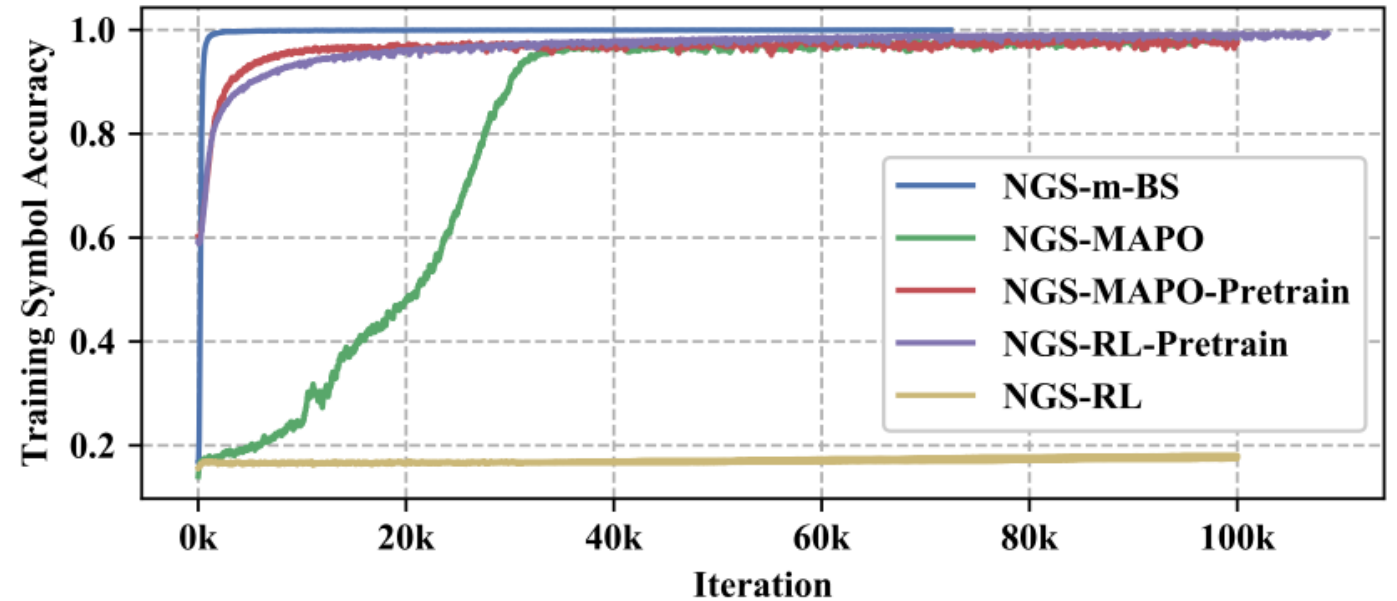
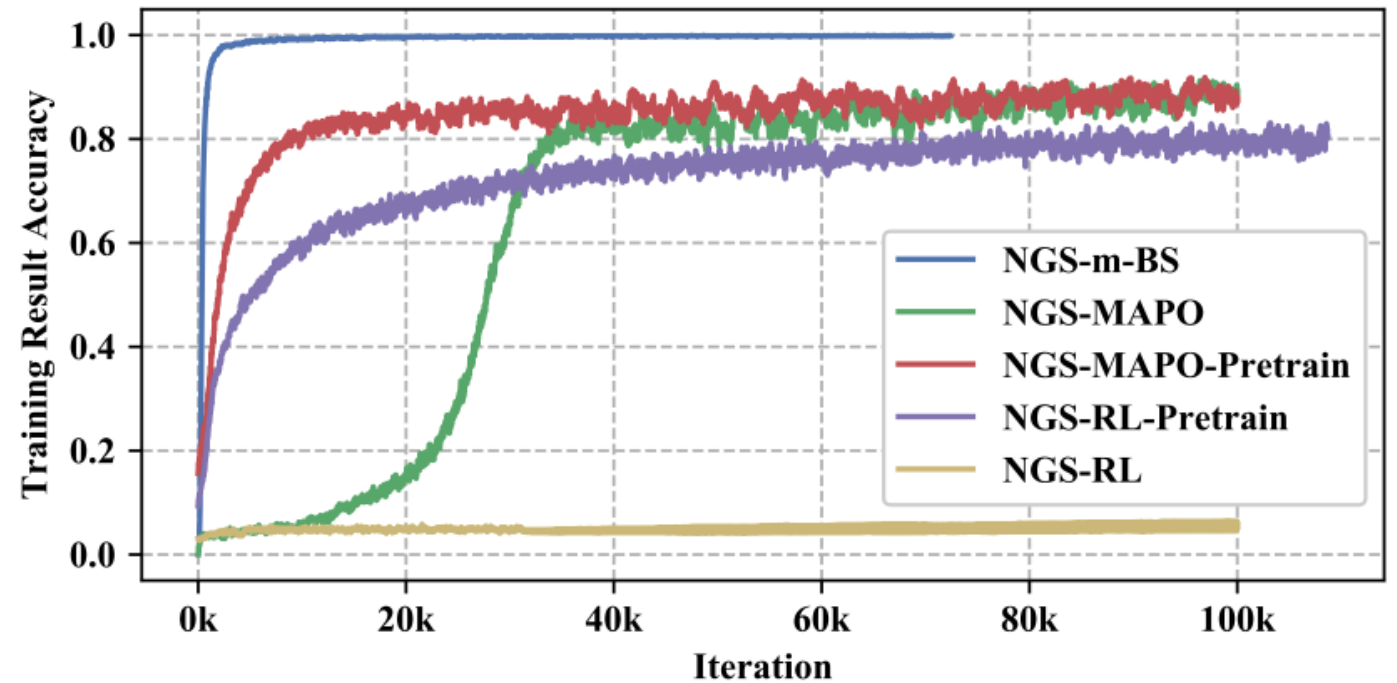
Hand-written Formula Recognition

- Dataset
 - Built from the [CROHME](#) challenge
 - 10k expressions for training, 2k expressions for testing
 - Evaluation
 - Symbol accuracy, Result Accuracy
 - Models
 - NGS-RL, NGS-RL-Pretrained
 - NGS-MAPO*, NGS-MAPO-Pretrained
 - NGS-BS
- *Pretrain NN on a set of fully-supervised data
*Memory-Augmented Policy Optimization [1]



[1] Liang, Chen, et al. "Memory augmented policy optimization for program synthesis and semantic parsing." *Advances in Neural Information Processing Systems*. 2018.

Learning curves

1. NGS-RL fails without pretraining
2. NGS-MAPO works without pretraining but takes a long time to start improving (cold start).
3. Both NGS-RL and NGS-MAPO have noisy learning curves.
4. NGS-BS doesn't suffer from the cold start.
5. NGS-BS converges much faster and the learning curve is smooth.
6. NGS-BS achieves nearly perfect accuracy.





Data efficiency







	25%	50 %	75 %	100%
NGS-RL	0.035	0.036	0.034	0.034
NGS-MAPO	0.051	0.095	0.305	0.717
NGS-RL-Pretrained	0.534	0.621	0.663	0.685
NGS-MAPO-Pretrained	0.687	0.773	0.893	0.956
NGS-BS	0.933	0.957	0.975	0.985

	25%	50 %	75 %	100%
NGS-RL	0.170	0.170	0.170	0.170
NGS-MAPO	0.316	0.481	0.785	0.967
NGS-RL-Pretrained	0.916	0.945	0.959	0.964
NGS-MAPO-Pretrained	0.962	0.983	0.985	0.991
NGS-BS	0.988	0.992	0.995	0.997

Examples



Examples




→ Random Walk

Neural Symbolic VQA


- NS-VQA on CLEVR [1]
- Replace the Seq2Seq question parser with Pointer Network

[1] Yi, Kexin, et al. "Neural-symbolic VQA: Disentangling reasoning from vision and language understanding." NeurIPS 2018.



Examples

Q: Are there any cylinders behind the brown ball? **A:** no




exist	Error
filter_color[brown]	Error
relate[behind]	Error
unique	Error
filter_shape[sphere]	{}
filter_shape[cylinder]	{3}
scene	{1, 2, 3, 4, 5, 6}

1-BS →

exist	no
filter_shape[cylinder]	{}
relate[behind]	{}
unique	4
filter_shape[sphere]	{4}
filter_color[brown]	{4}
scene	{1, 2, 3, 4, 5, 6}

Q: How many cubes are either tiny blue objects or metal things? **A:** 0



count	1
filter_color[blue]	{4}
union	{3, 4}
filter_shape[cube]	{}
filter_size[small]	{2, 4}
scene	{1, 2, 3, 4}
filter_material[metal]	{3, 4}
scene	{1, 2, 3, 4}

1-BS →

count	0
filter_shape[cube]	{}
union	{3, 4}
filter_color[blue]	{4}
filter_size[small]	{2, 4}
scene	{1, 2, 3, 4}
filter_material[metal]	{3, 4}
scene	{1, 2, 3, 4}

Conclusions & Future works

- RL is inefficient for weakly-supervised neural-symbolic learning.
- Back-Search boosts neural-symbolic learning.
- m-BS is a Metropolis-Hastings sampler for the posterior distribution.
- Back-search might be applied to a variety of neural-symbolic tasks, such as semantic parsing, math word problem.
- How to incorporate grammar learning and logic induction is still an open problem.

Thank you!

Project: <https://liqing-ustc.github.io/NGS/>



Code: <https://github.com/liqing-ustc/NGS>

