

# NEURAL PROGRAMMER-INTERPRETERS <sup>1</sup>

Presented by

Fethiye Irmak Doğan

---

<sup>1</sup>Reed, S., and De Freitas, N. (2015). Neural programmer-interpreters. arXiv.

# Central Challenges of AI

- Teaching machine to learn new programs

# Central Challenges of AI

- Teaching machine to learn new programs
- Execute these programs automatically

# Neural Programmer-Interpreters (NPI)

Neural Programmer-Interpreters is a recurrent and compositional neural network that learns how to

- represent a program

# Neural Programmer-Interpreters (NPI)

Neural Programmer-Interpreters is a recurrent and compositional neural network that learns how to

- represent a program
- execute a program (as an **interpreter**)

# Neural Programmer-Interpreters (NPI)

Neural Programmer-Interpreters is a recurrent and compositional neural network that learns how to

- represent a program
- execute a program (as an **interpreter**)
- generate new program embeddings (as a **programmer**)

# Compositional architecture of NPI

**Task agnostic recurrent core** : LSTM based sequence model which is a single core module with the shared parameters across all tasks

# Compositional architecture of NPI

**Task agnostic recurrent core** : LSTM based sequence model which is a single core module with the shared parameters across all tasks

**Persistent key-value program memory** : Learnable key-value memory of program embeddings which provides learning and reusing programs



# Compositional architecture of NPI

**Task agnostic recurrent core** : LSTM based sequence model which is a single core module with the shared parameters across all tasks

**Persistent key-value program memory** : Learnable key-value memory of program embeddings which provides learning and reusing programs

**Domain-specific encoders**: encoder that enables NPI to operate in diverse environments

**Curriculum Learning**<sup>2</sup>: Start small, learn easier aspects of the task or easier subtasks, and then gradually increase the difficulty level.

---

<sup>2</sup>Bengio, Y., Louradour, J., Collobert, R., and Weston, J. (2009). Curriculum learning. In Proceedings of the 26th annual international conference on machine learning, pages 41-48

**Curriculum Learning**<sup>2</sup>: Start small, learn easier aspects of the task or easier subtasks, and then gradually increase the difficulty level.

**Rich Supervision**: Rather than using large number of relatively weak labels, exploit from the fewer fully supervised execution traces

---

<sup>2</sup>Bengio, Y., Louradour, J., Collobert, R., and Weston, J. (2009). Curriculum learning. In Proceedings of the 26th annual international conference on machine learning, pages 41-48

## Related Work

### Dynamically Programmable Networks

- activations of one network become the weights of a second network

# Related Work

## Dynamically Programmable Networks

- activations of one network become the weights of a second network

## Neural Turing Machine

- learning and executing simple programs

# Related Work

## Dynamically Programmable Networks

- activations of one network become the weights of a second network

## Neural Turing Machine

- learning and executing simple programs

## Program Induction

- inducing a program given example input and output pairs

## Novelties of NPI

- being trained on execution traces instead of input and output pairs

## Novelties of NPI

- being trained on execution traces instead of input and output pairs
- incorporating **compositional structure** into the network using a program memory



## Novelties of NPI

- being trained on execution traces instead of input and output pairs
- incorporating **compositional structure** into the network using a program memory
- learning new programs by combining sub-programs

# NPI Core

NPI Core acts as a **router** between programs and there is a single inference core shared by arbitrary programs

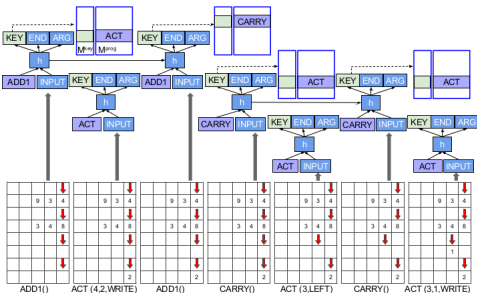


Figure: Example execution trace of single-digit addition

# NPI Core




Figure: Example execution trace of single-digit addition

NPI Core is conditioned on

- current state observations:

# NPI Core




Figure: Example execution trace of single-digit addition

NPI Core is conditioned on

- current state observations:
  - learnable program embedding, program arguments, feature representation of the environment

# NPI Core




Figure: Example execution trace of single-digit addition

NPI Core is conditioned on

- current state observations:
  - learnable program embedding, program arguments, feature representation of the environment
- previous hidden unit states

# NPI Core




Figure: Example execution trace of single-digit addition

NPI Core outputs

- key indicating what program to call next

# NPI Core




Figure: Example execution trace of single-digit addition

## NPI Core outputs

- key indicating what program to call next
- probability of ending the current program

# NPI Core




Figure: Example execution trace of single-digit addition

## NPI Core outputs

- key indicating what program to call next
- probability of ending the current program
- argument for the following program (passed by reference or value)



# Program Embedding Memory

Different programs correspond to different embeddings stored in a persistent memory




Figure: Example execution trace of single-digit addition

# Feed-Forward steps of program inference

$e_t$  : environment observation at time  $t$     $a_t$  : current program argument    $s_t$  : state encoding  
 $p_t$  : program embedding    $h_{t-1}$  : previous hidden unit    $c_{t-1}$  : previous cell unit  
 $r_t$  : end of program probability    $k_t$  : program key embedding    $a_t$  : output arguments at time  $t$   
 $f_{enc}$  : domain specific encoder    $f_{lstm}$  : LSTM mapping  
 $f_{end}$  : probability of finishing the program    $f_{prog}$  : key embedding for next program    $f_{arg}$  : arguments to next program

$$s_t = f_{enc}(e_t, a_t)$$

# Feed-Forward steps of program inference

$e_t$  : environment observation at time  $t$     $a_t$  : current program argument    $s_t$  : state encoding  
 $p_t$  : program embedding    $h_{t-1}$  : previous hidden unit    $c_{t-1}$  : previous cell unit  
 $r_t$  : end of program probability    $k_t$  : program key embedding    $a_t$  : output arguments at time  $t$   
 $f_{enc}$  : domain specific encoder    $f_{lstm}$  : LSTM mapping  
 $f_{end}$  : probability of finishing the program    $f_{prog}$  : key embedding for next program    $f_{arg}$  : arguments to next program

$$s_t = f_{enc}(e_t, a_t)$$

$$h_t = f_{lstm}(s_t, p_t, h_{t-1})$$

# Feed-Forward steps of program inference

$e_t$  : environment observation at time  $t$     $a_t$  : current program argument    $s_t$  : state encoding  
 $p_t$  : program embedding    $h_{t-1}$  : previous hidden unit    $c_{t-1}$  : previous cell unit  
 $r_t$  : end of program probability    $k_t$  : program key embedding    $a_t$  : output arguments at time  $t$   
 $f_{enc}$  : domain specific encoder    $f_{lstm}$  : LSTM mapping  
 $f_{end}$  : probability of finishing the program    $f_{prog}$  : key embedding for next program    $f_{arg}$  : arguments to next program

$$s_t = f_{enc}(e_t, a_t)$$

$$h_t = f_{lstm}(s_t, p_t, h_{t-1})$$

$$r_t = f_{end}(h_t), k_t = f_{prog}(h_t), a_{t+1} = f_{arg}(h_t)$$

# Program Embedding

$k_t$  : program key embedding     $i$  : program ID     $p_{t+1}$  : next program embedding  
 $M^{\text{key}}$  : key embeddings which stores all the program keys     $M^{\text{prog}}$  : program embeddings

$$i^* = \arg \max_{i=1..N} (M_{i,:}^{\text{key}})^T k_t \quad , \quad p_{t+1} = M_{i^*,:}^{\text{prog}}$$

# Environmental State

$e_t$  : environment observation at time  $t$     $p_t$  : program embedding    $a_t$  : output arguments at time  $t$   
 $f_{env}$  : domain specific transition mapping    $e_{t+1}$  : next environmental state

$$e_{t+1} \sim f_{env}(e_t, p_t, a_t)$$

# Inference Algorithm

---

**Algorithm 1** Neural programming inference

---

- 1: **Inputs:** Environment observation  $e$ , program id  $i$ , arguments  $a$ , stop threshold  $\alpha$
  - 2: **function** RUN( $i, a$ )
  - 3:      $h \leftarrow \mathbf{0}, r \leftarrow 0, p \leftarrow M_{i,:}^{\text{prog}}$  ▷ Init LSTM and return probability.
  - 4:     **while**  $r < \alpha$  **do**
  - 5:          $s \leftarrow f_{enc}(e, a), h \leftarrow f_{lstm}(s, p, h)$  ▷ Feed-forward NPI one step.
  - 6:          $r \leftarrow f_{end}(h), k \leftarrow f_{prog}(h), a_2 \leftarrow f_{arg}(h)$
  - 7:          $i_2 \leftarrow \arg \max_{j=1..N} (M_{j,:}^{\text{key}})^T k$  ▷ Decide the next program to run.
  - 8:         **if**  $i == \text{ACT}$  **then**  $e \leftarrow f_{env}(e, p, a)$  ▷ Update the environment based on ACT.
  - 9:         **else** RUN( $i_2, a_2$ ) ▷ Run subprogram  $i_2$  with arguments  $a_2$
-

# Inference Algorithm

---

**Algorithm 1** Neural programming inference

---

```
1: Inputs: Environment observation  $e$ , program id  $i$ , arguments  $a$ , stop threshold  $\alpha$ 
2: function RUN( $i, a$ )
3:    $h \leftarrow \mathbf{0}, r \leftarrow 0, p \leftarrow M_{i,:}^{\text{prog}}$  ▷ Init LSTM and return probability.
4:   while  $r < \alpha$  do
5:      $s \leftarrow f_{\text{enc}}(e, a), h \leftarrow f_{\text{lstn}}(s, p, h)$  ▷ Feed-forward NPI one step.
6:      $r \leftarrow f_{\text{end}}(h), k \leftarrow f_{\text{prog}}(h), a_2 \leftarrow f_{\text{arg}}(h)$ 
7:      $i_2 \leftarrow \arg \max_{j=1..N} (M_{j,:}^{\text{key}})^T k$  ▷ Decide the next program to run.
8:     if  $i == \text{ACT}$  then  $e \leftarrow f_{\text{env}}(e, p, a)$  ▷ Update the environment based on ACT.
9:     else RUN( $i_2, a_2$ ) ▷ Run subprogram  $i_2$  with arguments  $a_2$ 
```

---

- actions are encapsulated into ACT program shared across tasks and indicated by the NPI-generated arguments  $a_t$



# Inference Algorithm

---

**Algorithm 1** Neural programming inference

---

```
1: Inputs: Environment observation  $e$ , program id  $i$ , arguments  $a$ , stop threshold  $\alpha$ 
2: function RUN( $i, a$ )
3:    $h \leftarrow \mathbf{0}, r \leftarrow 0, p \leftarrow M_{i,:}^{\text{prog}}$  ▷ Init LSTM and return probability.
4:   while  $r < \alpha$  do
5:      $s \leftarrow f_{enc}(e, a), h \leftarrow f_{lstm}(s, p, h)$  ▷ Feed-forward NPI one step.
6:      $r \leftarrow f_{end}(h), k \leftarrow f_{prog}(h), a_2 \leftarrow f_{arg}(h)$ 
7:      $i_2 \leftarrow \arg \max_{j=1..N} (M_{j,:}^{\text{key}})^T k$  ▷ Decide the next program to run.
8:     if  $i == \text{ACT}$  then  $e \leftarrow f_{env}(e, p, a)$  ▷ Update the environment based on ACT.
9:     else RUN( $i_2, a_2$ ) ▷ Run subprogram  $i_2$  with arguments  $a_2$ 
```

---

- actions are encapsulated into ACT program shared across tasks and indicated by the NPI-generated arguments  $a_t$
- core module is completely agnostic to the data modality used in the state encoding

# Training

$\epsilon_t^{inp} : \{e_t, i_t, a_t\}$  and  $\epsilon_t^{out} : \{i_{t+1}, a_{t+1}, r_t\}$  are the execution traces

# Training

$\epsilon_t^{inp} : \{e_t, i_t, a_t\}$  and  $\epsilon_t^{out} : \{i_{t+1}, a_{t+1}, r_t\}$  are the execution traces

$i_t$  and  $i_{t+1}$  are program IDs and row indices in  $M^{key}$   $M^{prog}$  of the programs to run at time  $t$  and  $t+1$

# Training

$\varepsilon_t^{inp} : \{e_t, i_t, a_t\}$  and  $\varepsilon_t^{out} : \{i_{t+1}, a_{t+1}, r_t\}$  are the execution traces

$i_t$  and  $i_{t+1}$  are program IDs and row indices in  $M^{key}$   $M^{prog}$  of the programs to run at time  $t$  and  $t+1$

$$\theta^* = \arg \max_{\theta} \sum_{(\xi^{inp}, \xi^{out})} \log P(\xi^{out} | \xi^{inp}; \theta)$$

# Training

$\varepsilon_t^{inp} : \{e_t, i_t, a_t\}$  and  $\varepsilon_t^{out} : \{i_{t+1}, a_{t+1}, r_t\}$  are the execution traces

$i_t$  and  $i_{t+1}$  are program IDs and row indices in  $M^{key}$   $M^{prog}$  of the programs to run at time  $t$  and  $t+1$

$$\theta^* = \arg \max_{\theta} \sum_{(\xi^{inp}, \xi^{out})} \log P(\xi^{out} | \xi^{inp}; \theta)$$

since traces are **variable length** above equation can be written as:

# Training

$\varepsilon_t^{inp} : \{e_t, i_t, a_t\}$  and  $\varepsilon_t^{out} : \{i_{t+1}, a_{t+1}, r_t\}$  are the execution traces

$i_t$  and  $i_{t+1}$  are program IDs and row indices in  $M^{key}$   $M^{prog}$  of the programs to run at time  $t$  and  $t+1$

$$\theta^* = \arg \max_{\theta} \sum_{(\xi^{inp}, \xi^{out})} \log P(\xi^{out} | \xi^{inp}; \theta)$$

since traces are **variable length** above equation can be written as:

$$\log P(\xi_{out} | \xi_{inp}; \theta) = \sum_{t=1}^T \log P(\xi_t^{out} | \xi_1^{inp}, \dots, \xi_t^{inp}; \theta)$$

# Training

$\varepsilon_t^{inp} : \{e_t, i_t, a_t\}$  and  $\varepsilon_t^{out} : \{i_{t+1}, a_{t+1}, r_t\}$  are the execution traces

$i_t$  and  $i_{t+1}$  are program IDs and row indices in  $M^{key}$   $M^{prog}$  of the programs to run at time  $t$  and  $t+1$

$$\theta^* = \arg \max_{\theta} \sum_{(\xi^{inp}, \xi^{out})} \log P(\xi^{out} | \xi^{inp}; \theta)$$

since traces are **variable length** above equation can be written as:

$$\log P(\xi_{out} | \xi_{inp}; \theta) = \sum_{t=1}^T \log P(\xi_t^{out} | \xi_1^{inp}, \dots, \xi_t^{inp}; \theta)$$

since **hidden unit activations** are capable of capturing temporal dependencies, right hand side can be written as:

# Training

$\varepsilon_t^{inp} : \{e_t, i_t, a_t\}$  and  $\varepsilon_t^{out} : \{i_{t+1}, a_{t+1}, r_t\}$  are the execution traces

$i_t$  and  $i_{t+1}$  are program IDs and row indices in  $M^{key}$   $M^{prog}$  of the programs to run at time  $t$  and  $t+1$

$$\theta^* = \arg \max_{\theta} \sum_{(\xi^{inp}, \xi^{out})} \log P(\xi^{out} | \xi^{inp}; \theta)$$

since traces are **variable length** above equation can be written as:

$$\log P(\xi_{out} | \xi_{inp}; \theta) = \sum_{t=1}^T \log P(\xi_t^{out} | \xi_1^{inp}, \dots, \xi_t^{inp}; \theta)$$

since **hidden unit activations** are capable of capturing temporal dependencies, right hand side can be written as:

$$\log P(\xi_t^{out} | \xi_1^{inp}, \dots, \xi_t^{inp}) = \log P(i_{t+1} | h_t) + \log P(a_{t+1} | h_t) + \log P(r_t | h_t)$$



# Training

- **Adaptive curriculum** : sample frequency of a program is determined by model's current prediction error in that program

# Training

- **Adaptive curriculum** : sample frequency of a program is determined by model's current prediction error in that program
  - forces the model to focus on learning the program worst in execution

# Training

- **Adaptive curriculum** : sample frequency of a program is determined by model's current prediction error in that program
  - forces the model to focus on learning the program worst in execution
- Memory advantage thanks to **parallel execution** in sub-programs

# Addition

- **Task:** read in the digits of two base-10 numbers and produce the digits of the answer

---

<sup>3</sup><https://en.wikipedia.org/wiki/One-hot>

# Addition

- **Task:** read in the digits of two base-10 numbers and produce the digits of the answer
- Four pointers: one for each of the two input numbers, one for the carry, and another to write the output

---

<sup>3</sup><https://en.wikipedia.org/wiki/One-hot>

# Addition

- **Task:** read in the digits of two base-10 numbers and produce the digits of the answer
- Four pointers: one for each of the two input numbers, one for the carry, and another to write the output
- Model sees the current values at each pointer locations as 1-of-K encodings<sup>3</sup> (K=10)

---

<sup>3</sup><https://en.wikipedia.org/wiki/One-hot>

# Addition

- **Task:** read in the digits of two base-10 numbers and produce the digits of the answer
- Four pointers: one for each of the two input numbers, one for the carry, and another to write the output
- Model sees the current values at each pointer locations as 1-of-K encodings<sup>3</sup> (K=10)

$$f_{enc}(Q, i_1, i_2, i_3, i_4, a_t) = MLP([Q(1, i_1), Q(2, i_2), Q(3, i_3), Q(4, i_4), a_t(1), a_t(2), a_t(3)])$$

$Q \in R^{4 \times N \times K}$  is the scratch pad, first dimension of Q corresponds to scratch pad rows, N is the number of columns (digits) and K is the one-hot encoding dimension

---


<sup>3</sup><https://en.wikipedia.org/wiki/One-hot>

# Addition

Program	Descriptions	Calls
ADD	Perform multi-digit addition	ADD1, LSHIFT
ADD1	Perform single-digit addition	ACT, CARRY
CARRY	Mark a 1 in the carry row one unit left	ACT
LSHIFT	Shift a specified pointer one step left	ACT
RSHIFT	Shift a specified pointer one step right	ACT
ACT	Move a pointer or write to the scratch pad	-



(a) Example scratch pad and pointers used for computing “ $96 + 125 = 221$ ”. Carry step is being implemented.



(b) Actual trace of addition program generated by our model on the problem shown to the left. Note that we substituted the ACT calls in the trace with more human-readable steps.

Figure: Illustration of the addition environment



# Sorting

- **Task:** comparing each pair of adjacent items and swaps them if they are in the wrong order (Bubble Sort <sup>4</sup> )

---

<sup>4</sup>[https://en.wikipedia.org/wiki/Bubble\\_sort](https://en.wikipedia.org/wiki/Bubble_sort)

# Sorting

- **Task:** comparing each pair of adjacent items and swaps them if they are in the wrong order (Bubble Sort <sup>4</sup> )

$$f_{enc}(Q, i_1, i_2, a_t) = MLP([Q(1, i_1), Q(1, i_2), a_t(1), a_t(2), a_t(3)])$$


$Q \in \mathbb{R}^{1 \times N \times K}$  is the scratch pad, N is the array length and K is the array entry embedding dimension

---


<sup>4</sup>[https://en.wikipedia.org/wiki/Bubble\\_sort](https://en.wikipedia.org/wiki/Bubble_sort)

# Sorting

BUBBLESORT	Perform bubble sort (ascending order)	BUBBLE, RESET
BUBBLE	Perform one sweep of pointers left to right	ACT, BSTEP
RESET	Move both pointers all the way left	LSHIFT
BSTEP	Conditionally swap and advance pointers	COMPSPWAP, RSHIFT
COMPSPWAP	Conditionally swap two elements	ACT
LSHIFT	Shift a specified pointer one step left	ACT
RSHIFT	Shift a specified pointer one step right	ACT
ACT	Swap two values at pointer locations or move a pointer	-



(a) Example scratch pad and pointers used for sorting. Several steps of the BUBBLE subprogram are shown.



(b) Excerpt from the trace of the learned bubblesort program.

Figure: Illustration of the sorting environment

## Canonicalizing 3D Models

- **Task:** learn a visual program that canonicalizes the model with respect to its pose

## Canonicalizing 3D Models

- **Task:** learn a visual program that canonicalizes the model with respect to its pose
- Nontrivial problem: different starting positions and different car models

# Canonicalizing 3D Models

- **Task:** learn a visual program that canonicalizes the model with respect to its pose
- Nontrivial problem: different starting positions and different car models

$$f_{enc}(Q, x, i_1, i_2, a_t) = MLP([Q(1, i_1), Q(2, i_2), f_{CNN}(x), a_t(1), a_t(2), a_t(3)])$$

$x \in R^{H \times W \times 3}$  is the car rendering and  $Q \in R^{2 \times 1 \times K}$  is the scratch pad, first dimension of  $Q$  corresponds to  $i_1, i_2$  (fixed at 1) which are the pointer locations of the azimuth and elevation and  $K(=24)$  is the one-hot encoding dimension of pose coordinates


# Canonicalizing 3D Models

GOTO	Change 3D car pose to match the target	HGOTO, VGOTO
HGOTO	Move horizontally to the target angle	LGOTO, RGOTO
LGOTO	Move left to match the target angle	ACT
RGOTO	Move right to match the target angle	ACT
VGOTO	Move vertically to the target elevation	UGOTO, DGOTO
UGOTO	Move up to match the target elevation	ACT
DGOTO	Move down to match the target elevation	ACT
ACT	Move camera 15° up, down, left or right	-


GOTO 1 2  
 HGOTO  
 RGOTO  
 ACT (RIGHT)  
 VGOTO  
 UGOTO  
 ACT (UP)



GOTO 1 2  
 HGOTO  
 RGOTO  
 ACT (RIGHT)  
 ACT (RIGHT)  
 ACT (RIGHT)  
 VGOTO  
 DGOTO  
 ACT (DOWN)  
 ACT (DOWN)



GOTO 1 2  
 HGOTO  
 LGOTO  
 ACT (LEFT)  
 ACT (LEFT)  
 ACT (LEFT)  
 ACT (LEFT)  
 ACT (LEFT)  
 VGOTO  
 UGOTO  
 ACT (UP)



GOTO 1 2  
 HGOTO  
 LGOTO  
 ACT (LEFT)  
 VGOTO  
 DGOTO  
 ACT (DOWN)



Figure: canonicalization of several different test set cars

## Sample Complexity on Bubble Sort Problem

- Memory requirements is reduced from  $O(n^2)$  to  $O(n)$  thanks to compositional structure of the model



## Sample Complexity on Bubble Sort Problem

- Memory requirements is reduced from  $O(n^2)$  to  $O(n)$  thanks to compositional structure of the model
- Number of required training samples are also reduced:




Figure: Test accuracy by the varying sample complexity

## Generalization on Bubble Sort Problem

- Training the model with variable-sized input (single-digit numbers from length 2 to length 20)

## Generalization on Bubble Sort Problem

- Training the model with variable-sized input (single-digit numbers from length 2 to length 20)
- Adding a third pointer that acts as a counter to handle variable-sized inputs

## Generalization on Bubble Sort Problem

- Training the model with variable-sized input (single-digit numbers from length 2 to length 20)
- Adding a third pointer that acts as a counter to handle variable-sized inputs
- Checking the success of the model on the inputs of previously unseen size to check how much the problem is learned

## Generalization on Bubble Sort Problem

- Training the model with variable-sized input (single-digit numbers from length 2 to length 20)
- Adding a third pointer that acts as a counter to handle variable-sized inputs
- Checking the success of the model on the inputs of previously unseen size to check how much the problem is learned




Figure: Strong vs. weak generalization

## Generalization on 3D Canonicalization Problem

- NPI is able to canonicalize cars of varying appearance from **multiple starting positions**

## Generalization on 3D Canonicalization Problem

- NPI is able to canonicalize cars of varying appearance from **multiple starting positions**
- NPI can generalize to car appearances **not encountered in the training**

# Generalization on 3D Canonicalization Problem

- NPI is able to canonicalize cars of varying appearance from **multiple starting positions**
- NPI can generalize to car appearances **not encountered in the training**




Figure: canonicalization of several different test set cars



## Learning New Programs with a Fixed Core

- Fixing all the weights of core routing module

## Learning New Programs with a Fixed Core

- Fixing all the weights of core routing module
- Only updating memory slots of the new programs

## Prevent Existing Programs from Calling Subsequently Added Programs

- Looking back at the training data for known programs

## Prevent Existing Programs from Calling Subsequently Added Programs

- Looking back at the training data for known programs
- Allowing addition of new programs

# Solving Multiple Tasks with a Single Network

Task	Single	Multi	+ Max
Addition	100.0	97.0	97.0
Sorting	100.0	100.0	100.0
Canon. seen car	89.5	91.4	91.4
Canon. unseen	88.7	89.9	89.9
Maximum	-	-	100.0

Per-sequence % accuracy

- NPI learns MAX perfectly **without forgetting the other tasks**

# Solving Multiple Tasks with a Single Network

Task	Single	Multi	+ Max
Addition	100.0	97.0	97.0
Sorting	100.0	100.0	100.0
Canon. seen car	89.5	91.4	91.4
Canon. unseen	88.7	89.9	89.9
Maximum	-	-	100.0

Per-sequence % accuracy

- NPI learns MAX perfectly **without forgetting the other tasks**
- One multi-task NPI can learn all three programs with comparable accuracy compared to each single-task NPI

# Conclusion

## Neural Programmer-Interpreters (NPI)

- learns several programs by using a **single core model**

# Conclusion

## Neural Programmer-Interpreters (NPI)

- learns several programs by using a **single core model**
- reduces sample complexity



# Conclusion

## Neural Programmer-Interpreters (NPI)

- learns several programs by using a **single core model**
- reduces sample complexity
- provides strong generalization

# Conclusion

## Neural Programmer-Interpreters (NPI)

- learns several programs by using a **single core model**
- reduces sample complexity
- provides strong generalization
- works for dissimilar environments

# Conclusion

## Neural Programmer-Interpreters (NPI)

- learns several programs by using a **single core model**
- reduces sample complexity
- provides strong generalization
- works for dissimilar environments
- learns new programs without forgetting already learned ones

# Thank you!

