MLSys | 2023

Towards Cognitive AI Systems: A Survey and Prospective on Neuro-Symbolic AI

Zishen Wan, Che-Kai Liu*, Hanchen Yang*, Chaojian Li*, Haoran You*, Yonggan Fu, Cheng Wan, Tushar Krishna, Yingyan (Celine) Lin, Arijit Raychowdhury

Georgia Institute of Technology, Atlanta, GA

State of AI / Landscape

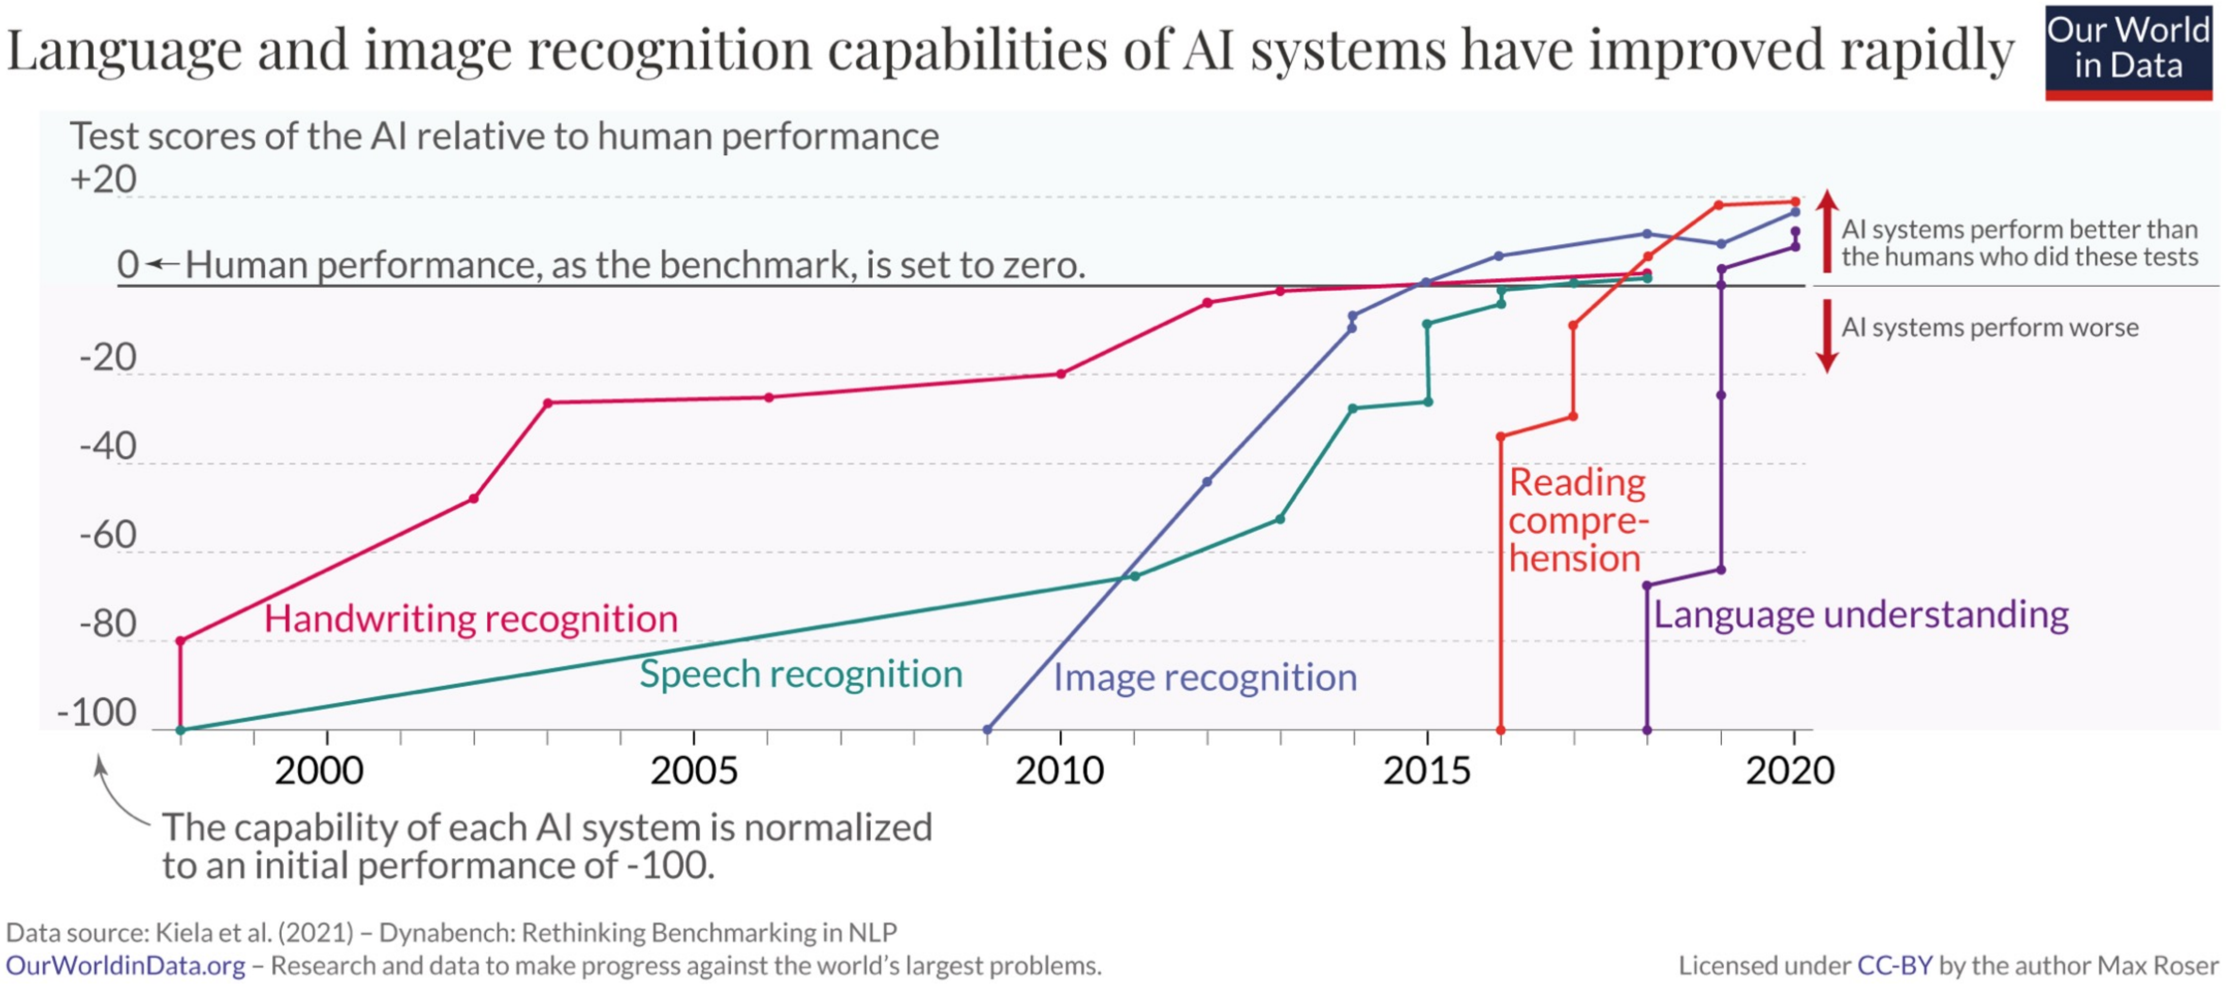

AI Challenges

AI Challenges

- Unsustainable compute trajectory

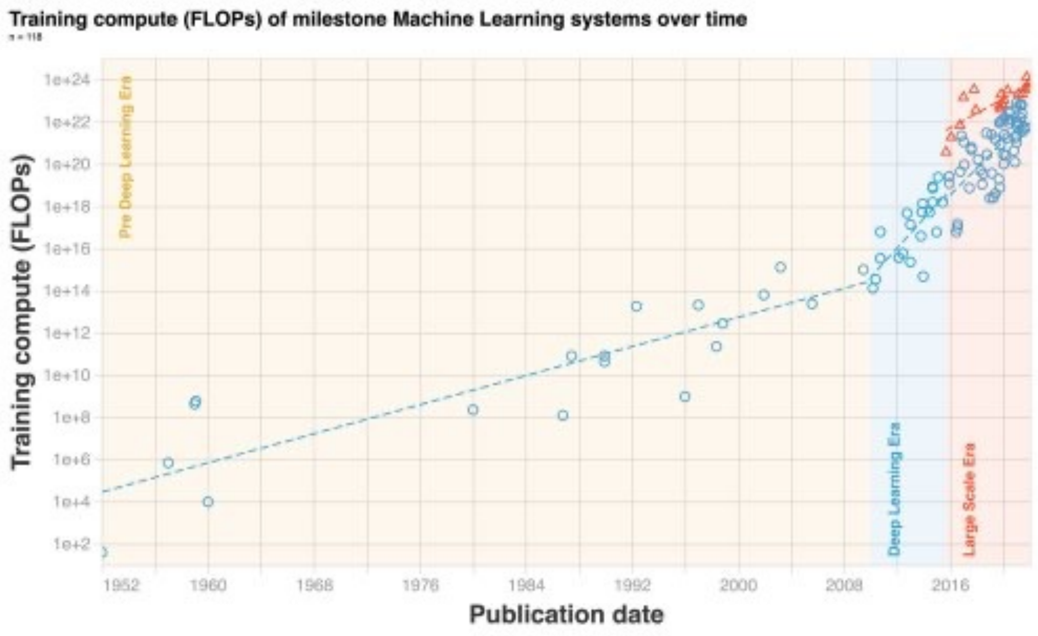

AI Challenges

- Unsustainable compute trajectory
- Lack of explainability and transparency

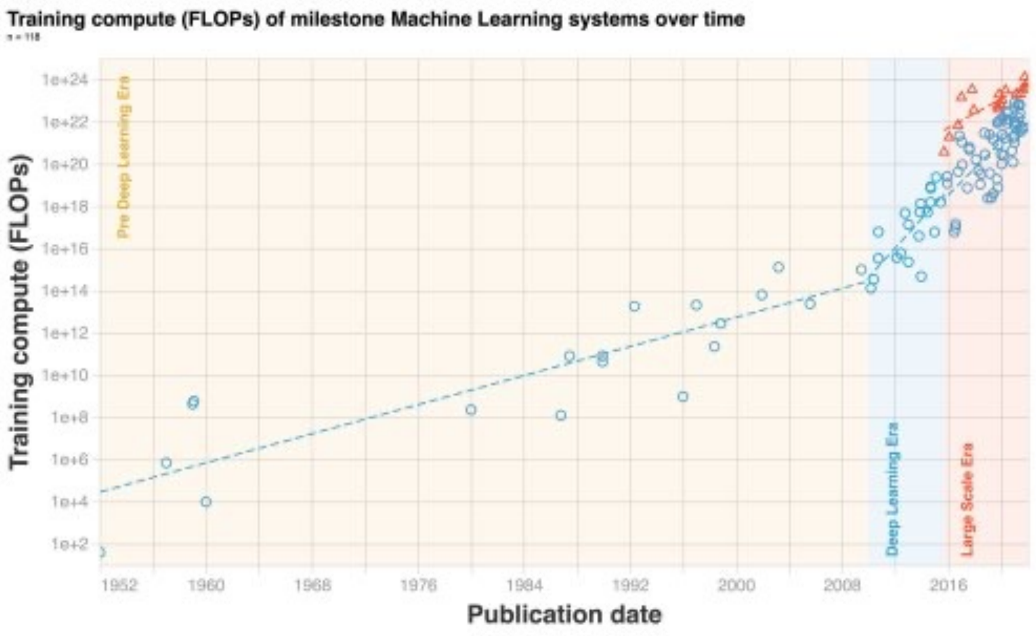

AI Challenges

- Unsustainable compute trajectory
- Lack of explainability and transparency
- Lack of robustness and reliability

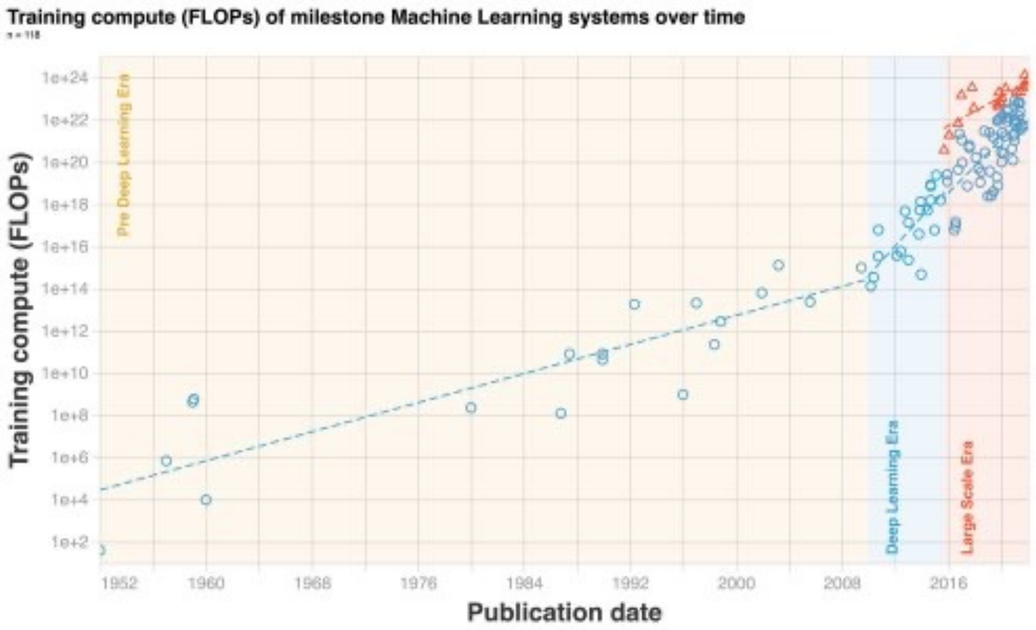

AI Challenges

- Unsustainable compute trajectory
- Lack of explainability and transparency
- Lack of robustness and reliability
- Struggle in some tasks

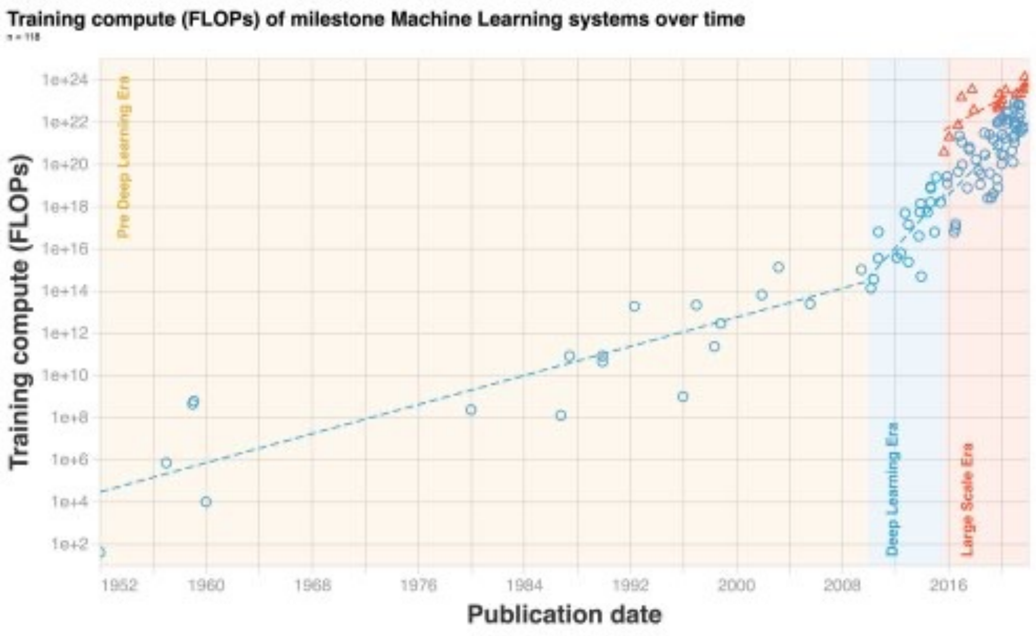

AI Challenges

- Unsustainable compute trajectory
- Lack of explainability and transparency
- Lack of robustness and reliability
- Struggle in some tasks

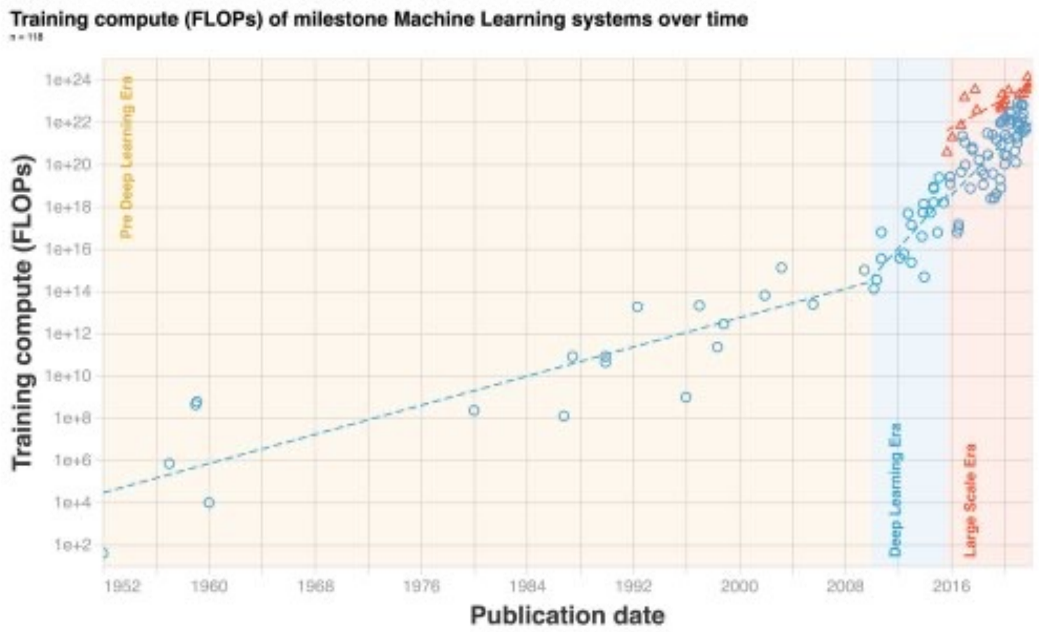

Are there more trees than animals?

AI Challenges

- Unsustainable compute trajectory
- Lack of explainability and transparency
- Lack of robustness and reliability
- Struggle in some tasks

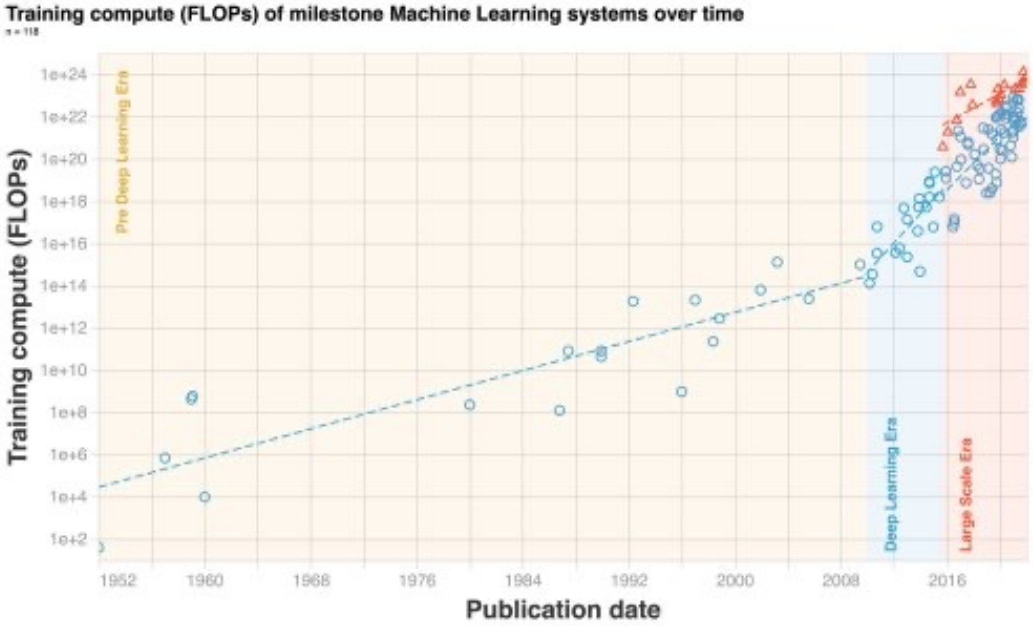

Are there more trees than animals?

What's the shape of object closest to large cylinder

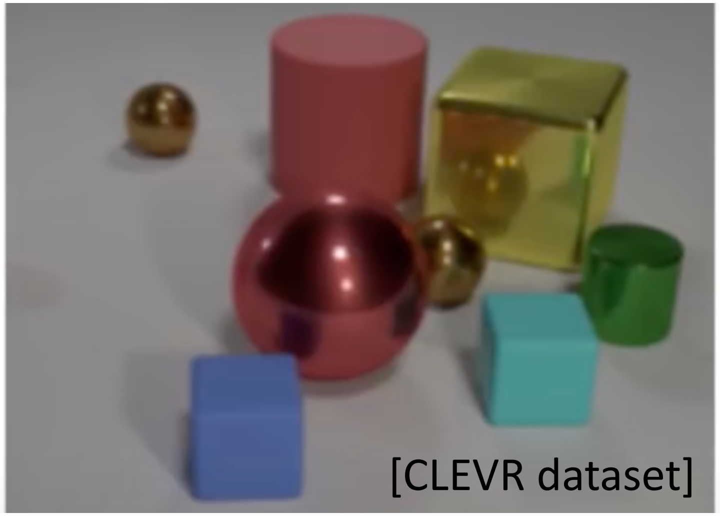

[CLEVR dataset]

Neural Networks / Deep Learning

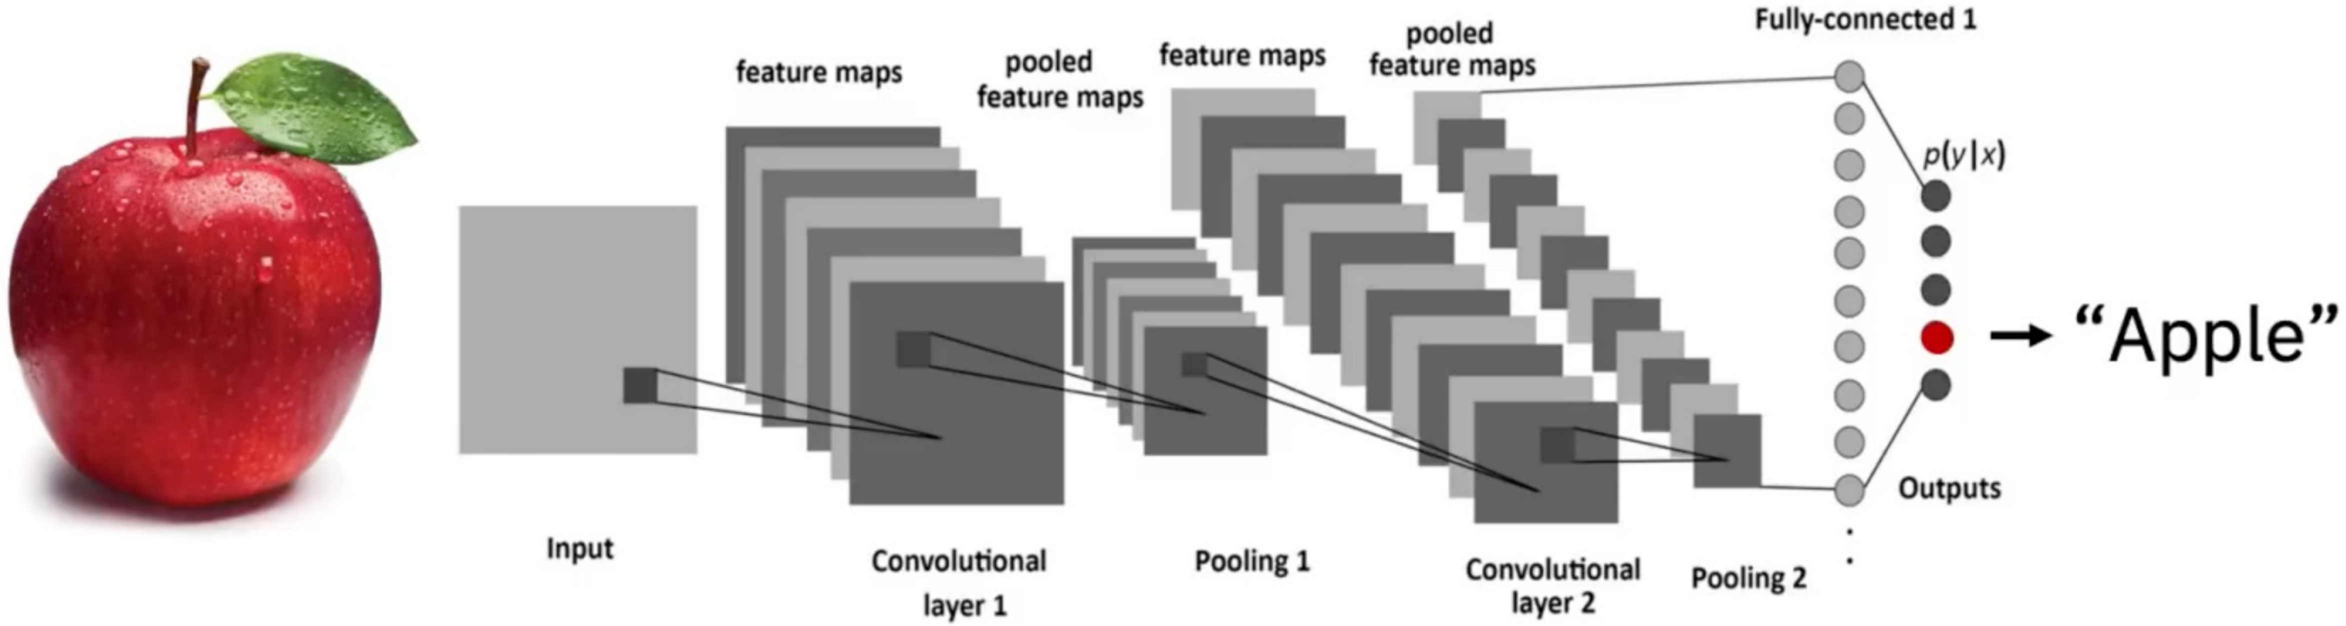

[Credit to MIT 6.S191, David Cox]

Symbolic AI

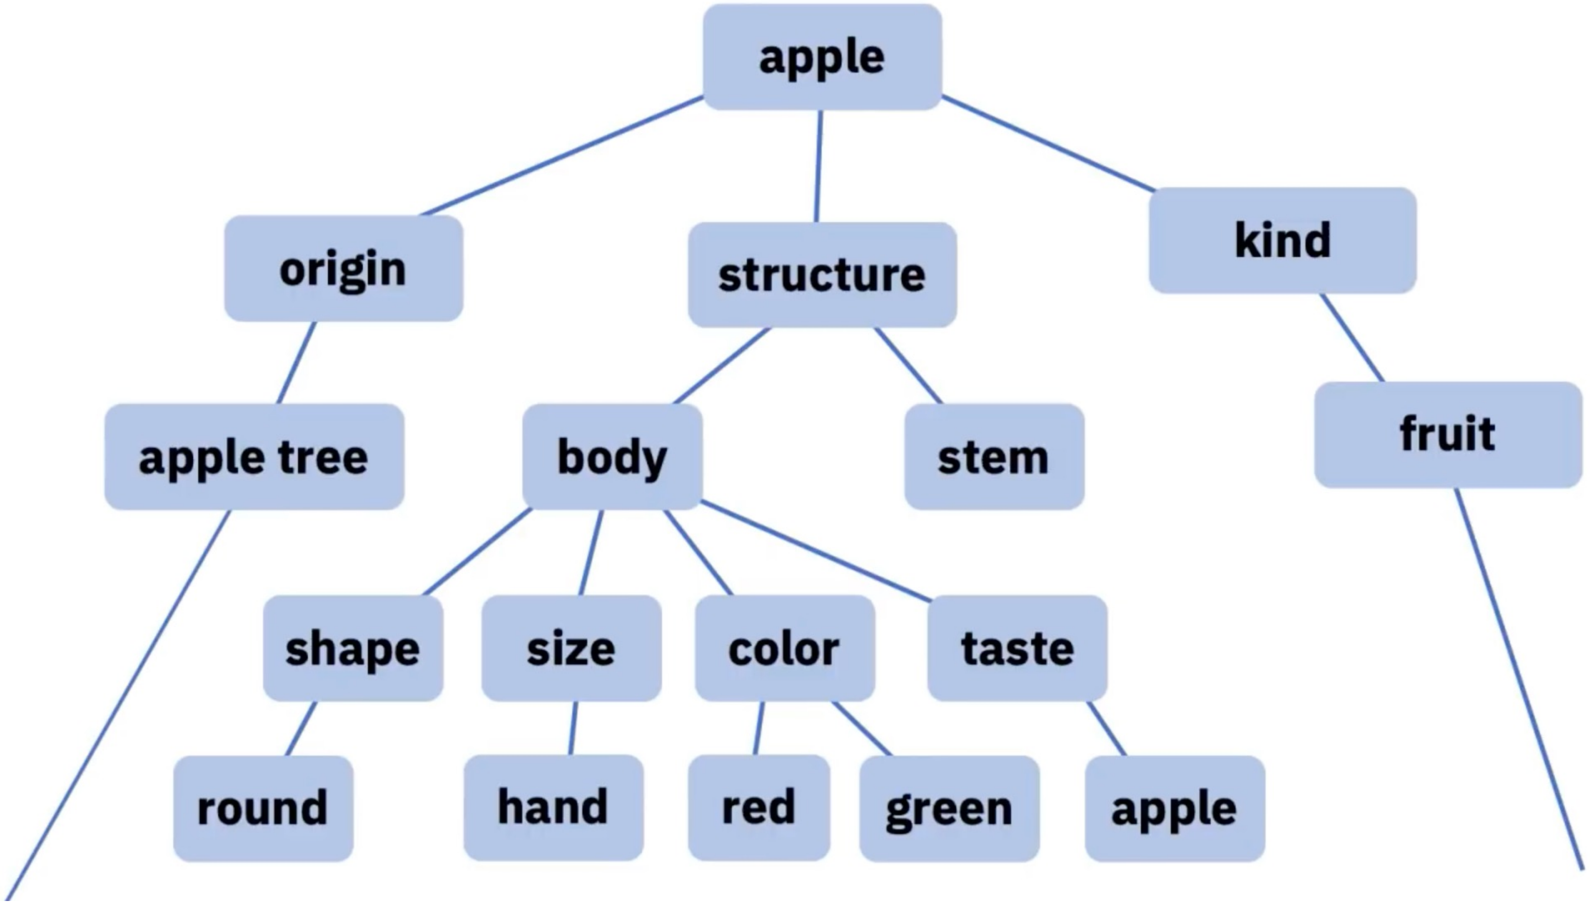

[Credit to MIT 6.S191, David Cox]

Neuro-Symbolic AI

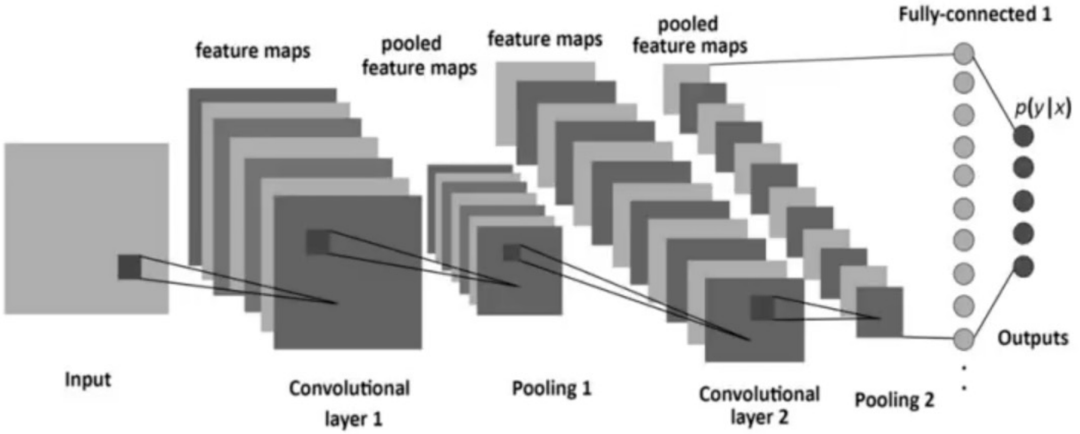

Neural Network

+

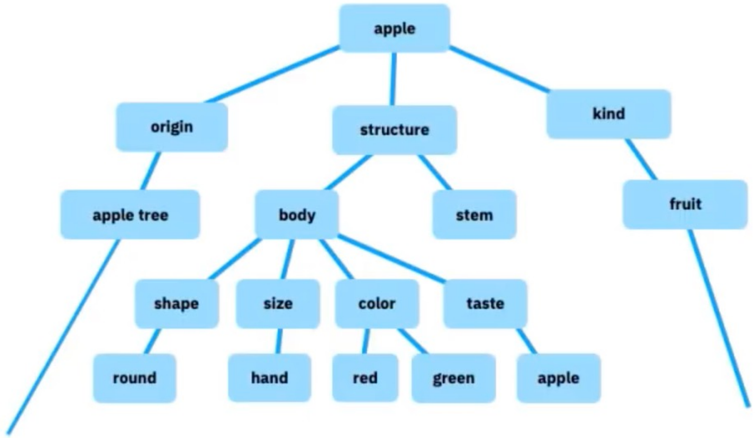


```
(:action pickup
 :parameters (?b1 ?b2 - block)
 :precondition (and (on ?b1 ?b2)
                    (hand-clear))
 :effect (and (not (hand-clear))
              (not (on ?b1 ?b2))
              (holding ?b1))
)
```

Symbolic AI

Neuro-Symbolic AI

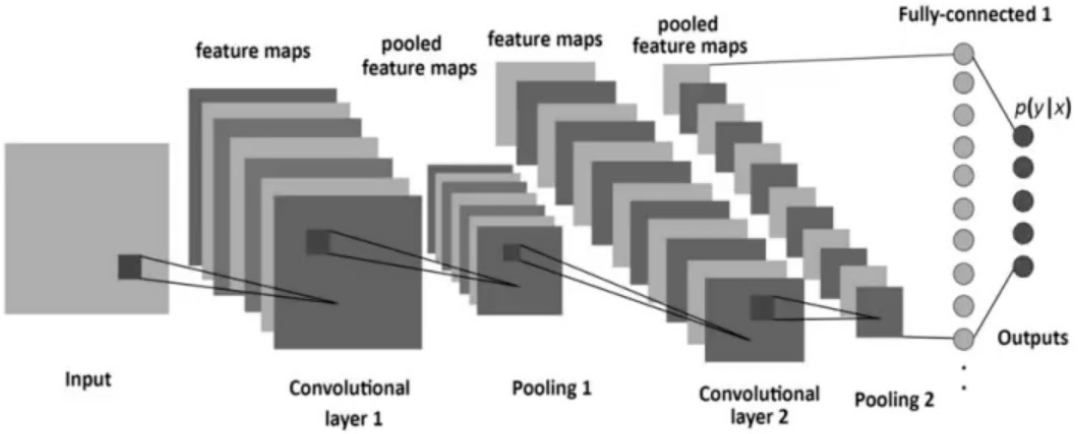

Neural Network

\+

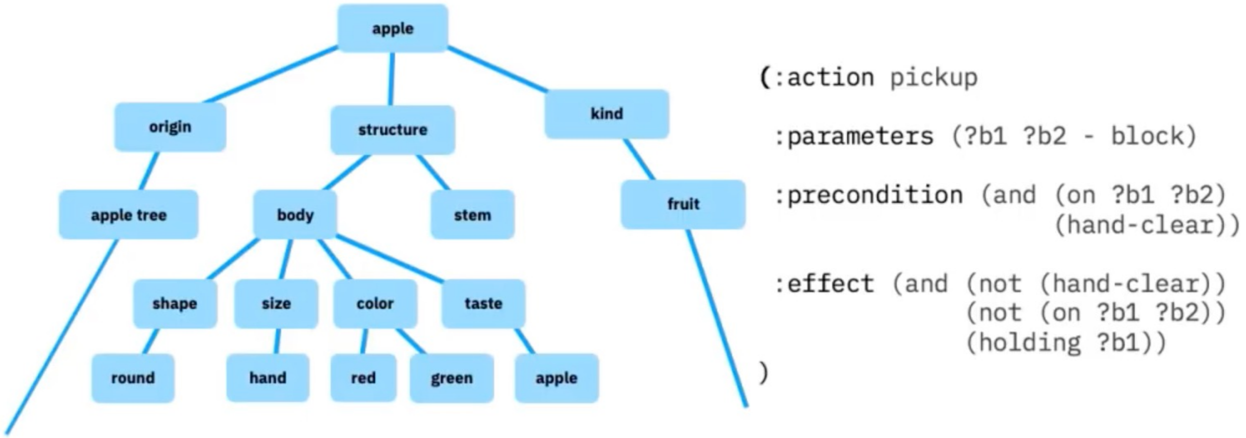

Symbolic AI

- Features
  - **Neuro**: scalable, flexible, handle inconsistency
  - **Symbolic**: interpretable, explainable, data-efficient

Neuro-Symbolic AI

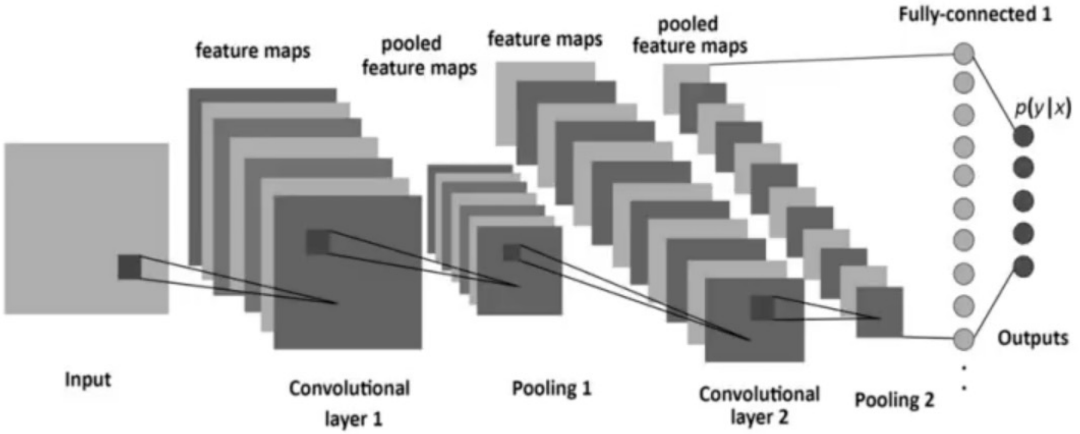

Neural Network

+

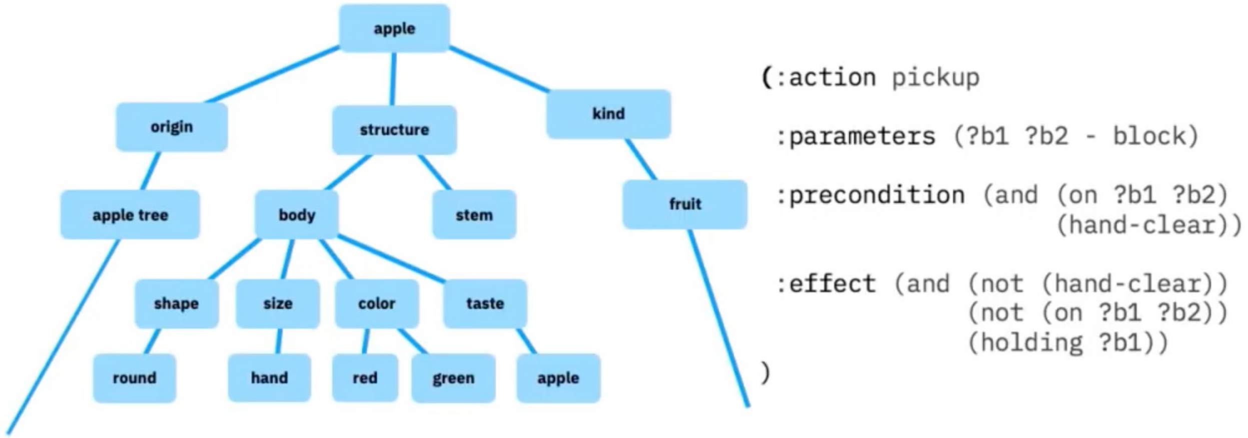

Symbolic AI

- Features
  - **Neuro**: scalable, flexible, handle inconsistency
  - **Symbolic**: interpretable, explainable, data-efficient

- Advantages:
  - Improve efficiency, robustness, and explainability
  - Human-like **reasoning** and **cognition** capability

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

Step 1: Categorizing Neuro-Symbolic AI **Algorithm**

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

Step 1: Categorizing Neuro-Symbolic AI **Algorithm**

Step 2: Benchmarking selecting algorithms on current **Hardware**

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

Step 1: Categorizing Neuro-Symbolic AI **Algorithm**

Step 2: Benchmarking selecting algorithms on current **Hardware**

Step 3: Our view for Neuro-Symbolic AI **Challenges and Opportunities**

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

Step 1: Categorizing Neuro-Symbolic AI Algorithm

Step 2: Benchmarking selecting algorithms on current **Hardware**

Step 3: Our view for Neuro-Symbolic AI **Challenges and Opportunities**

Neuro-Symbolic AI Algorithms

- Classification Criterion:

How neuro-symbolic
integrated into a
cohesive system
(Henry Kautz's taxonomy)

Neuro-Symbolic AI Algorithms

- Classification Criterion:

How neuro-symbolic integrated into a cohesive system (Henry Kautz's taxonomy)

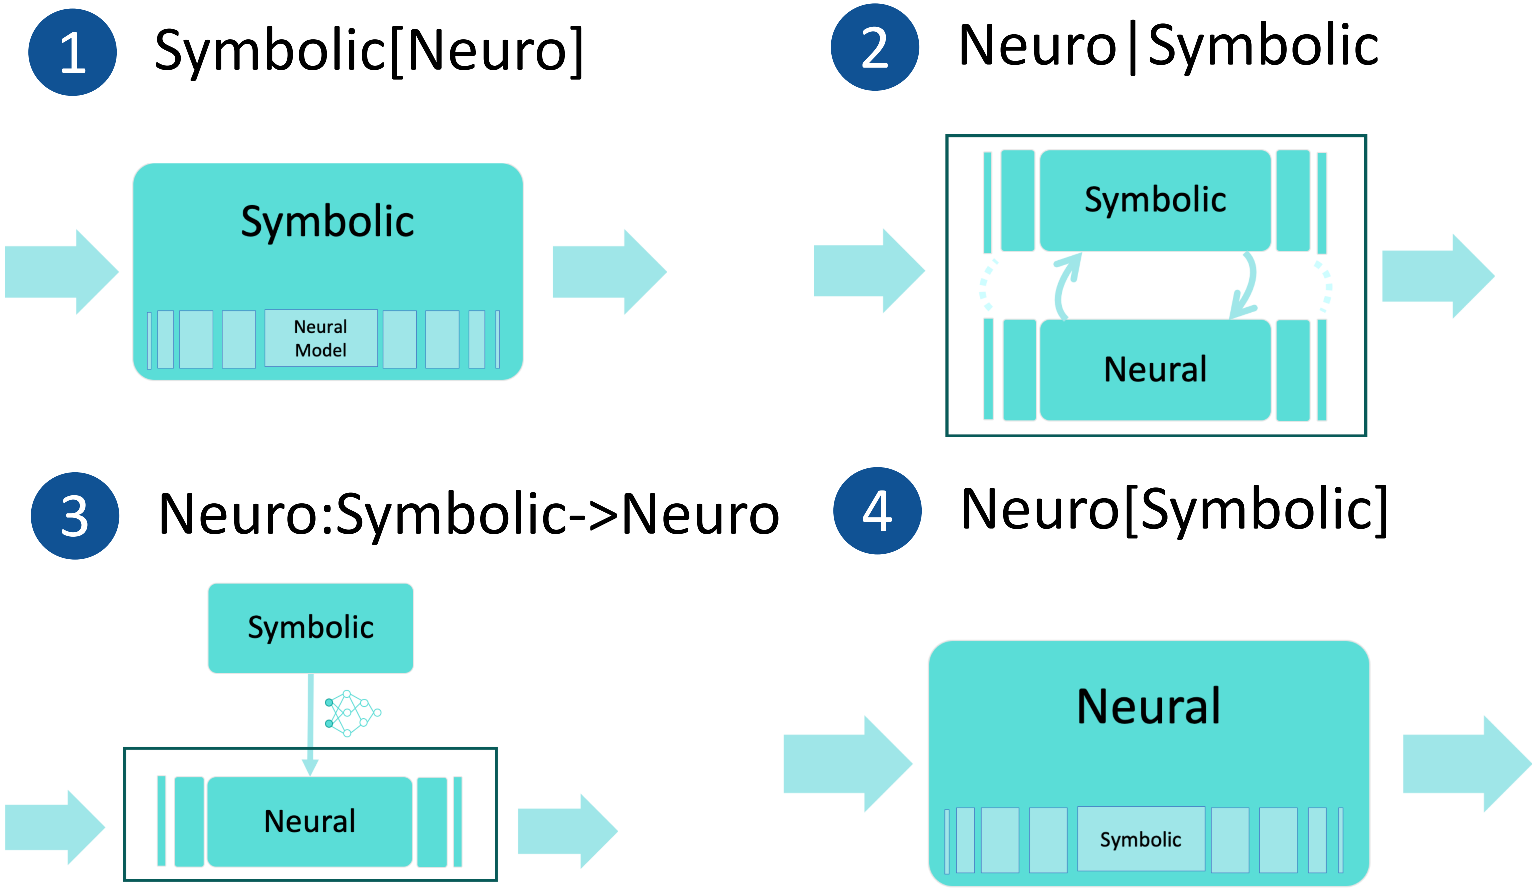

Neuro-Symbolic AI Algorithms

1 Symbolic[Neuro]

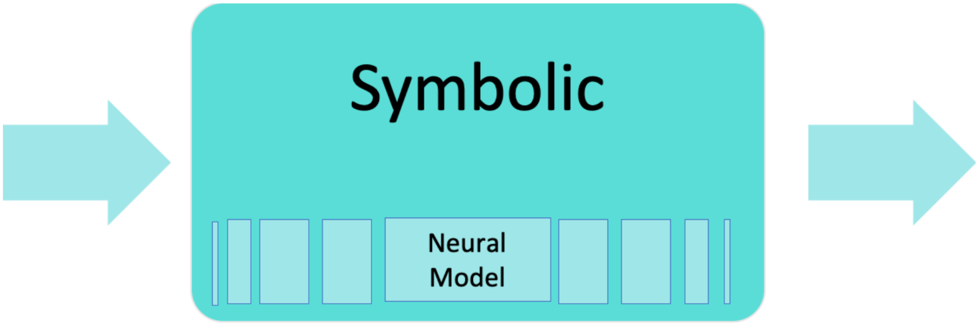

Example:
AlphaGo[1]
AlphaZero[2]

[1] Nature 2017; [2] Nature 2020

Neuro-Symbolic AI Algorithms

1 Symbolic[Neuro]

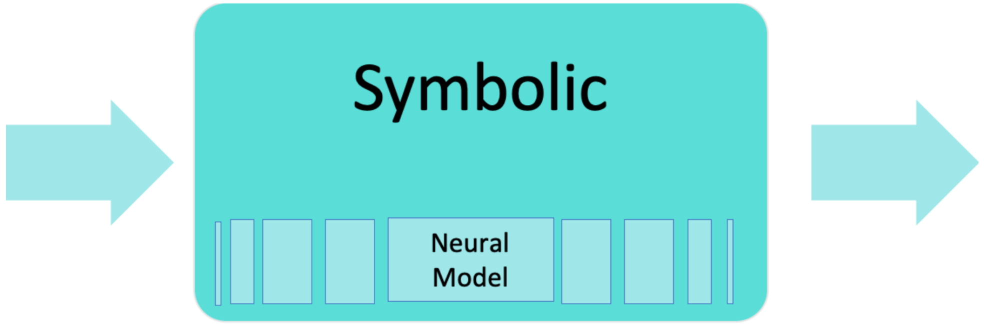

Example:
AlphaGo[1]
AlphaZero[2]

[1] Nature 2017; [2] Nature 2020

2 Neuro | Symbolic

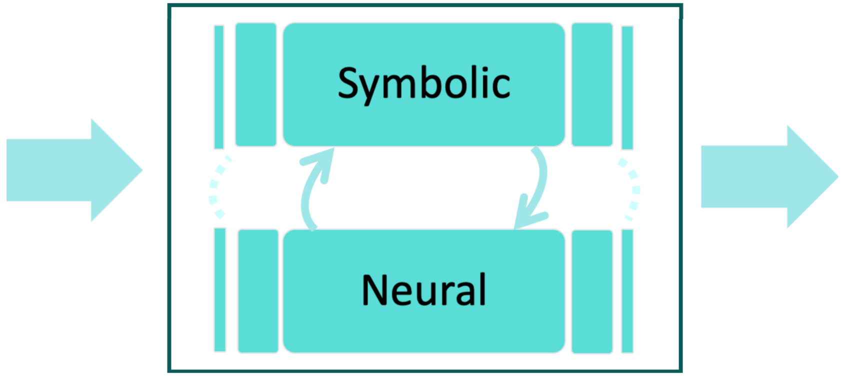

Example:
neuro-vector-symbolic architecture[3]
neuro-probabilistic logic programming[4]
neuro-symbolic dynamic reasoning[5]

[3] Nature 2023; [4] AI 2021; [5] ICLR 2020

Neuro | Symbolic Example

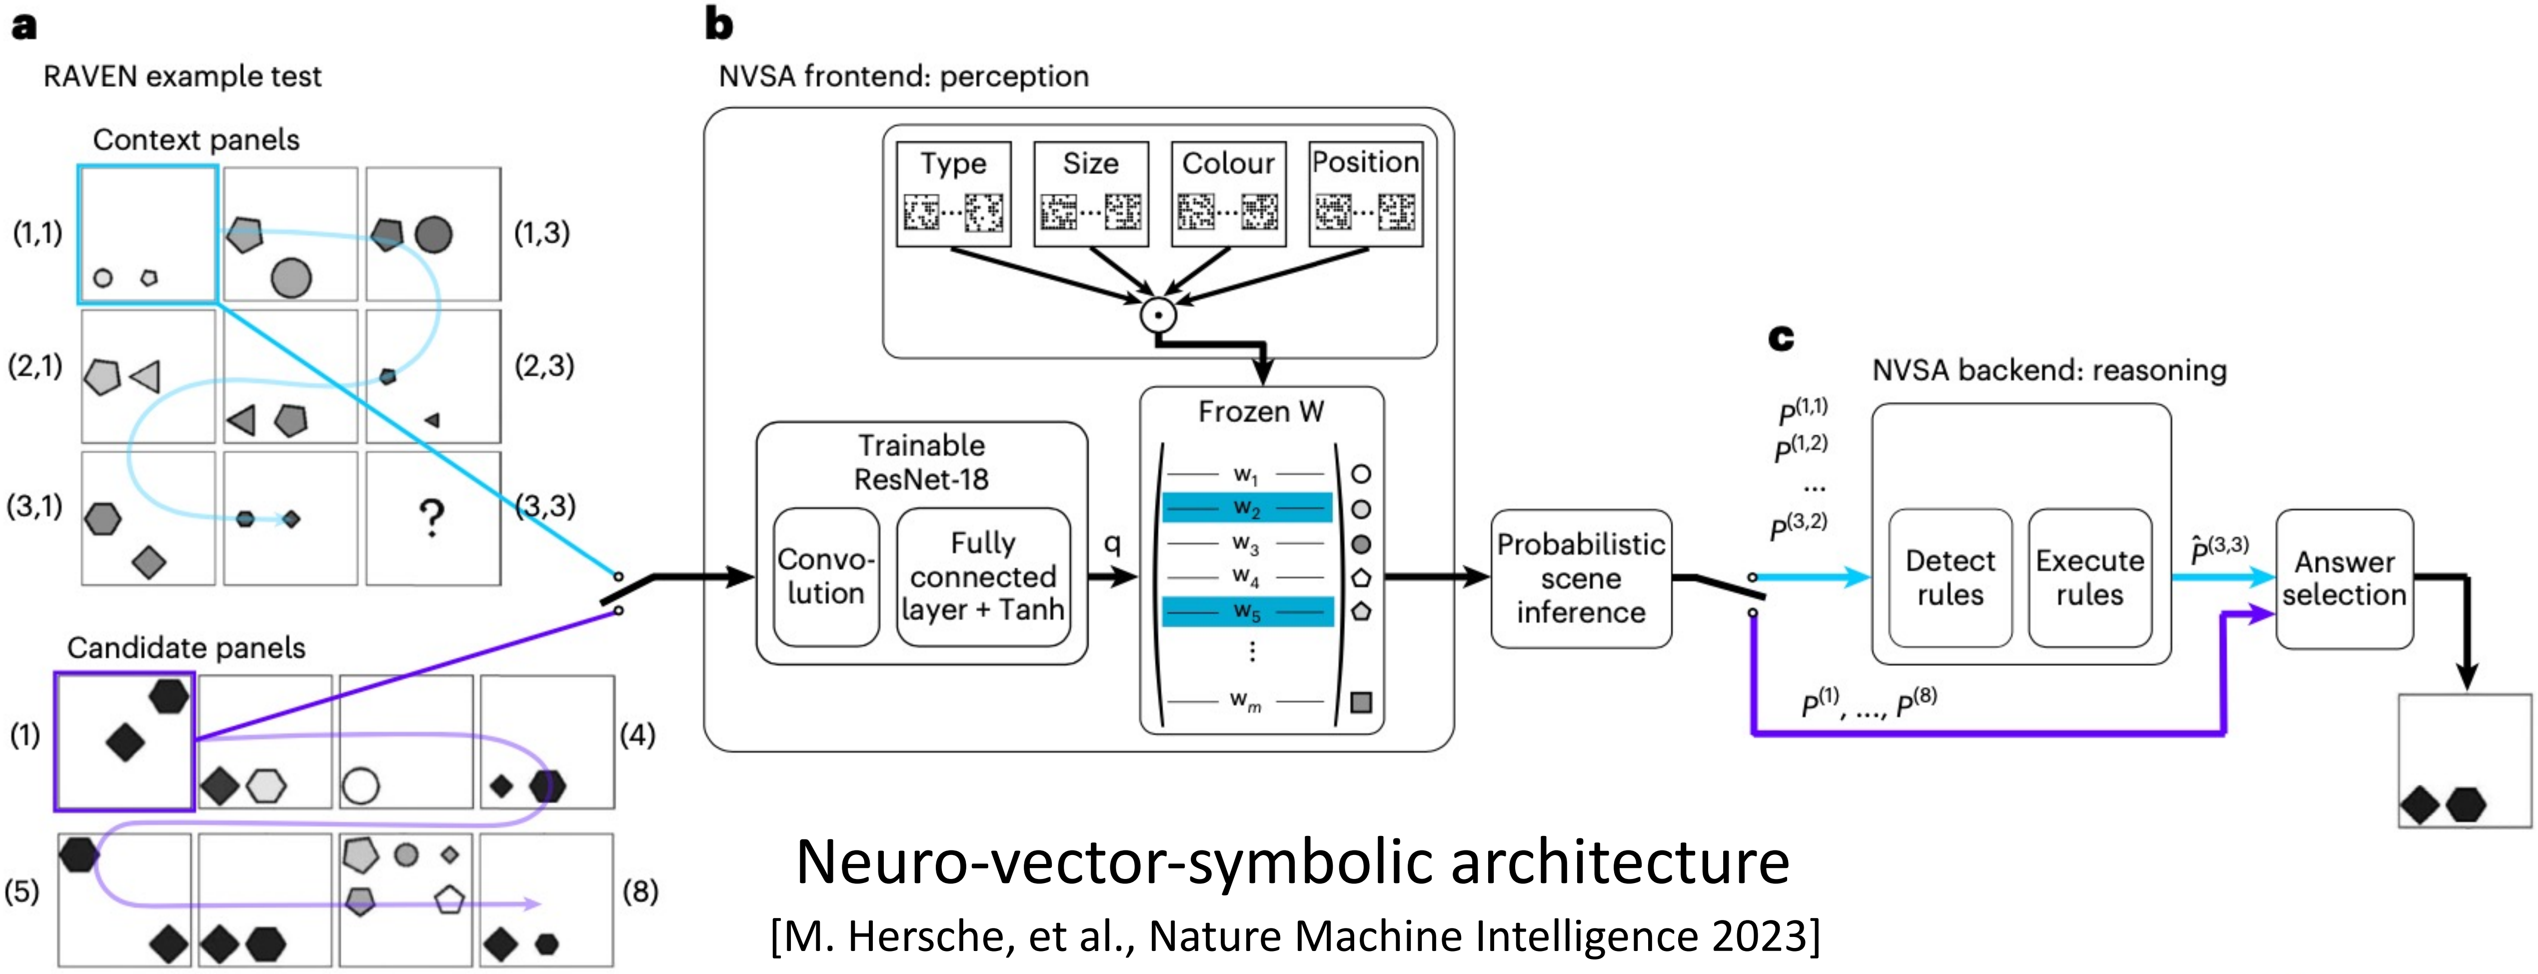

Neuro-vector-symbolic architecture

[M. Hersche, et al., Nature Machine Intelligence 2023]

Neuro | Symbolic Example

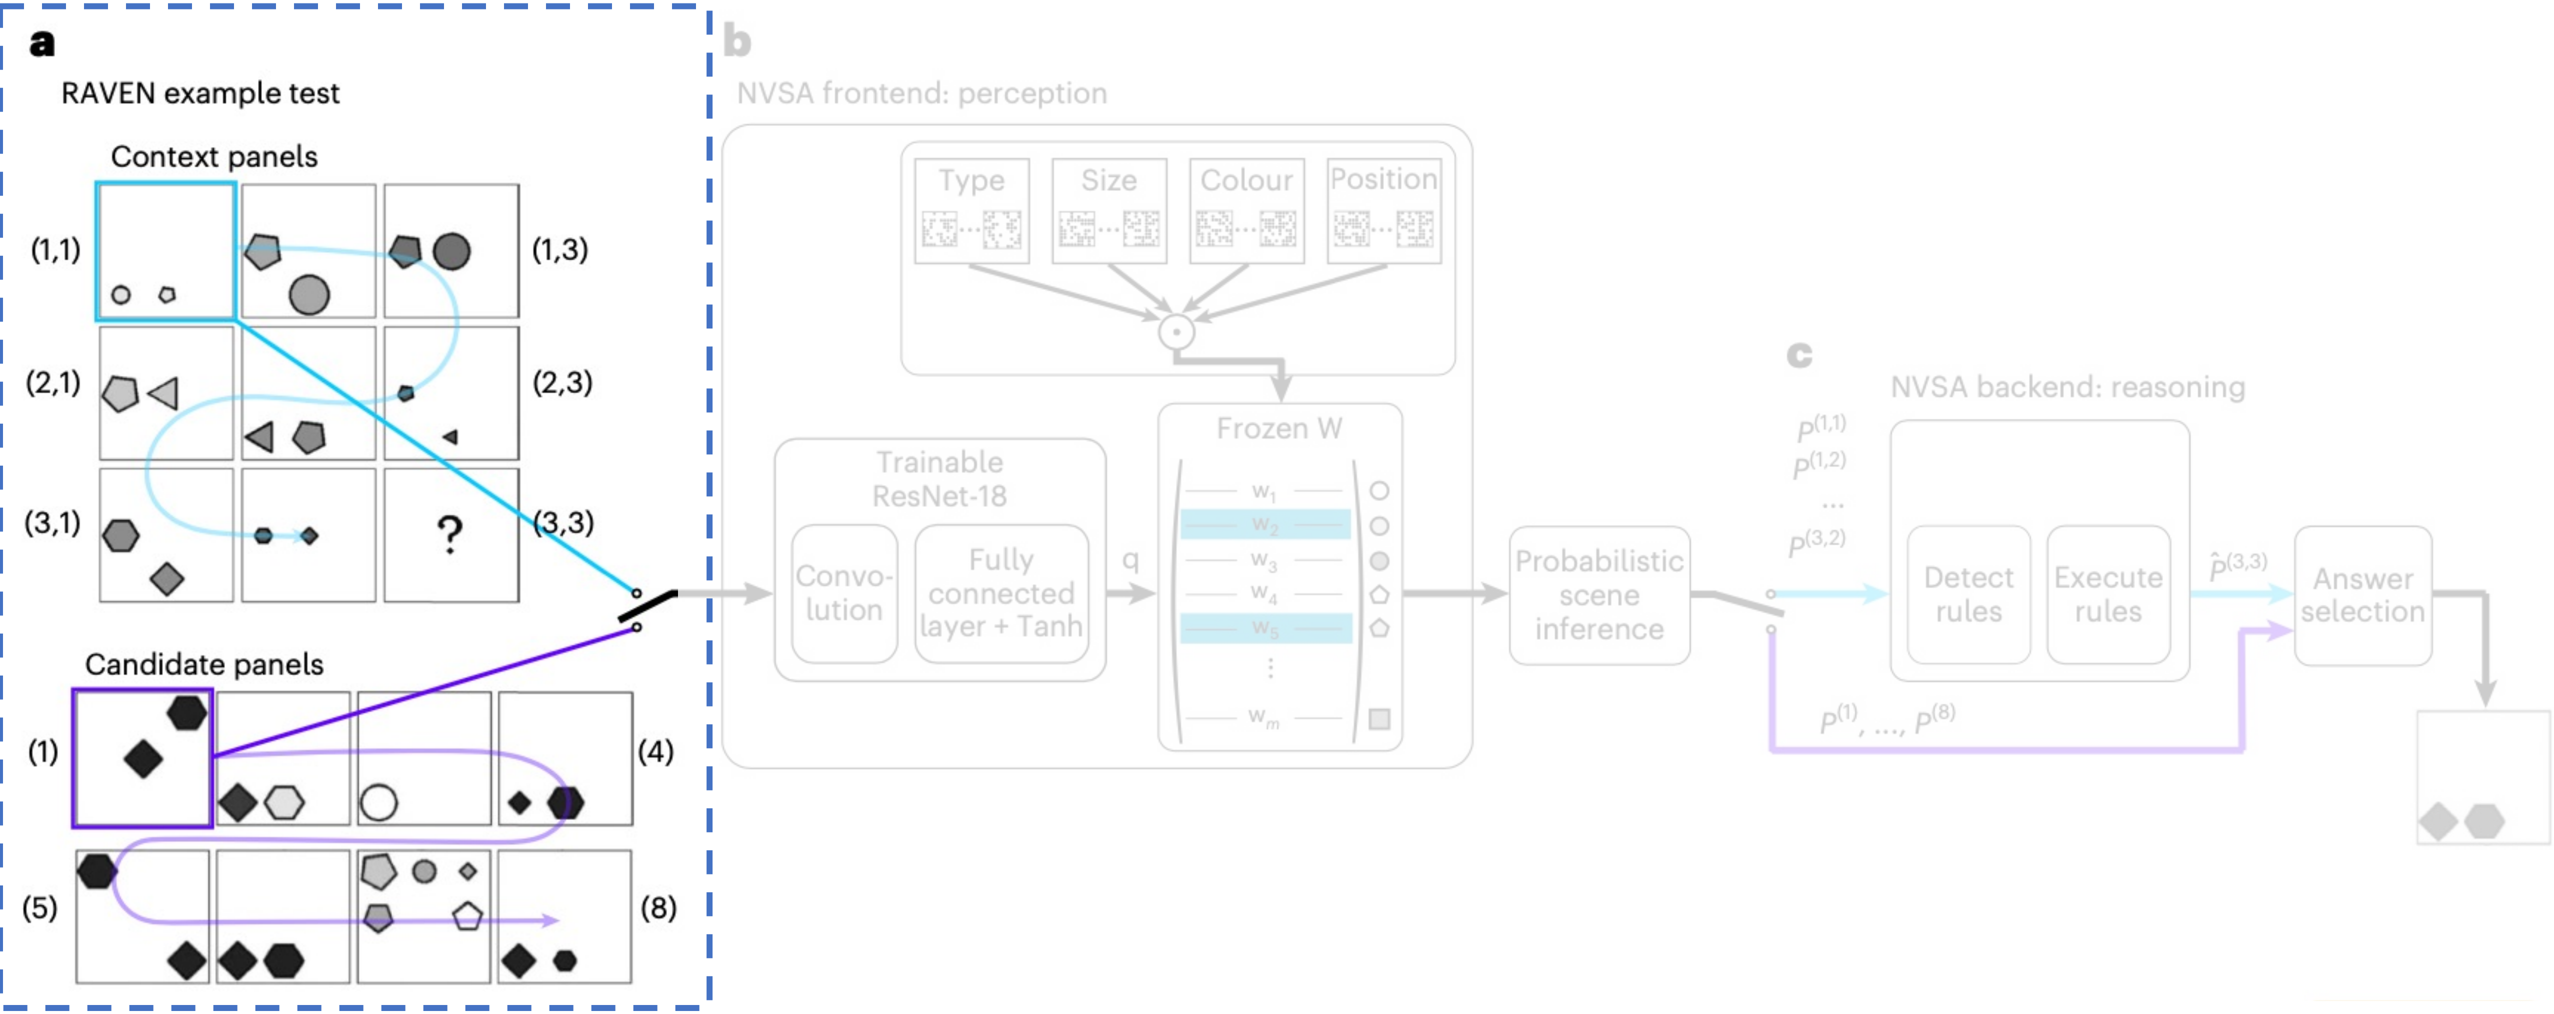

Neuro | Symbolic Example

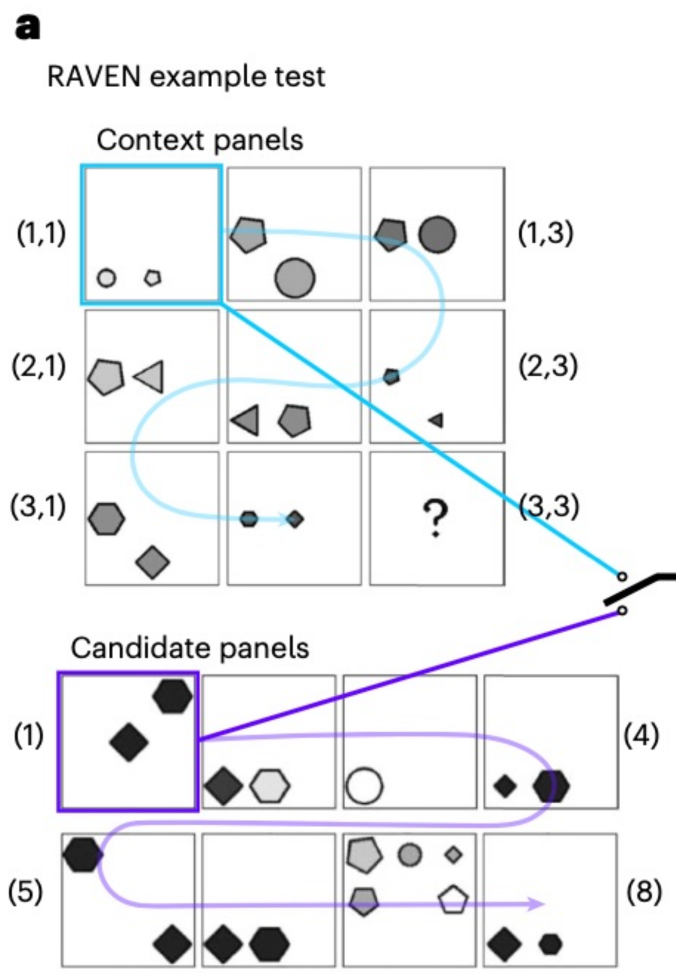

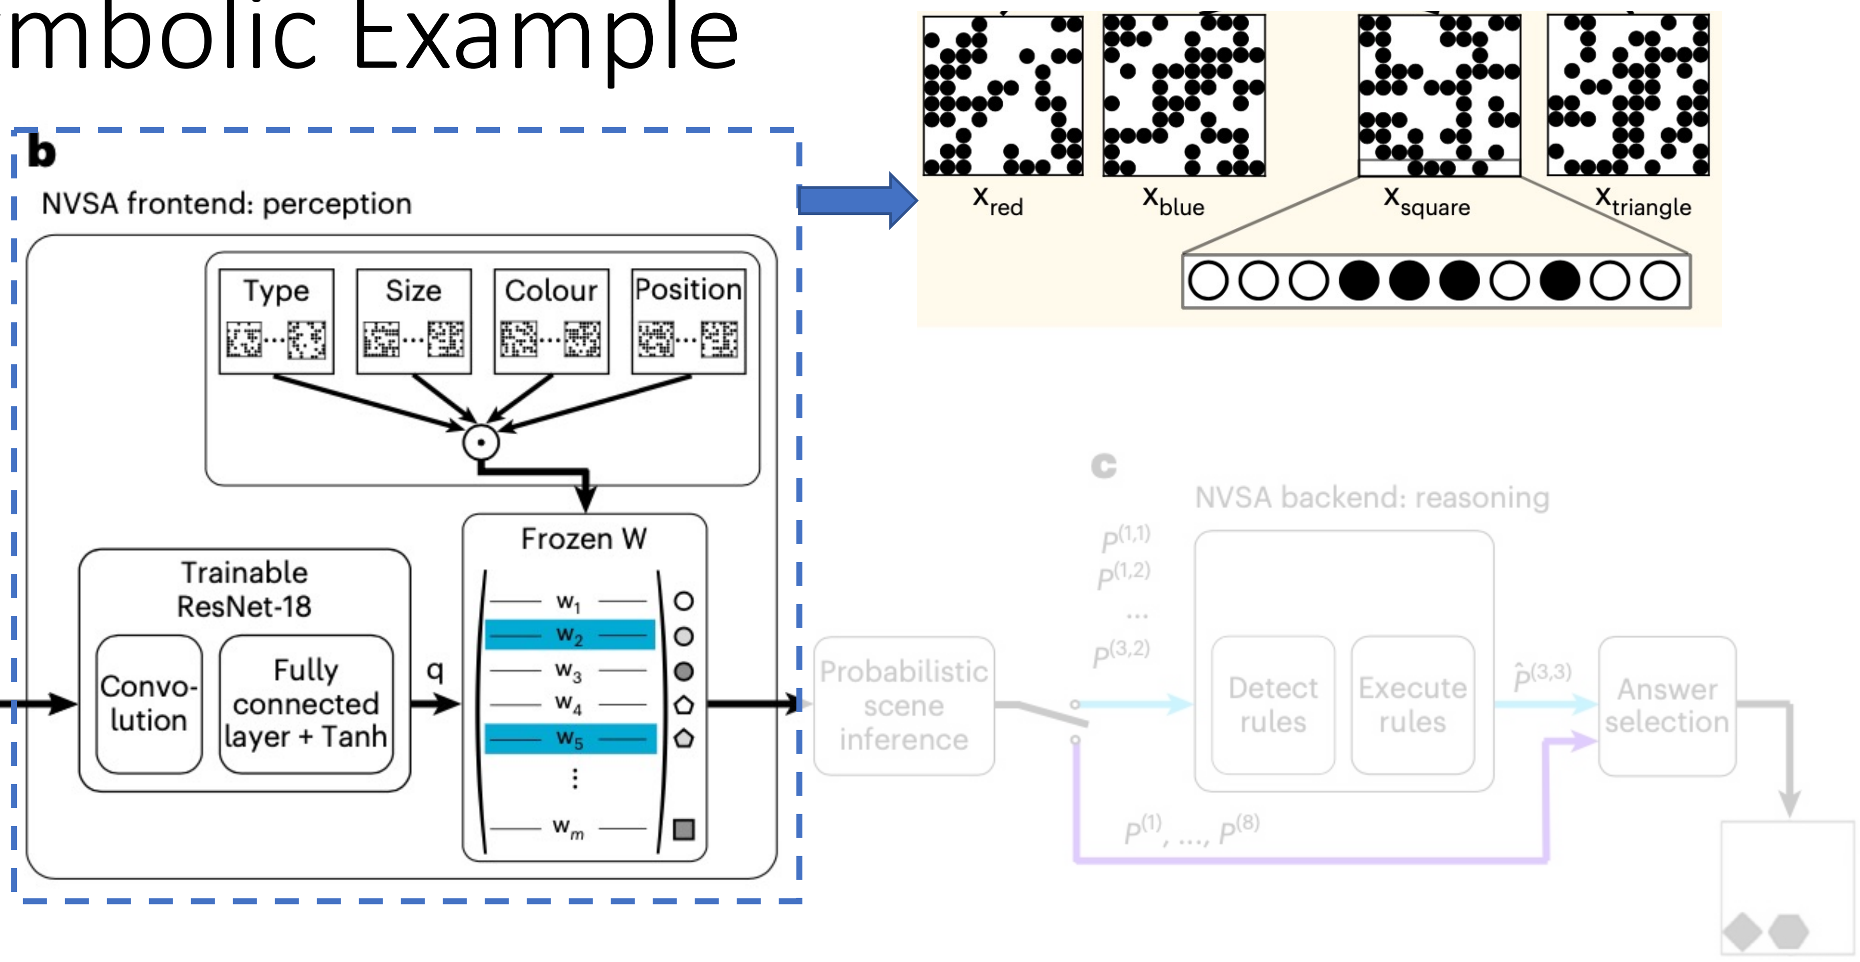

Neuro - Perception

Neuro | Symbolic Example

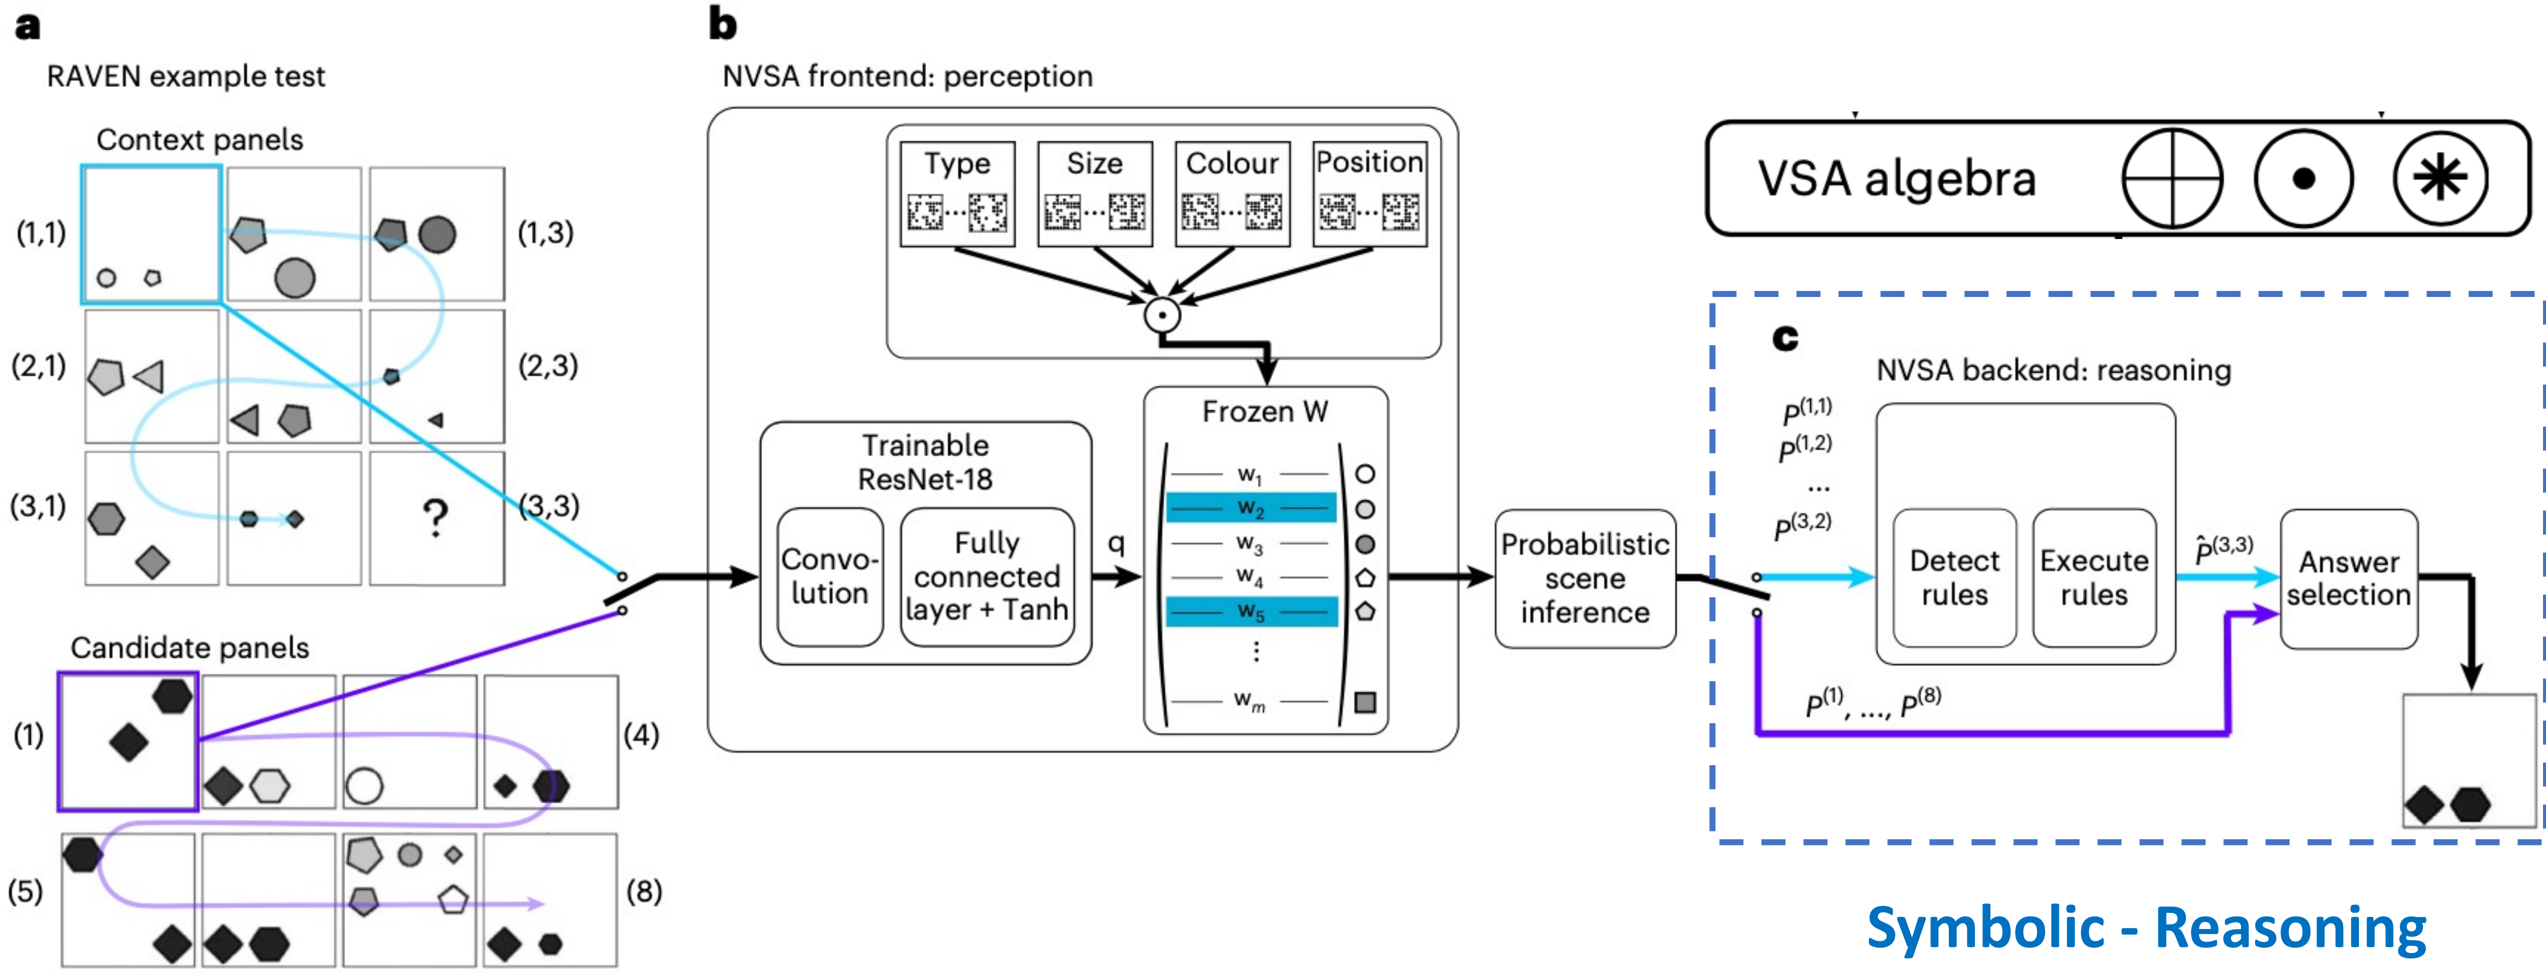

Symbolic - Reasoning

Neuro-Symbolic AI Algorithms

3 Neuro:Symbolic -> Neuro

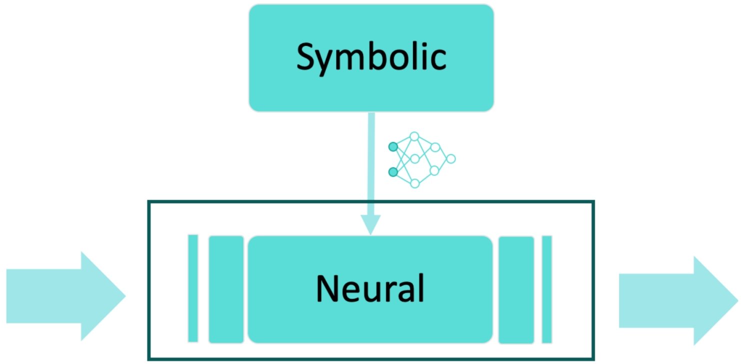

Example:
logical neural network[6]
Inductive logic programming[7]

Neuro-Symbolic AI Algorithms

3 Neuro:Symbolic -> Neuro

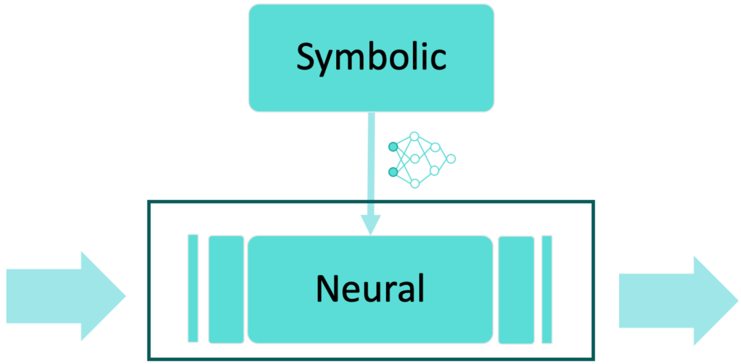

Example:
logical neural network[6]
Inductive logic programming[7]

[6] NeurIPS 2020; [7] JAIR 2018

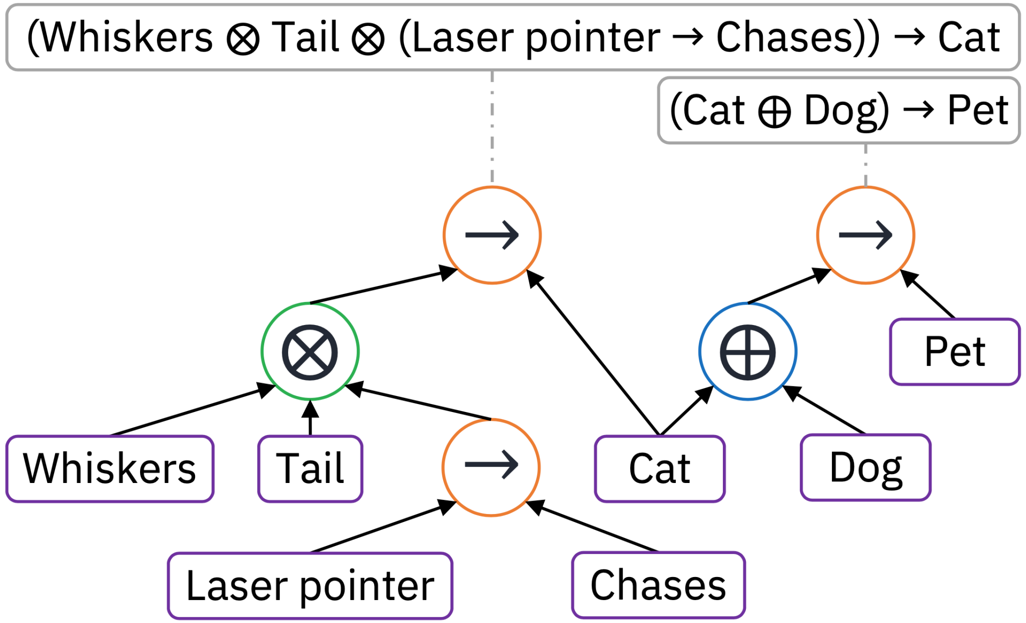

logical neural network[6]

Neuro-Symbolic AI Algorithms

4 Neuro[Symbolic]

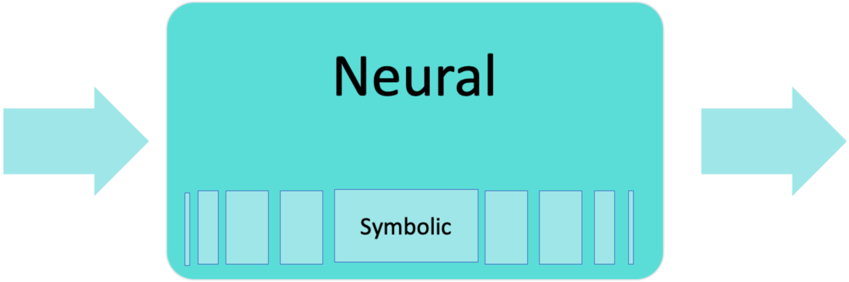

Example:
logical tensor network[8]
deep ontology network[9]

[8] AI 2022; [9] JAIR 2020

Neuro-Symbolic AI Algorithms

4 Neuro[Symbolic]

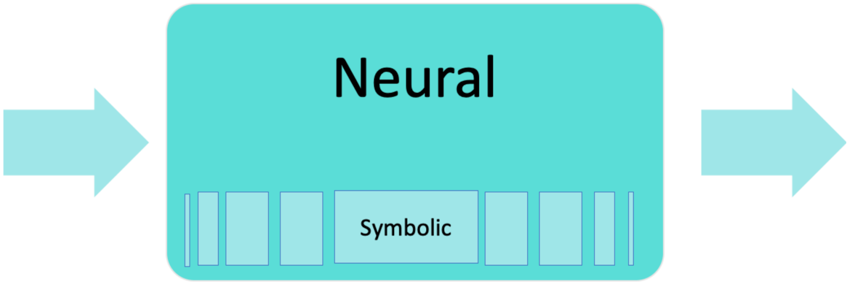

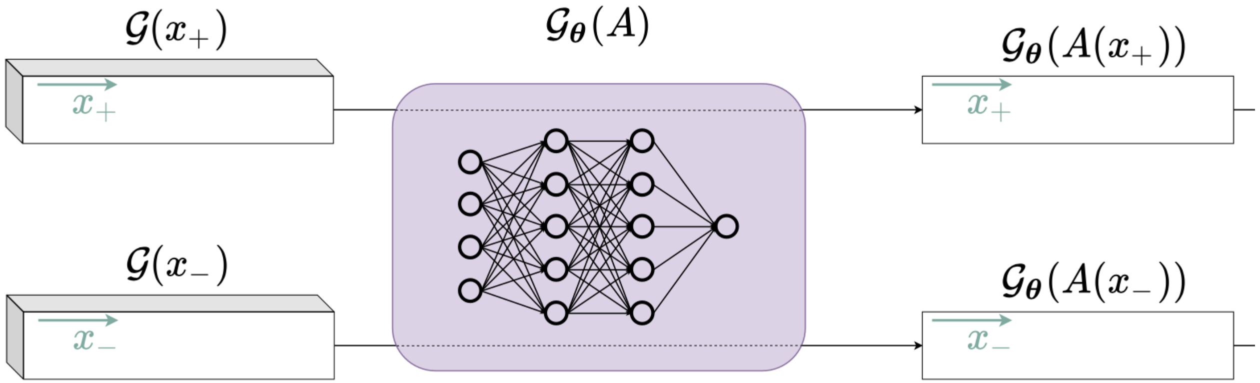

Example:
logical tensor network[8]
deep ontology network[9]

logical tensor network[8]

[8] AI 2022; [9] JAIR 2020

Neuro-Symbolic AI Algorithms

- **Classification Criterion**:

How neuro-symbolic integrated into a cohesive system (Henry Kautz's taxonomy)

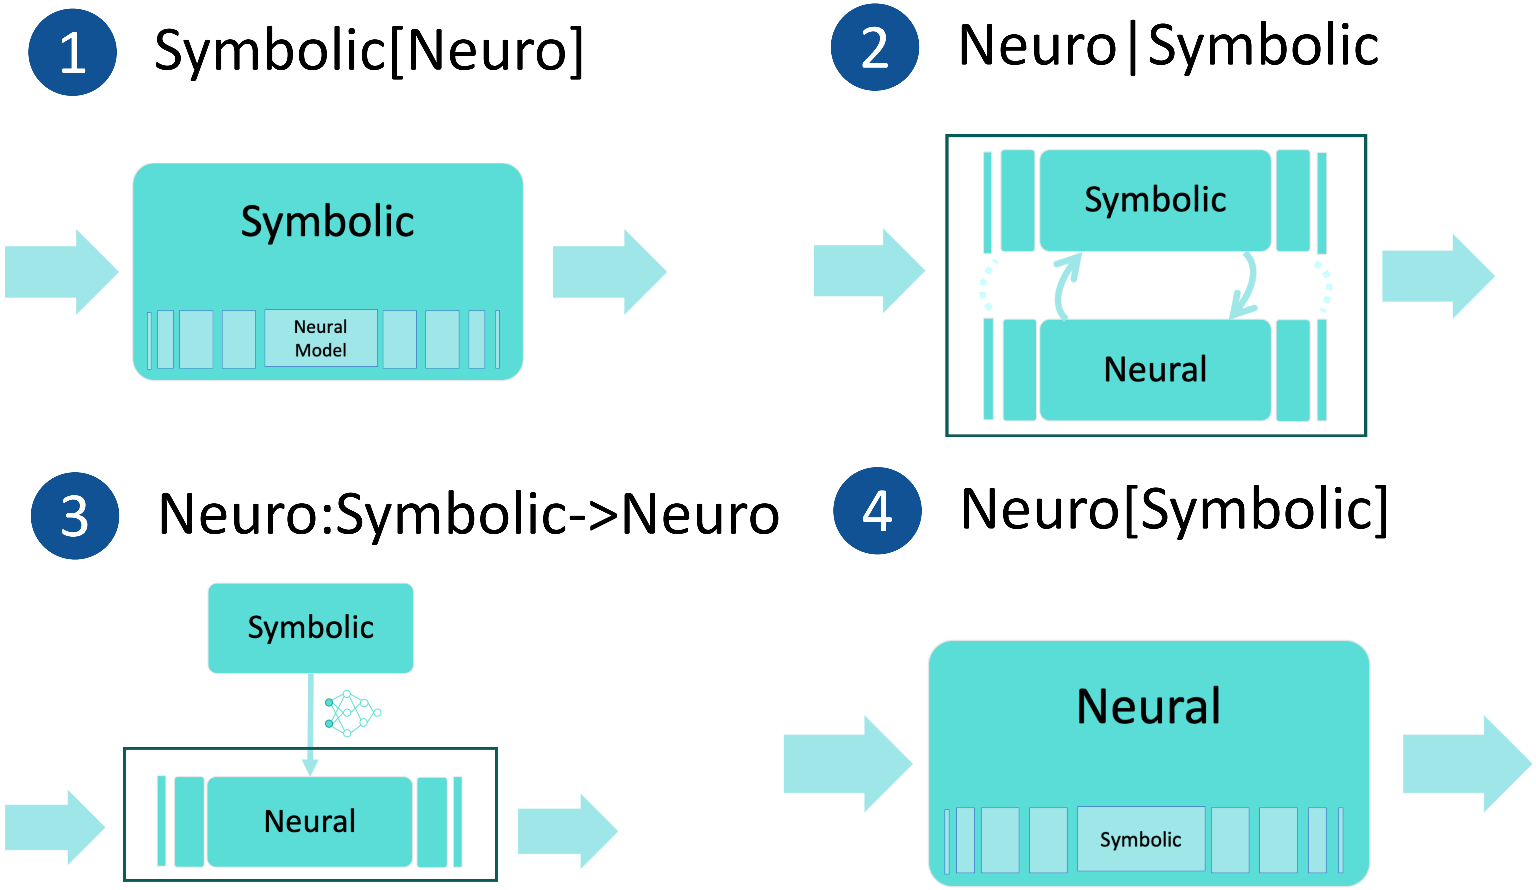

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

Step 1: Categorizing Neuro-Symbolic AI **Algorithm**

Step 2: Benchmarking selecting algorithms on current Hardware

Step 3: Our view for Neuro-Symbolic AI **Challenges and Opportunities**

Selected Neuro-Symbolic AI Models

- Logical Neural Network

Neuro:Symbolic->Neuro

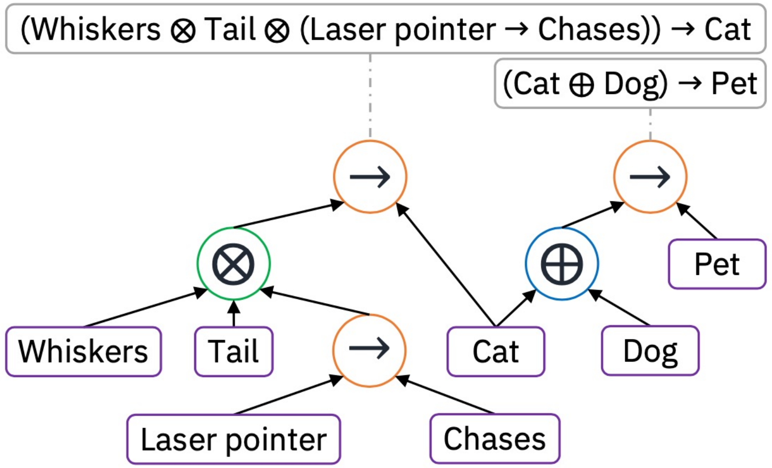

- Logical Tensor Network

Neuro[Symbolic]

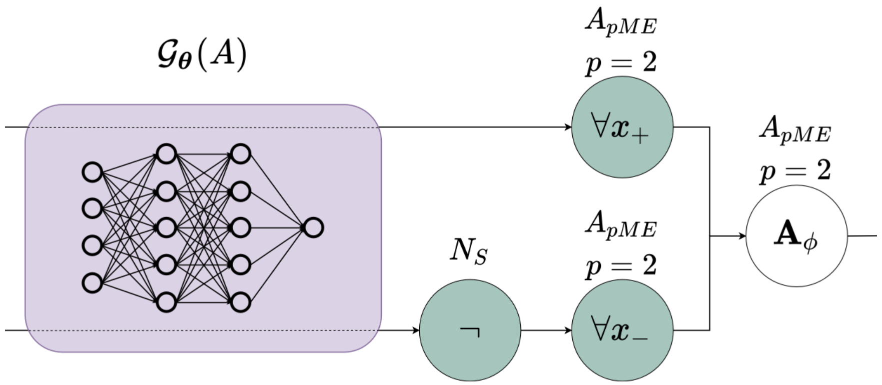

- Neuro-Vector-Symbolic Architecture

Neuro | Symbolic

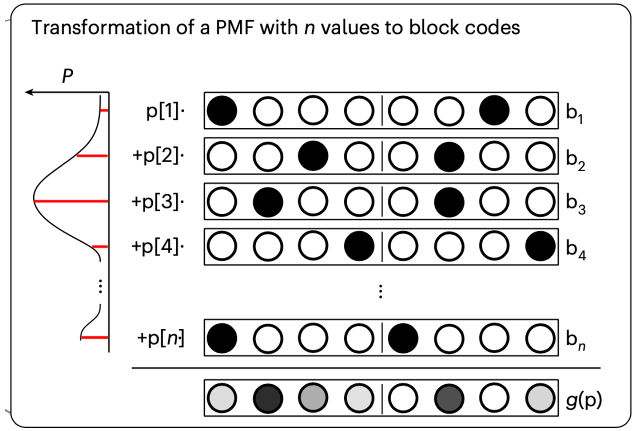

Compute Runtime Analysis

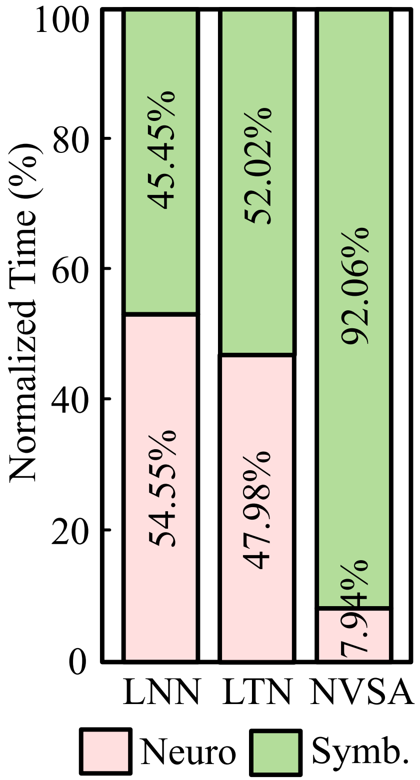

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

- Measurement Method: Pytorch Profiler
- Hardware: Intel Xeon 4114 CPU, Nvidia RTX 2080 Ti GPU

Compute Runtime Analysis

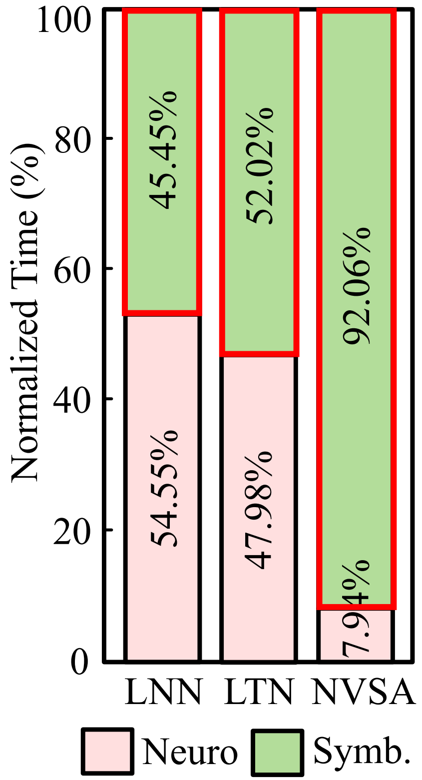

- Observations
  - **Runtime Breakdown**: symbolic workloads are not negligible and may become a system bottleneck.

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

- Measurement Method: Pytorch Profiler
- Hardware: Intel Xeon 4114 CPU, Nvidia RTX 2080 Ti GPU

Compute Runtime Analysis

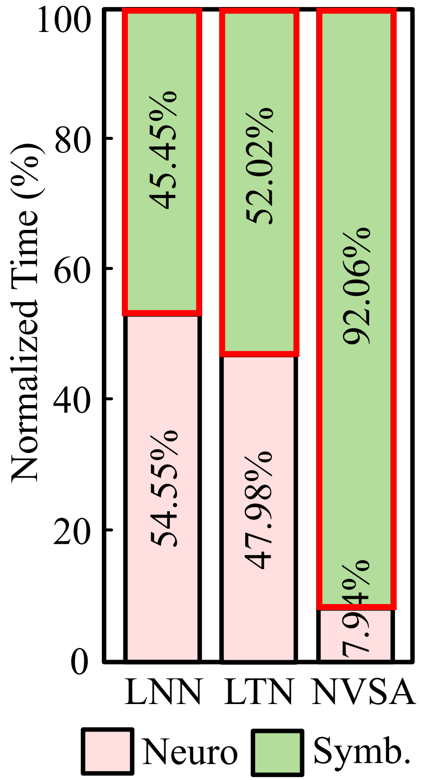

- Observations
  - **Runtime Breakdown**: symbolic workloads are not negligible and may become a system bottleneck.
  - **Runtime Breakdown**: for same test size, neuro vs. symbolic runtime ratio remains stable.

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

- Measurement Method: Pytorch Profiler
- Hardware: Intel Xeon 4114 CPU, Nvidia RTX 2080 Ti GPU

Compute Runtime Analysis

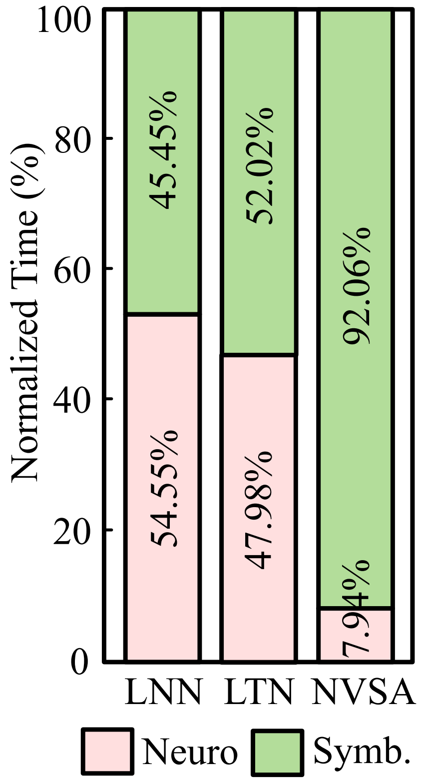

- Observations
  - **Runtime Breakdown**: symbolic workloads are not negligible and may become a system bottleneck.
  - **Runtime Breakdown**: for same test size, neuro vs. symbolic runtime ratio remains stable.
  - **Runtime Scalability**: when test size increase from 2 to 3, runtime increase 5.02x.

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

- Measurement Method: Pytorch Profiler
- Hardware: Intel Xeon 4114 CPU, Nvidia RTX 2080 Ti GPU

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

Compute Operator Analysis

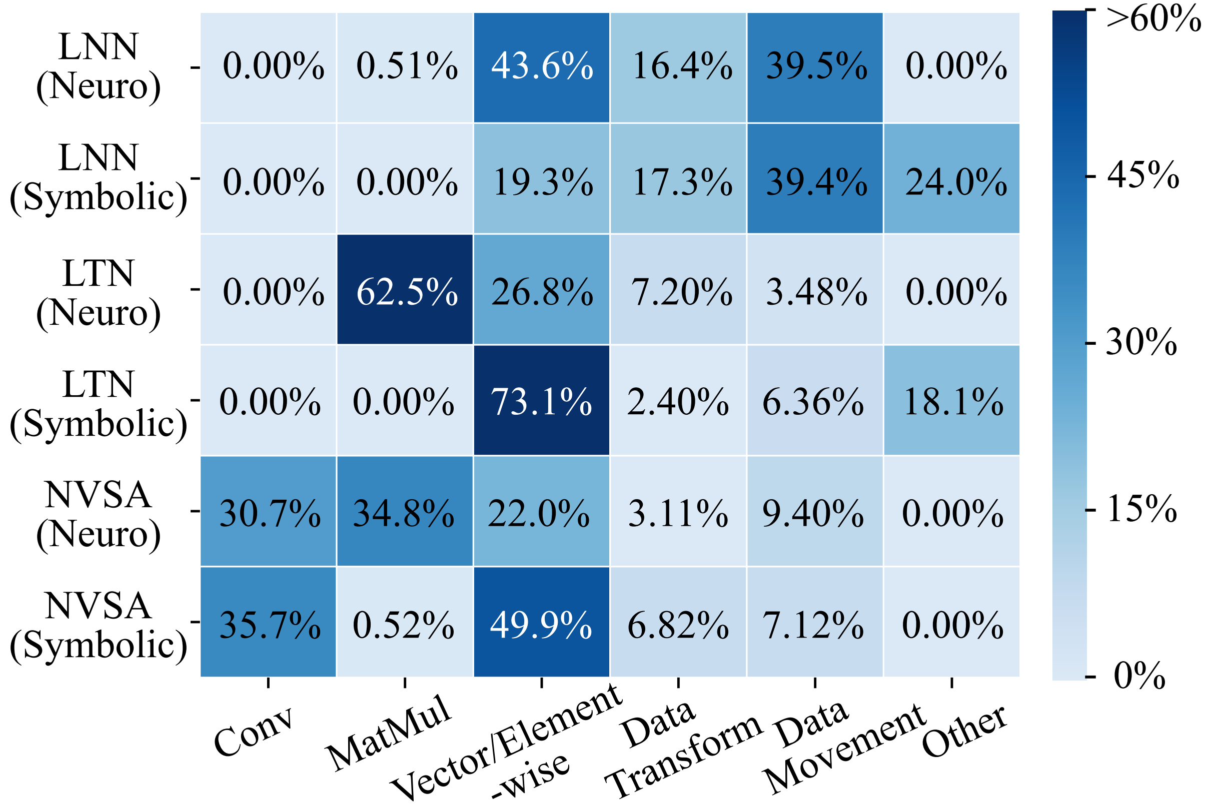

- **Six operators**: convolution, matrix multiplication (MatMul), vector/element-wise operation, data transformation, data movement, others (logic)

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

Compute Operator Analysis

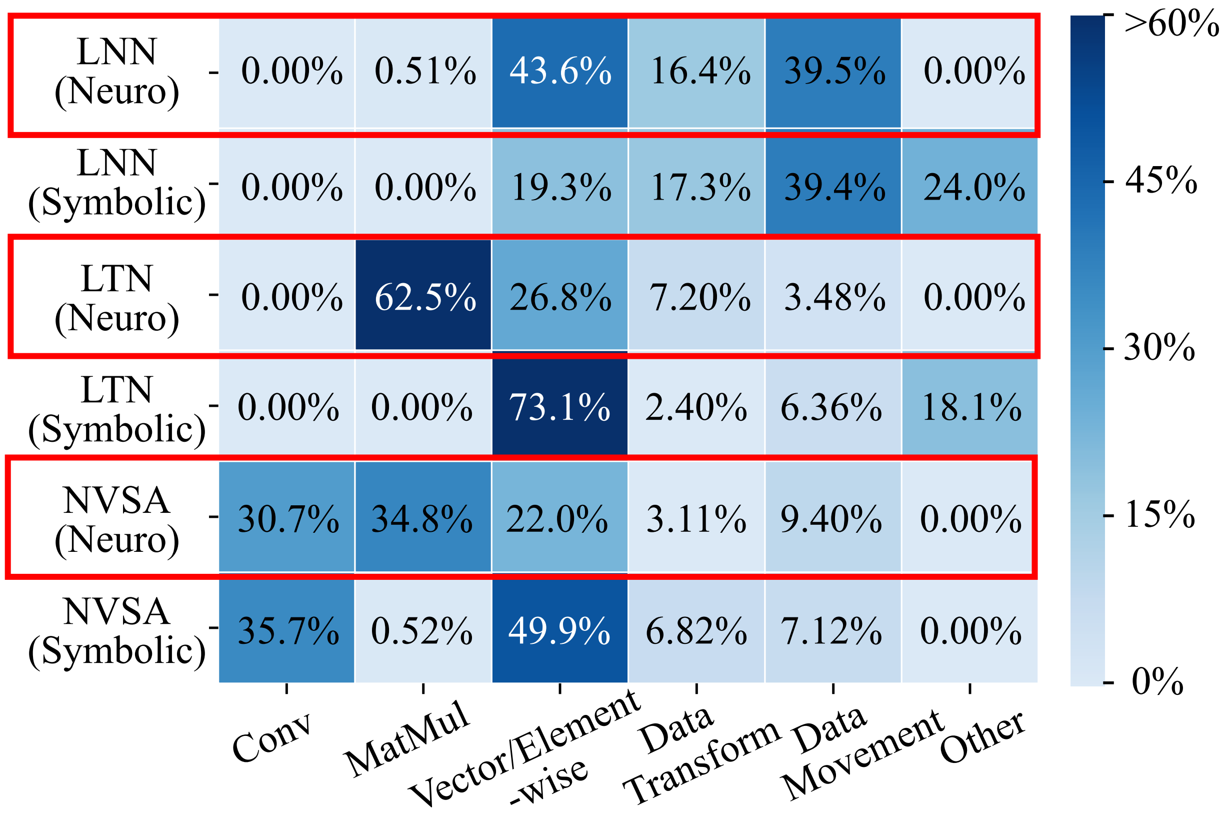

- Observations
  - **Neuro Workload**: dominated by MatMul and activation operations

- **Six operators**: convolution, matrix multiplication (MatMul), vector/element-wise operation, data transformation, data movement, others (logic)

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

Compute Operator Analysis

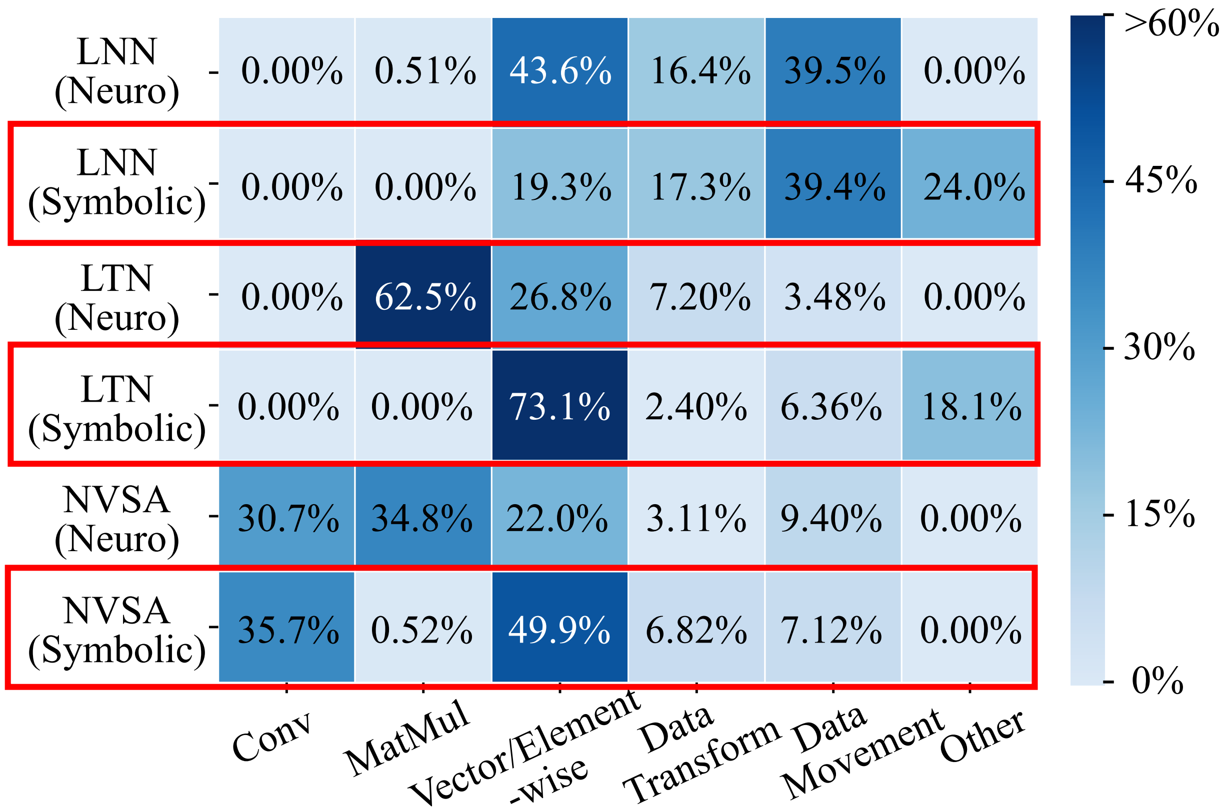

- Observations
  - **Neuro Workload**: dominated by MatMul and activation operations
  - **Symbolic Workload**: dominated by vector and scalar operations - low operational intensity and complex control flows (inefficient on GPUs)

- **Six operators**: convolution, matrix multiplication (MatMul), vector/element-wise operation, data transformation, data movement, others (logic)

LNN: Logical Neural Network
LTN: Logical Tensor Network
NVSA: Neuro-vector-symbolic architecture

Compute Operator Analysis

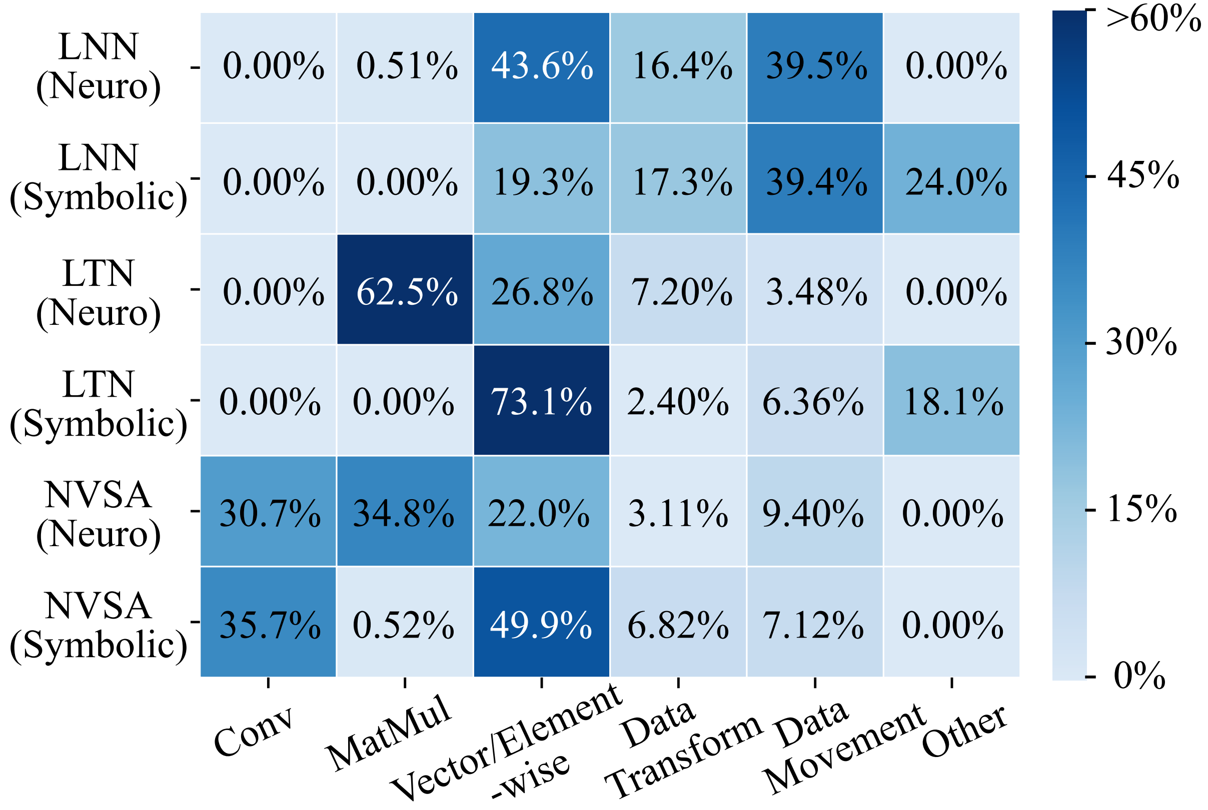

- Observations
  - **Neuro Workload**: dominated by MatMul and activation operations
  - **Symbolic Workload**: dominated by vector and scalar operations - low operational intensity and complex control flows (inefficient on GPUs)
- **Accelerating computation** becomes important

- **Six operators**: convolution, matrix multiplication (MatMul), vector/element-wise operation, data transformation, data movement, others (logic)

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

Step 1: Categorizing Neuro-Symbolic AI **Algorithm**

Step 2: Benchmarking selecting algorithms on current **Hardware**

Step 3: Our view for Neuro-Symbolic AI Challenges and Opportunities

Challenge and Opportunity

Data

🎯 Lack of cognitive datasets

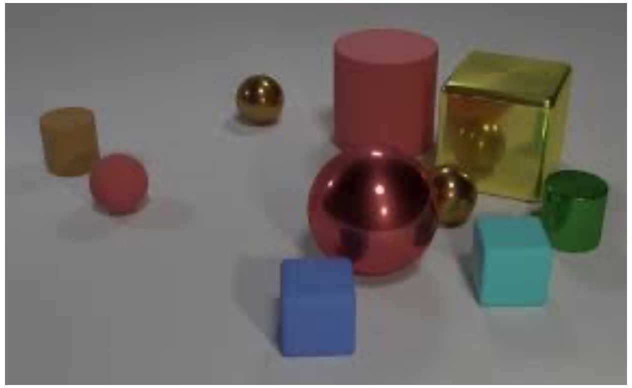

CLEVRER Dataset

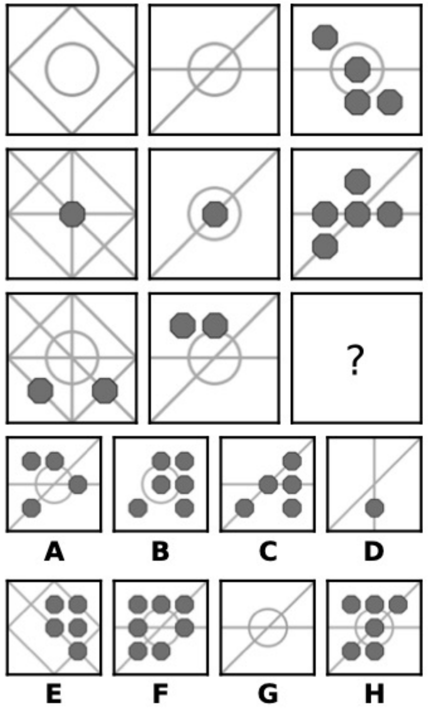

RAVEN Dataset

Challenge and Opportunity

Data | Lack of cognitive datasets | Building ImageNet-like NSAI datasets

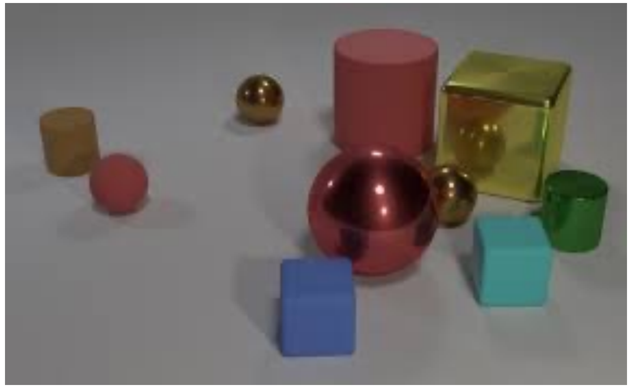

CLEVRER Dataset

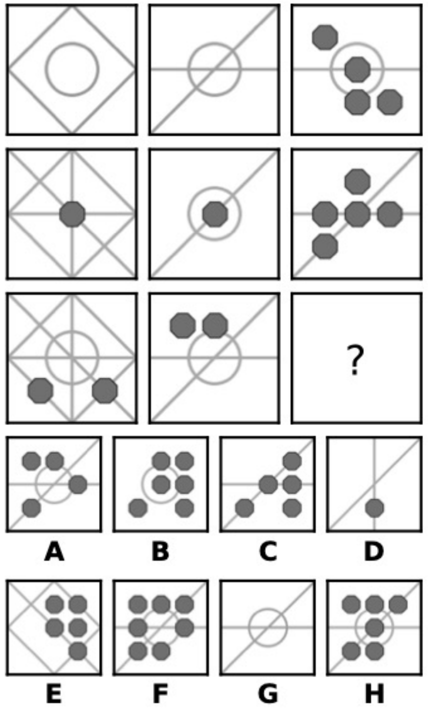

RAVEN Dataset

Human-like AI

Metacognition
Interpretability
Deductive Reasoning
Systematicity
Compositionality
Counterfactual thinking
…

Challenge and Opportunity

	Challenge	Opportunity
Data	Lack of cognitive datasets	Building ImageNet-like NSAI datasets
Model	Nascent integration	Unifying neuro-symbolic-prob models

Neural Network
Scalable
Flexible
Handle inconsistency

\+

Symbolic
Interpretable
Explainable
Data-efficient

\+

Probabilistic
Robust to uncertainty

Challenge and Opportunity

Data	Lack of cognitive datasets	Building ImageNet-like NSAI datasets
Model	Nascent integration	Unifying neuro-symbolic-prob models
Software	Modularity & extensibility	Developing efficient SW framework

Underlying Operations	Examples
Fuzzy logic (LTN)	$F = \forall x(isCarnivor(s)) \rightarrow (isMammal(x))$ $\{isCarnivor(s){:}[0,1], isMammal(x){:}[1,0]\} \rightarrow F = [1,0]$
Mul and Add (NVSA)	$X_i \in \{+1,-1\}^d \rightarrow (X_i \cdot X_j)/(X_i + X_j)$
Pre-defined objects (NSVQA)	$\texttt{equal\_color:}\ (entry, entry) \rightarrow Boolean$ $\texttt{equal\_integer:}\ (number, number) \rightarrow Boolean$

Efficient NSAI Software

- Encompass broad set of reasoning logics
- Fast
- Memory-efficient
- …

Challenge and Opportunity

	Challenge	Opportunity
Data	Lack of cognitive datasets	Building ImageNet-like NSAI datasets
Model	Nascent integration	Unifying neuro-symbolic-prob models
Software	Modularity & extensibility	Developing efficient SW framework
System	Diverse NSAI algo and HW	Benchmarking NSAI workloads

Challenge and Opportunity

	Challenge	Opportunity
Data	Lack of cognitive datasets	Building ImageNet-like NSAI datasets
Model	Nascent integration	Unifying neuro-symbolic-prob models
Software	Modularity & extensibility	Developing efficient SW framework
System	Diverse NSAI algo and HW	Benchmarking NSAI workloads
Hardware	Heterogeneous kernels	Designing cognitive HW architecture

Towards Understanding the Computational Characteristics of Neuro-Symbolic Workloads

- Very little understanding exists of the computational characteristics of neuro-symbolic AI workloads

✅ Step 1: Categorizing Neuro-Symbolic AI **Algorithm**

✅ Step 2: Benchmarking selecting algorithms on current **Hardware**

✅ Step 3: Our view for Neuro-Symbolic AI **Challenges and Opportunities**

Prof. Tushar Krishna (tushar@ece.gatech.edu)
Prof. Celine Lin (ylin715@gatech.edu)
Prof. Arijit Raychowdhury (arijit.raychowdhury@ece.gatech.edu)

Paper available at https://arxiv.org/pdf/2401.01040.pdf