Neuro-Symbolic Learning and Reasoning: Challenges and Opportunities

Alvaro Velasquez

Program Manager
Information Innovation Office (I2O)
Defense Advanced Research Projects Agency (DARPA)
alvaro.velasquez@darpa.mil

February 7, 2022

Neuro-Symbolic AI

Symbolic Artificial Intelligence

- Prescriptive in nature
- Useful for automation
- Works well in known environments

Neural Machine Learning

- Predictive in nature
- Useful for autonomy
- Works well in unknown environments by learning from data/experience

Neuro-Symbolic AI

Symbolic Artificial Intelligence

- Prescriptive in nature
- Useful for automation
- Works well in known environments

Neural Machine Learning

- Predictive in nature
- Useful for autonomy
- Works well in unknown environments by learning from data/experience

Symbolic Reasoning
Low-capacity sparse representation

Neural Learning
High-capacity dense representation

Neurosymbolic Artificial Intelligence for improved:

- Reliability
- Safety
- Efficiency
- Robustness
- Trust (e.g. Transparency, verifiability, or explainability)
- Generalization to out-of-distribution data

Neuro-Symbolic AI

Symbolic Artificial Intelligence

- Prescriptive in nature
- Useful for automation
- Works well in known environments

Neural Machine Learning

- Predictive in nature
- Useful for autonomy
- Works well in unknown environments by learning from data/experience

Symbolic Reasoning

Low-capacity sparse representation

Neural Learning

High-capacity dense representation

Physics Inspired Neural Networks

Data

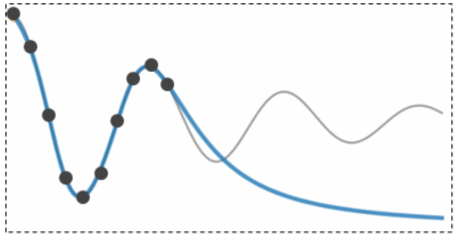

Neural network fitting to data

Symbolic knowledge

$$m\frac{d^2u}{dx^2} + \mu\frac{du}{dx} + ku = 0$$

Physics of 1D damped harmonic oscillator (parameters to be learned)

Results

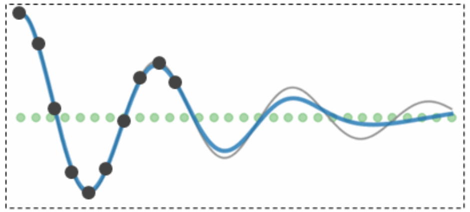

Source: https://benmoseley.blog/

Damping oscillator model included in loss function for training, demonstrates better generalization

The Need for Safe AI: A Quotidian Perspective

"The car assumed that the bus would yield when it attempted to merge back into traffic"

[1] A Google self-driving car caused a crash for the first time. http://www.theverge.com/2016/2/29/11134344/google-self-driving-car-crash-report. (2016).

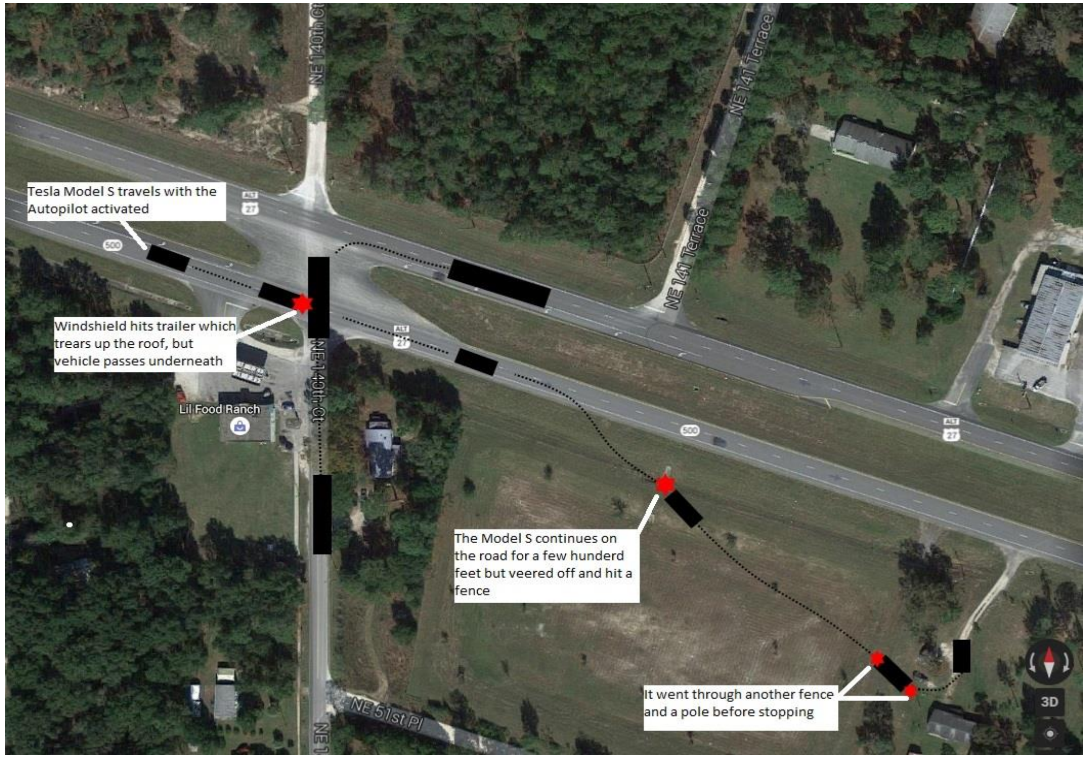

"The camera failed to recognize the white truck against a bright sky"

[2] Understanding the fatal Tesla accident on Autopilotand the NHTSA probe. https://electrek.co/2016/07/01/understanding-fatal-tesla-accident-autopilot-nhtsa-probe/. (2016).

The Need for Safe AI: Neuro-Symbolic Integration

"The car assumed that the bus would yield when it attempted to merge back into traffic"

[1] A Google self-driving car caused a crash for the first time. http://www.theverge.com/2016/2/29/11134344/google-self-driving-car-crash-report. (2016).

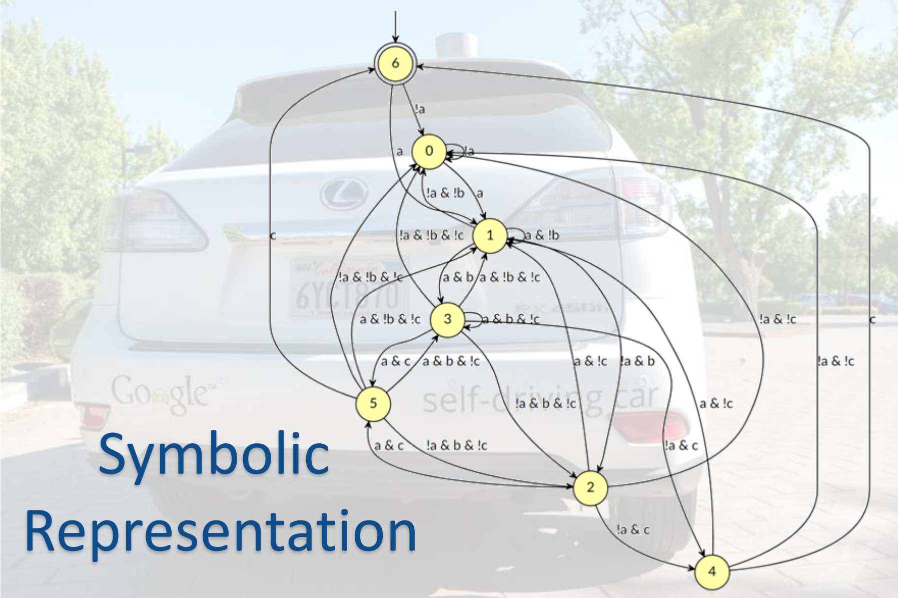

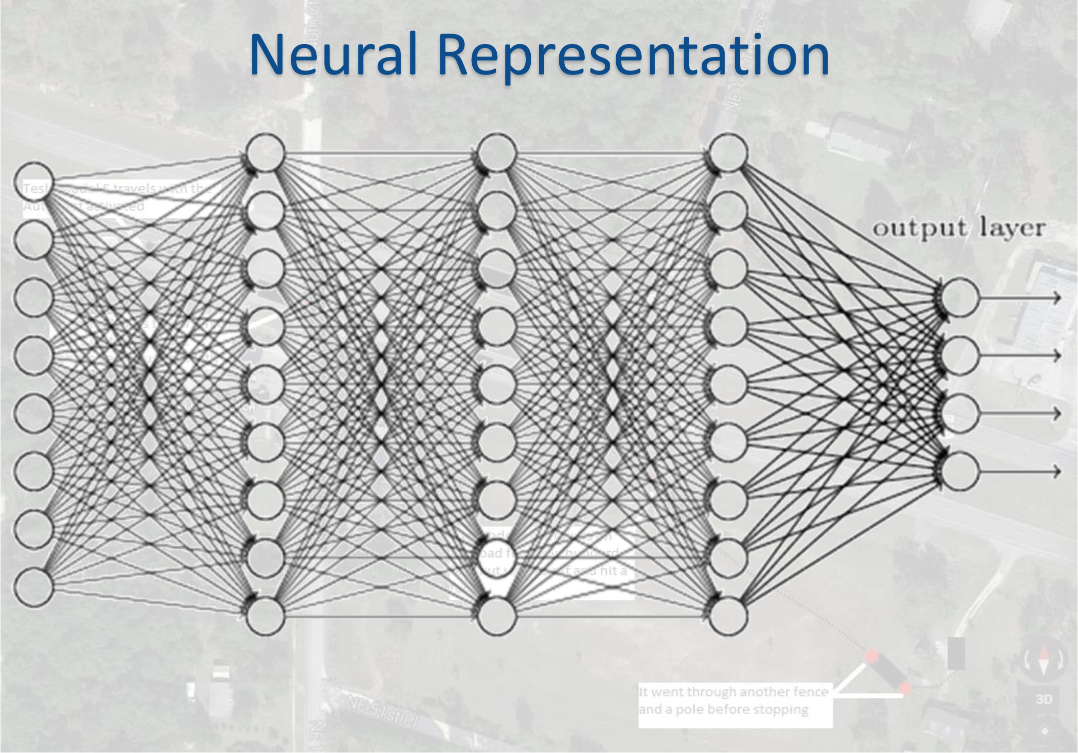

"The camera failed to recognize the white truck against a bright sky“

[2] Understanding the fatal Tesla accident on Autopilotand the NHTSA probe. https://electrek.co/2016/07/01/understanding-fatal-tesla-accident-autopilot-nhtsa-probe/. (2016).

The Need for Trusted AI: A Government Perspective

- "… [must] reduce barriers to the safe testing and deployment of AI technologies…"
- "The United States must foster public trust and confidence in AI technologies…"
- "Ensure that technical standards minimize vulnerability to attacks from malicious actors…"
- "…[identify] safety and security concerns, including those related to the association or compilation of data and models…"

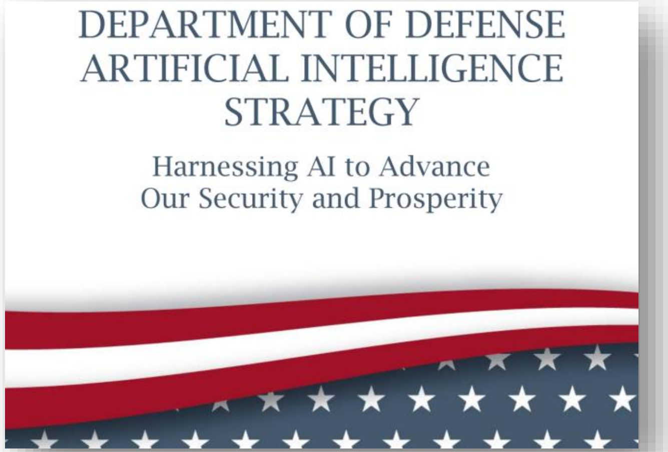

- "… [we must] lead the world in the development and adoption of transformative defense AI solutions that are safe, ethical, and secure."
- "The speed and scale of the change required are daunting…"

- "…[adaptive] systems will outgrow their initial verification and validation and will require more dynamic methods…"
- "…the [autonomous] machine must be auditable—in other words, be able to preserve and communicate an immutable, comprehensive record of the reasoning behind its decisions…"

- "We must anticipate how competitors and adversaries will employ new operational concepts and technologies…"

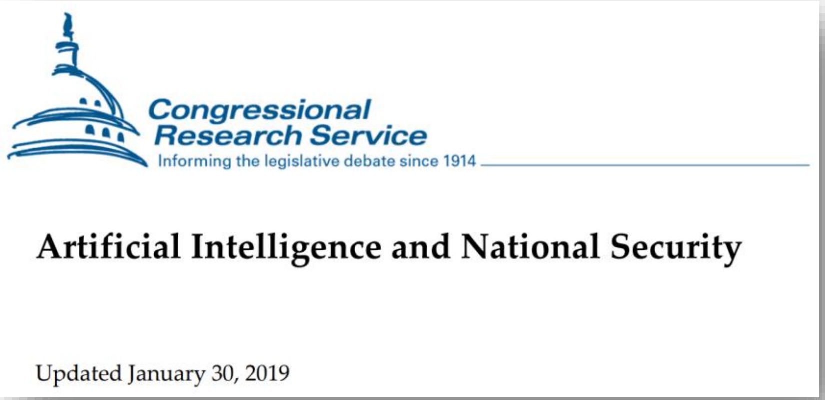

- "Increasing explainability will be key to humans building appropriate levels of trust in AI systems."
- "These [data] vulnerabilities highlight the need for robust data security, cybersecurity, and testing and evaluation processes…"

The Need for Efficient AI: Successes in Autonomy

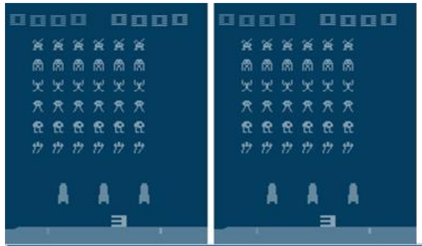

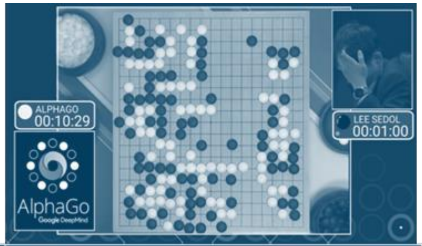

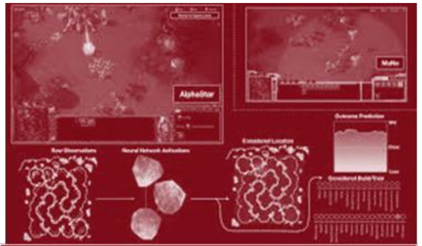

RL with DNN learns to play Atari Games and Tetris (2013)	AlphaGo beats Lee Sedol (2016); AlphaZero beats AllphaGo (2017)	Libratus defeats four top players in 120,000 hands no-limit Texas hold'em (2017)	DOTA bot beats champions in a more "tactical" strategy game (2017)	DeepMind AlphaStar beats "NaMa" and "TLO" in Starcraft II: a high dimensional RTS game (2019)
1v1	1v1	1v1	1v1	1v1
Reinforcement Learning (RL)	Ensemble Supervised Learning and/or RL	Game Theory + Deep RL	RL with Hand Feature Encoding	Supervised Learning + RL
Perfect Information	Perfect Information	Imperfect Information	Imperfect Information	Imperfect Information

The Need for Efficient AI: AlphaGo

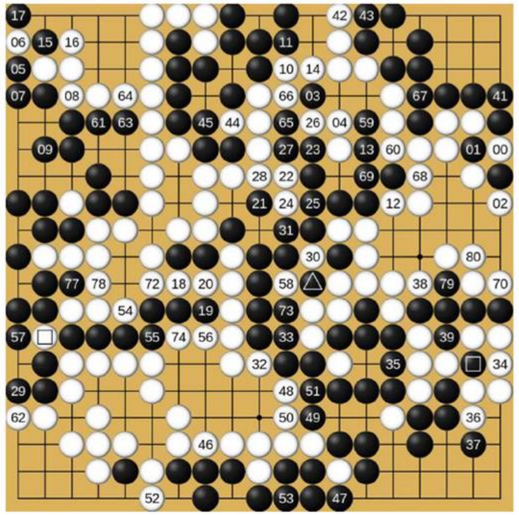

AlphaGo

$> 10,000,000$ games played to match human performance!

Silver, David, et al. "Mastering the game of Go with deep neural networks and tree search." *nature* 529.7587 (2016): 484.

The Need for Efficient AI: OpenAI Five

	OPENAI 1V1 BOT	OPENAI FIVE
CPUs	60,000 CPU cores on Azure	128,000 preemptible CPU cores on GCP
GPUs	256 K80 GPUs on Azure	256 P100 GPUs on GCP
Experience collected	~300 years per day	~180 years per day (~900 years per day counting each hero separately)
Size of observation	~3.3 kB	~36.8 kB
Observations per second of gameplay	10	7.5
Batch size	8,388,608 observations	1,048,576 observations
Batches per minute	~20	~60

OpenAI, "OpenAI Five", https://blog.openai.com/openai-five/, 2018

The OpenAI Five agents used to defeat world-class players at the game Defense Of The Ancients (DOTA) collected 900 years of real-time data per day for a total of 45000 years worth of data.

The Need for Efficient AI: AlphaStar

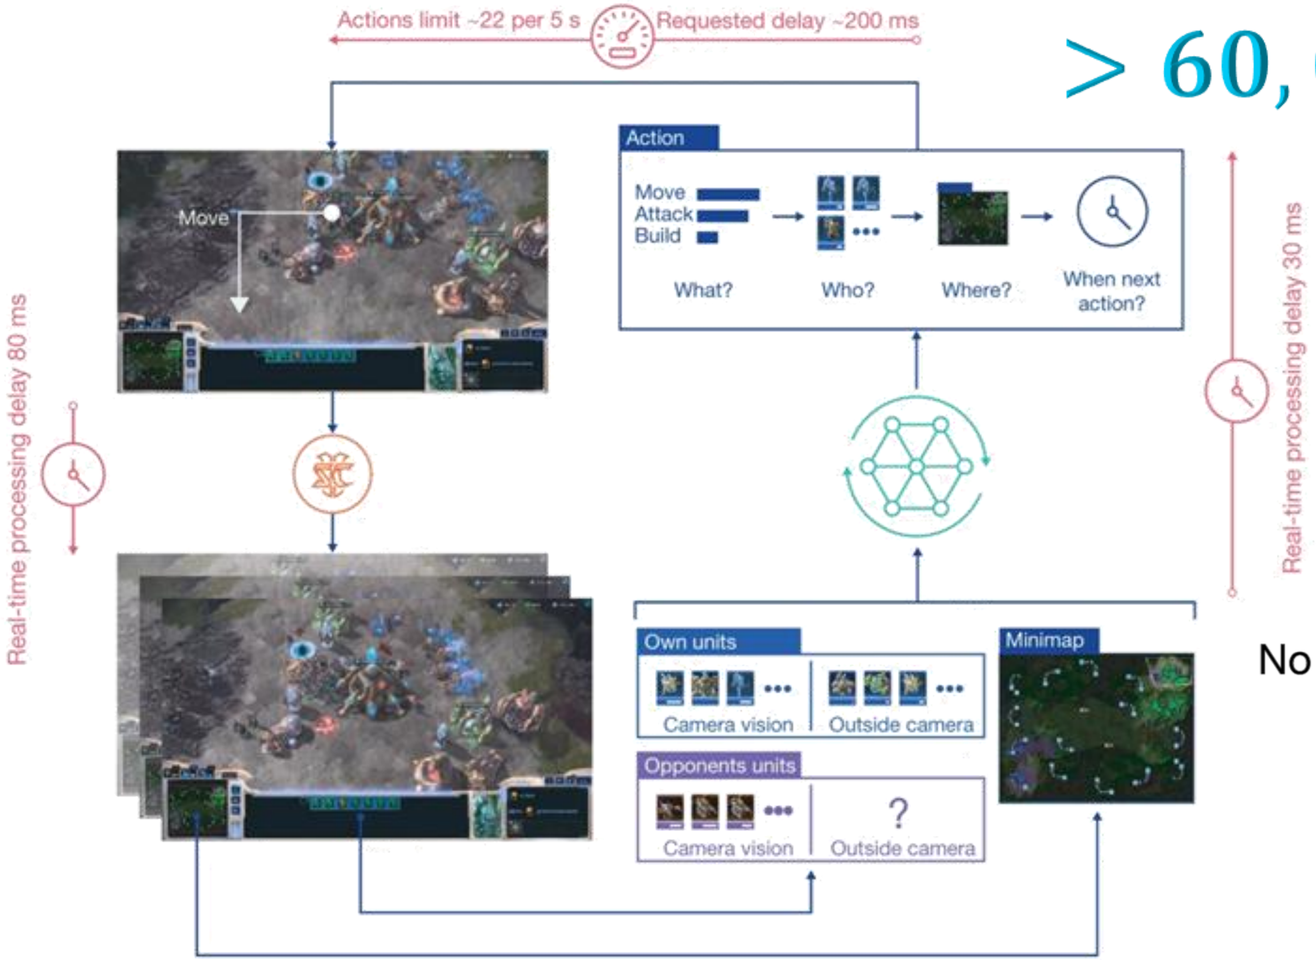

> 60, 000 years of games played!

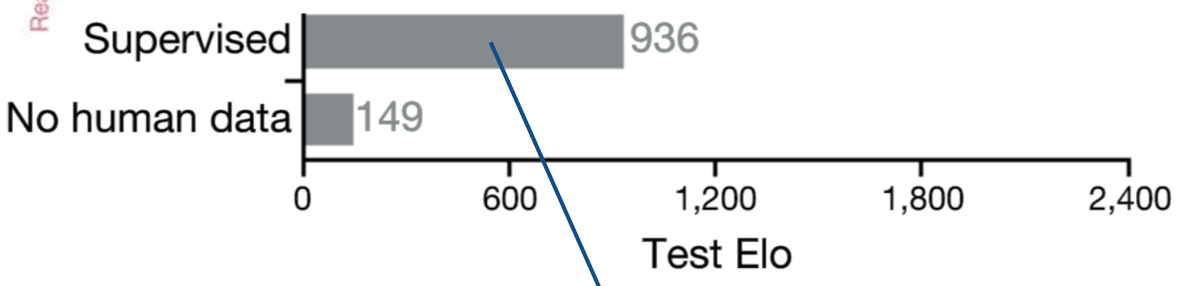

971k replays from top 22% of humans

The Need for Efficient AI: Slow High-Fidelity Platforms

AFSIM

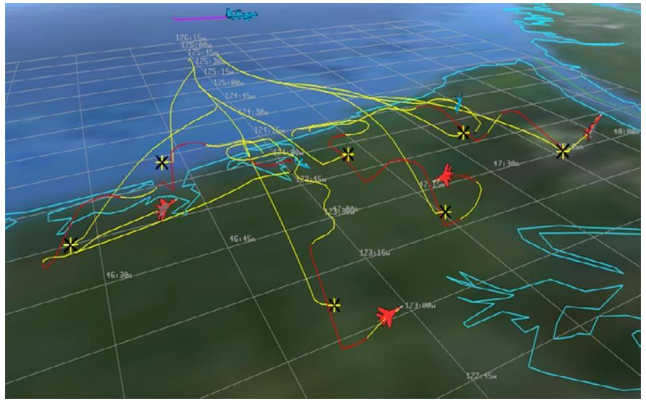

Runtime:
Minutes per execution

OPSIM

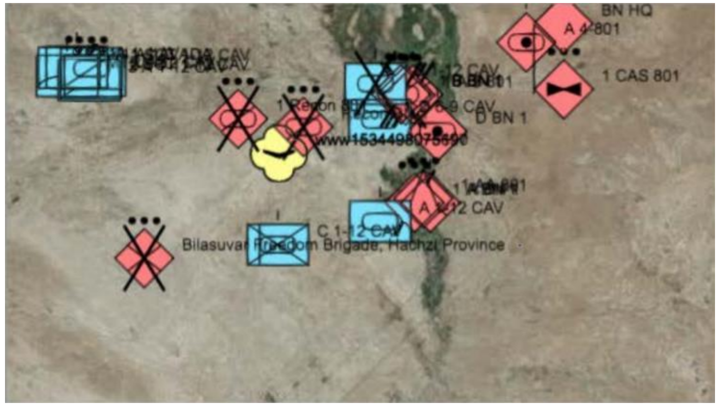

Runtime:
Minutes per execution

AWSIM

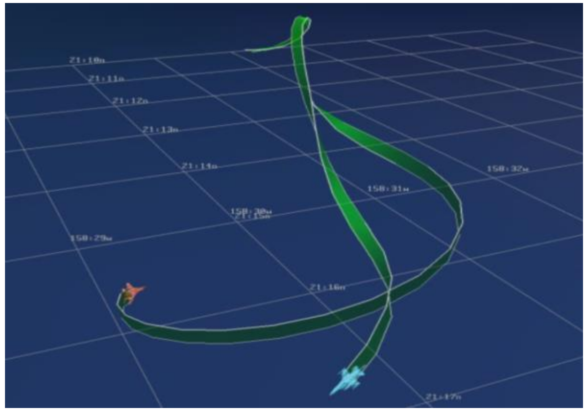

Runtime:
Hours per execution

10 million executions → years of runtime

The Need for Reliable AI: Sparse Rewards

Reward based on mission success
→ Sparse reinforcement

Reward shaping: Introduce artificial frequent reinforcement

The Need for Reliable AI: Reward Hacking

Reward based on mission success → Sparse reinforcement

Reward shaping can lead to reward hacking

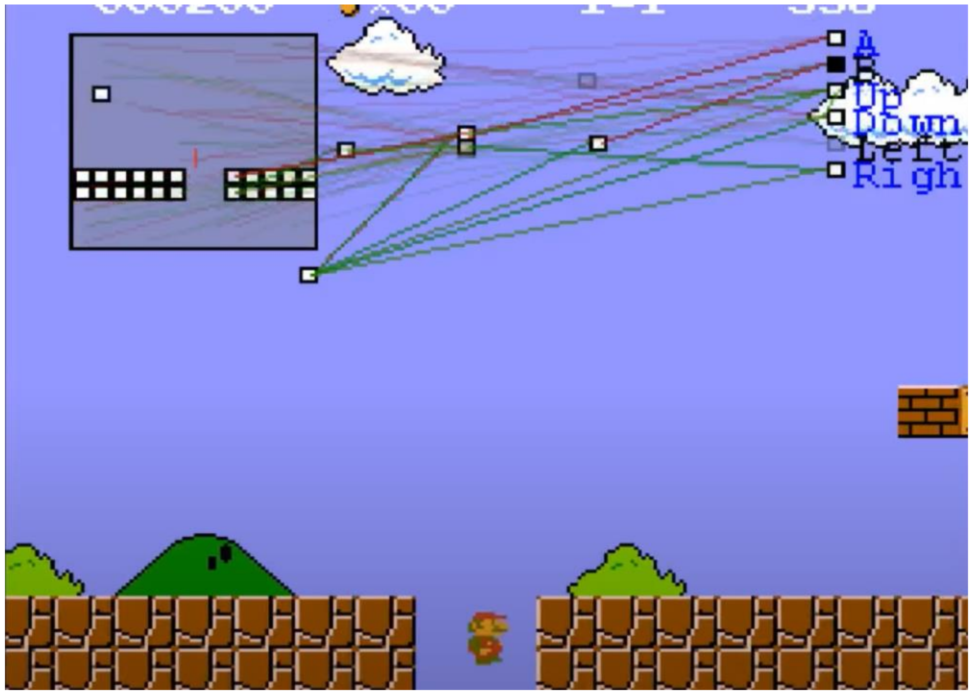

The Need for Robust AI: Adversarial Attacks

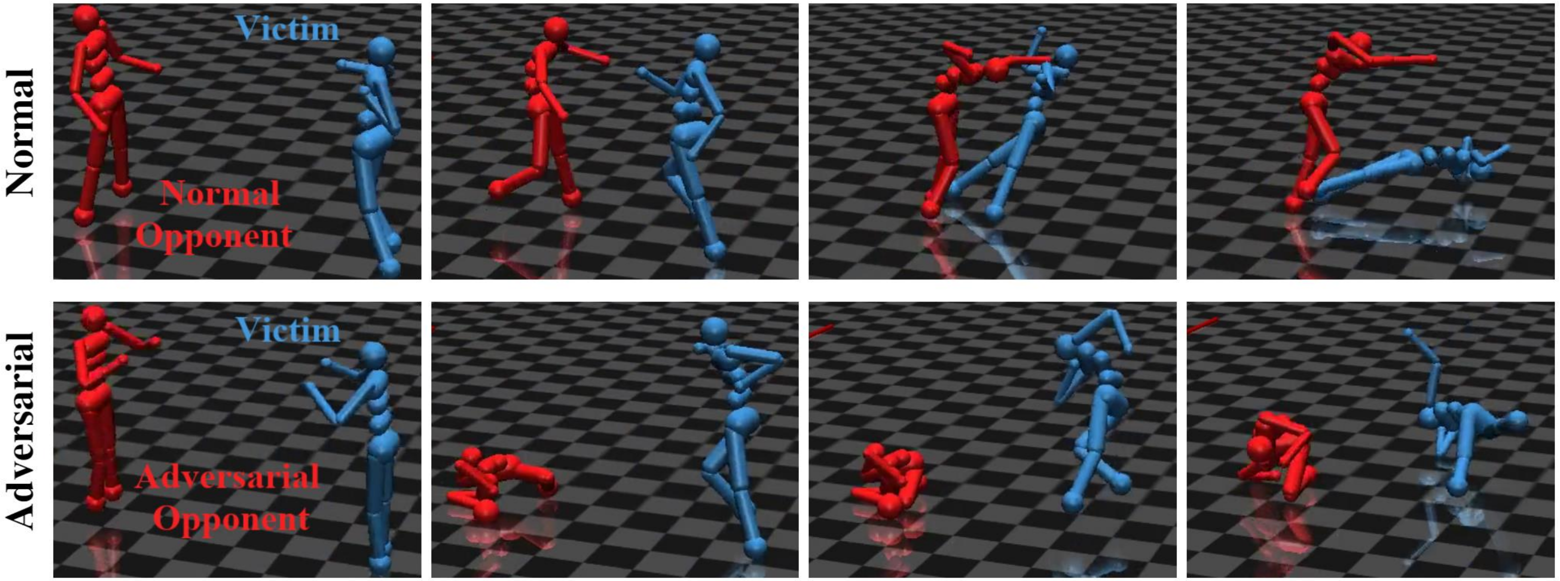

Adversarial Deep Reinforcement Learning

Adversary (in red) wins 86% of times by simply producing unexpected observations (falling down) for victim (in blue) in a “You Shall not Pass” game

Gleave, Adam, et al. "Adversarial policies: Attacking deep reinforcement learning." *arXiv preprint arXiv:1905.10615* (2019).

Opportunity: Early Success in Neuro-Symbolic AI

100x less data for classification [1]

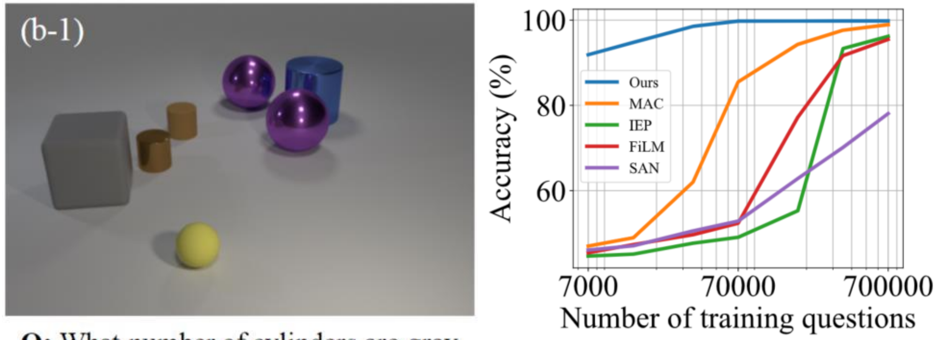

Q: What number of cylinders are gray objects or tiny brown matte objects?

Interpretable classification [2]

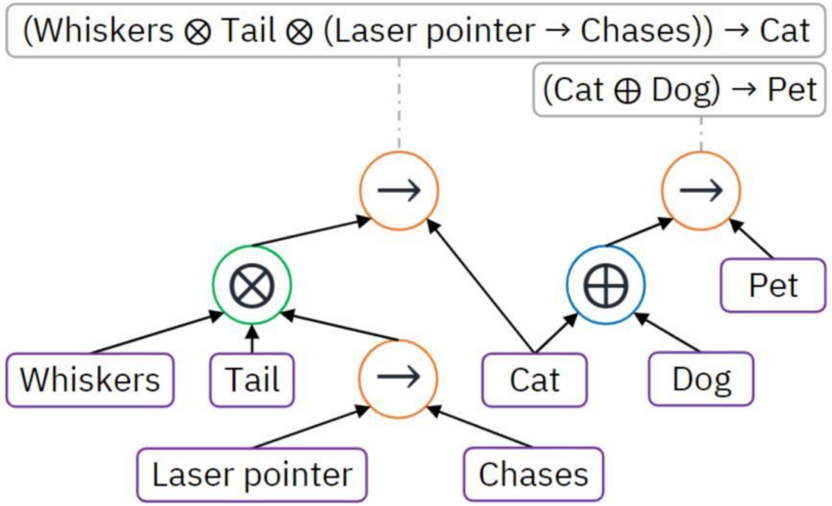

(a) The LNN graph structure reflects the formulae it represents.

Logically consistent learning [3]

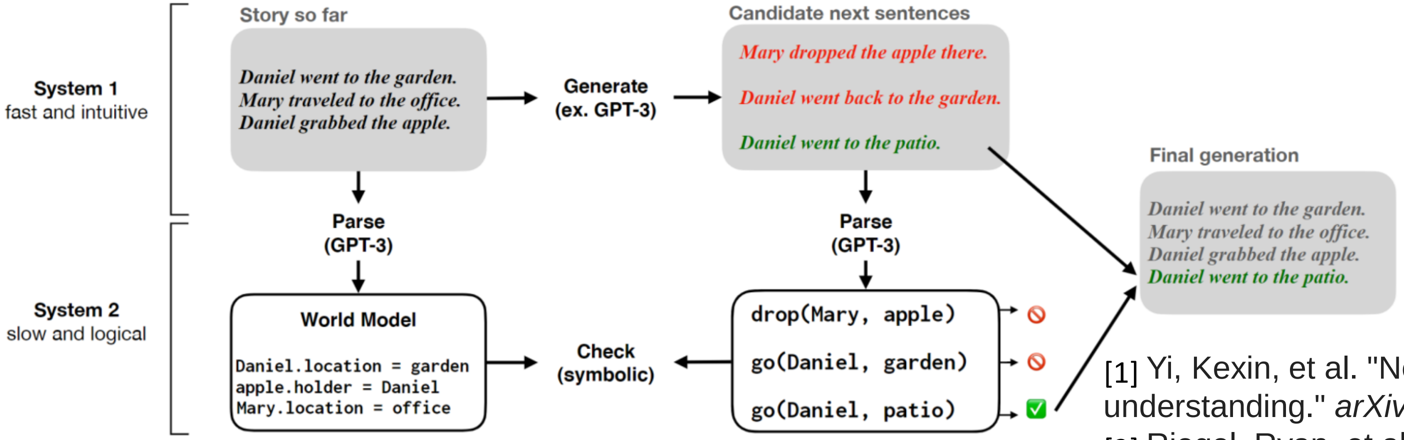

[1] Yi, Kexin, et al. "Neural-symbolic vqa: Disentangling reasoning from vision and language understanding." *arXiv preprint arXiv:1810.02338* (2018).
[2] Riegel, Ryan, et al. "Logical neural networks." *arXiv preprint arXiv:2006.13155* (2020).
[3] Nye, Maxwell, et al. "Improving coherence and consistency in neural sequence models with dual-system, neuro-symbolic reasoning." *Advances in Neural Information Processing Systems* 34 (2021).

Opportunity: Early Success in Neuro-Symbolic Autonomy

Sparse symbolic reward specifications

Automaton-Guided Tree Search (AGTS)

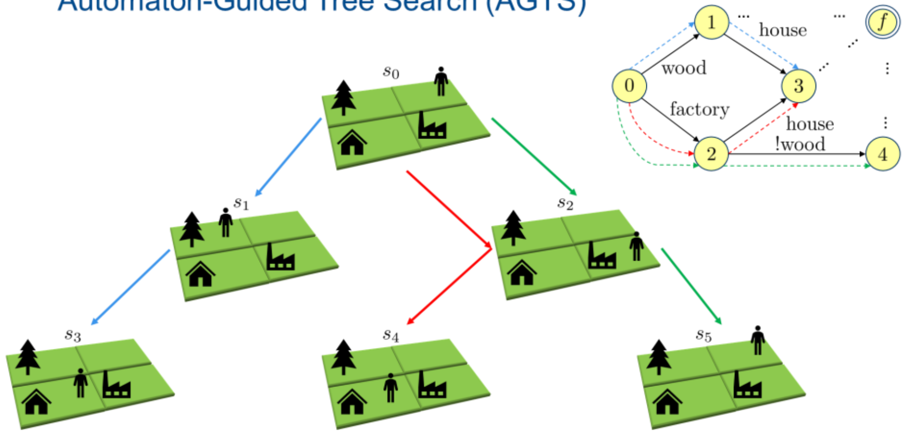

Velasquez, Alvaro, et al. "Dynamic Automaton-Guided Reward Shaping for Monte Carlo Tree Search." *AAAI* 2021.

Learning symbolic representations from observations

Hasanbeig, Mohammad hosein, et al. "DeepSynth: Automata synthesis for automatic task segmentation in deep reinforcement learning." *AAAI 2021.*

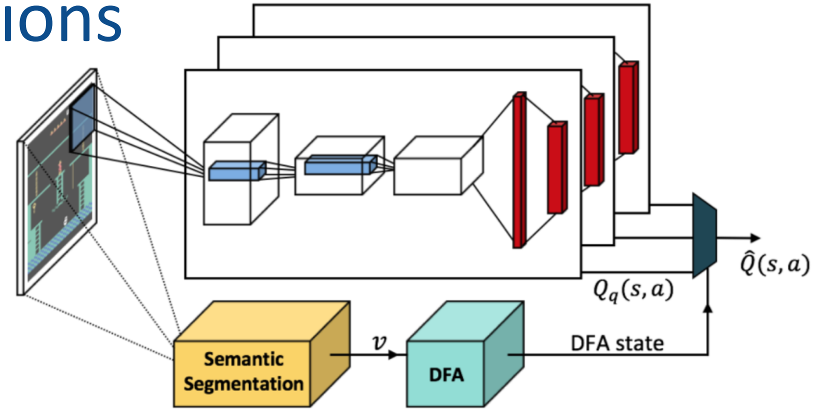

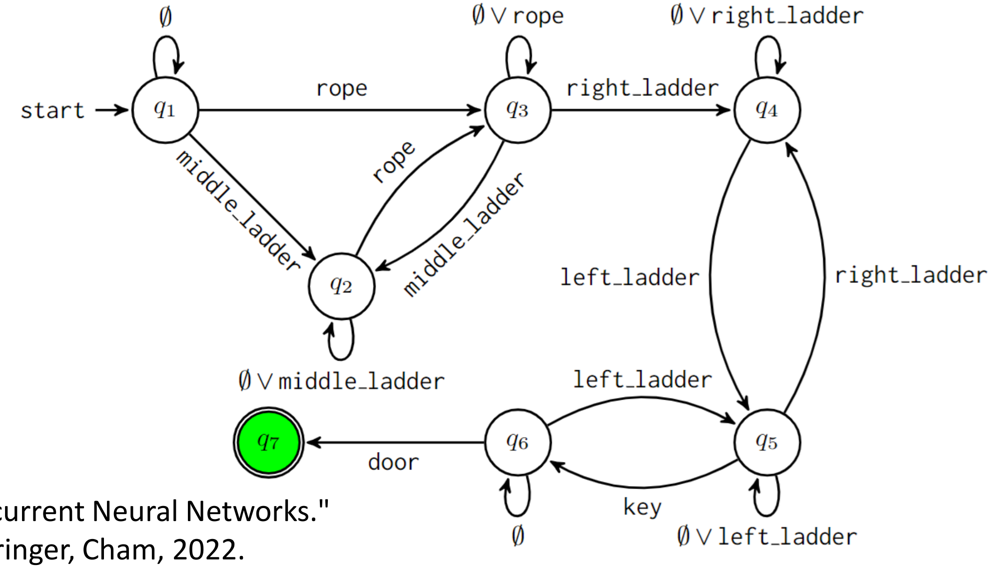

Learning symbolic representations from NN's

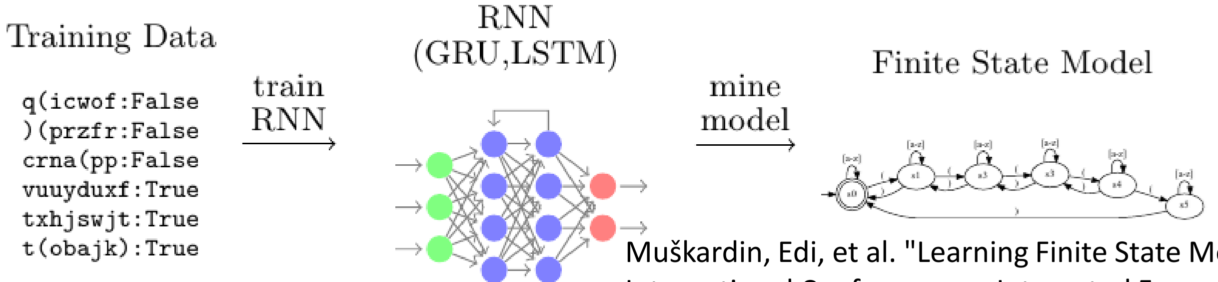

Muškardin, Edi, et al. "Learning Finite State Models from Recurrent Neural Networks." International Conference on Integrated Formal Methods. Springer, Cham, 2022.

From Neurons to Symbols

- Program synthesis using NNs.

Trivedi, Dweep, et al. "Learning to synthesize programs as interpretable and generalizable policies." *Advances in neural information processing systems* 34 (2021): 25146-25163.

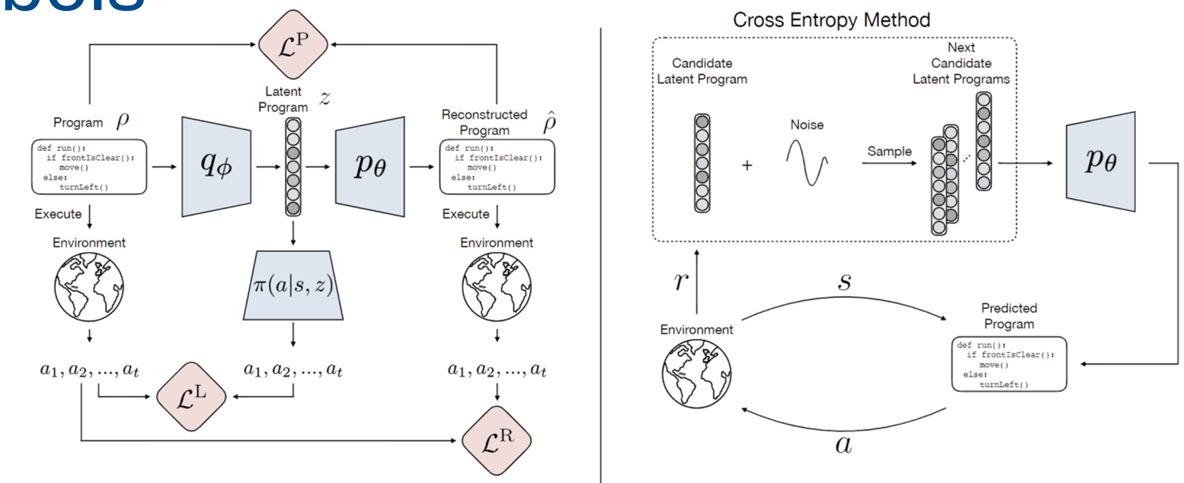

(a) Learning Program Embedding Stage (b) Latent Program Search Stage

- Symbolic finite-state controllers from RNNs.

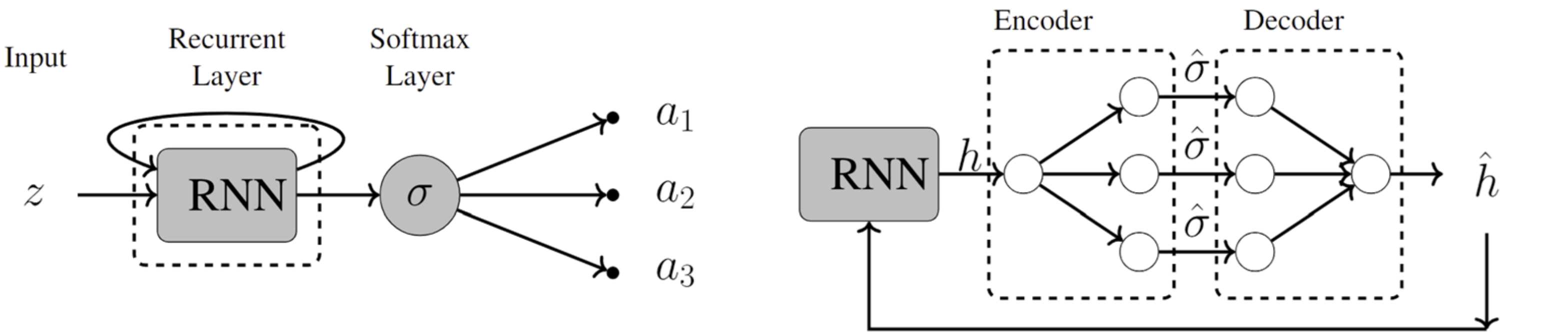

(a) RNN-based Policy $\hat{\pi}$.

(b) RNN block and associated QBN of $B_h = 3$ with quantized activation $\hat{\sigma}: \mathbb{R} \rightarrow \{-1, 0, 1\}$.

Carr, Steven, Nils Jansen, and Ufuk Topcu. "Task-aware verifiable RNN-based policies for partially observable Markov decision processes." Journal of Artificial Intelligence Research 72 (2021): 819-847.

From Symbols to Neurons

- Differentiable logics

$$\mathcal{L}_c(t_1 \leq t_2, i) = max(t_1(i) - t_2(i), 0)$$
$$\mathcal{L}_c(t_1 \neq t_2, i) = \zeta[t_1(i) = t_2(i)]$$
$$\mathcal{L}_c(t_1 < t_2, i) = \mathcal{L}_c(t_1 \leq t_2 \wedge t_1 \neq t_2, i)$$
$$\mathcal{L}_c(t_1 = t_2, i) = \mathcal{L}_c(t_1 \leq t_2 \wedge t_1 \geq t_2, i)$$

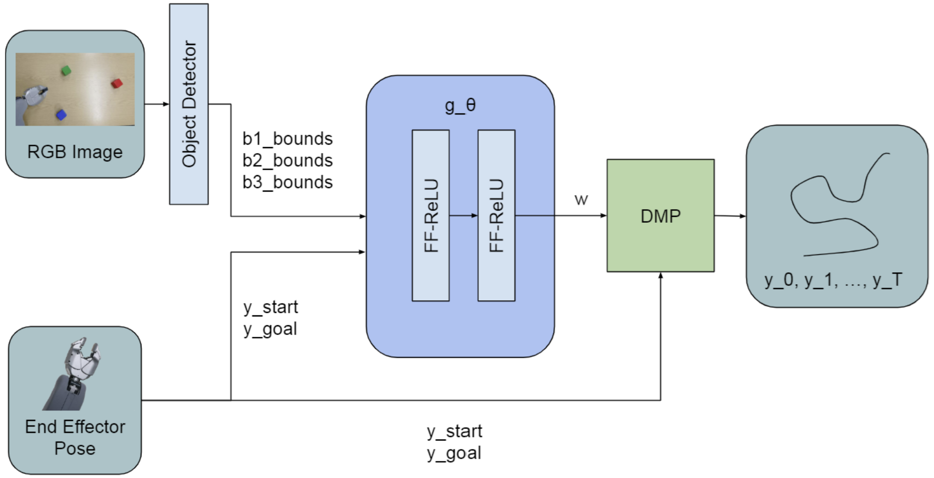

Innes, Craig, and Ram Ramamoorthy. "Elaborating on Learned Demonstrations with Temporal Logic Specifications." *Robotics: Science and Systems 2020*. 2020.

- Discrete optimization using neural networks

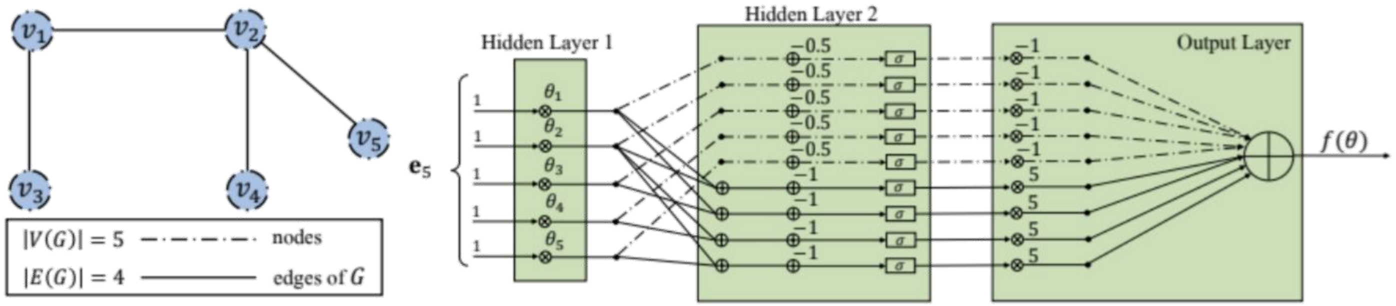

$$f(\theta) = -\sum_{v \in V} \sigma(\theta_v - 1/2) + n \sum_{(u,v) \in E} \sigma(\theta_u + \theta_v - 1)$$

Figure 2: An example of graph $G = (V = \{v_1, v_2, v_3, v_4, v_5\}, E = \{(v_1, v_2), (v_1, v_3), (v_2, v_4), (v_2, v_5)\})$ and its dNN construction f for the MIS problem.

Alkhouri, Ismail R., George K. Atia, and Alvaro Velasquez. "A differentiable approach to the maximum independent set problem using dataless neural networks." *Neural Networks* (2022).

Of Neurons and Symbols

Yu, Jiahui, et al. "Scaling autoregressive models for content-rich text-to-image generation." *arXiv preprint arXiv:2206.10789* (2022).

Neurosymbolic Representation Learning and Reasoning

Neural Representation ⟷ Symbolic Representation

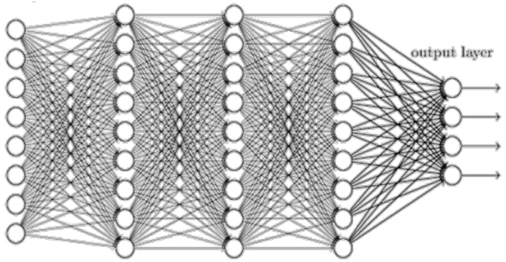

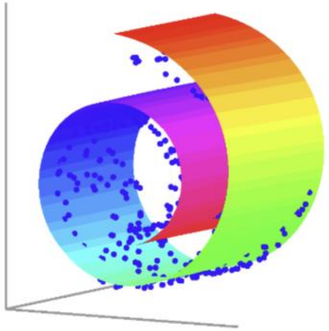

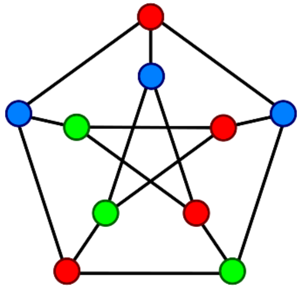

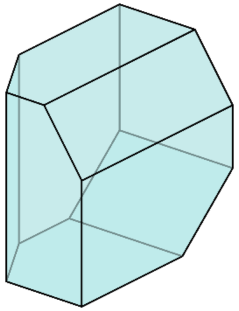

$$\mathcal{L}_c(t_1 \leq t_2, i) = max(t_1(i) - t_2(i), 0)$$
$$\mathcal{L}_c(t_1 \neq t_2, i) = \zeta[t_1(i) = t_2(i)]$$
$$\mathcal{L}_c(t_1 < t_2, i) = \mathcal{L}_c(t_1 \leq t_2 \wedge t_1 \neq t_2, i)$$
$$\mathcal{L}_c(t_1 = t_2, i) = \mathcal{L}_c(t_1 \leq t_2 \wedge t_1 \geq t_2, i)$$

Emerging Architectures

Prompt: “A punk rock squirrel in a studded leather jacket shouting into a microphone while standing on a stump and holding a beer on dark stage. dslr photo.”

Prompt: “A plate that has no bananas on it. There is a glass without orange juice next to it.”

Conclusion and Avenues of Inquiry

1. The integration of symbolic reasoning and neural pattern recognition has the potential to significantly improve efficiency, robustness, and trust in deep learning systems.

2. How can we make progress toward solving the variable binding problem of learning semantically meaningful symbols? How do we reason about neurons and symbols simultaneously?

3. What does the ImageNET moment look like for neuro-symbolic AI?

4. How do we measure the quality of neuro-symbolic AI? The need for metrics.

5. Powerful emerging architectures may pave the way for neuro-symbolic AI.