From Logic to Neurosymbolic AI

Luís Lamb

UFRGS, Brazil

MIT Sloan School of Management

twitter @luislamb

.inf UFRGS

Summary

A bit of evolution and history of the field (in 2min)

integrating learning and reasoning (in 2min)

Neurosymbolic AI & computing:
Vardi, Inspiration, and perspective (in 3min)

References & Afterword (in 1min)

Acknowledgements

Thanks to Moshe Vardi,
Artur Garcez, Dov Gabbay,
Rafael Borges, Krysia Broda, Alessandra
Russo, Marcelo Prates, Pedro Avelar,
Anderson Tavares, João Flach, Marcio
Nicolau, Marco Gori, and many others.

Or 20 years of
Neurosymbolic AI

Neurosymbolic AI: The 3^{rd} Wave

Artur d'Avila Garcez[1] and Luís C. Lamb[2]
[1] City, University of London, UK
a.garcez@city.ac.uk
[2] Federal University of Rio Grande do Sul, Brazil
luislamb@acm.org

December, 2020

https://arxiv.org/pdf/2012.05876.pdf

Luis Lamb @luislamb · Feb 13

It seems that the literature shows that Turing, McCulloch & Pitts, Kleene, von Neumann have shown interest in understanding the interplay of logic and neural networks. @frossi_t @AvilaGarcez @GaryMarcus @vardi #neurosymbolicAI

3 5 31

Gary Marcus @GaryMarcus · Feb 13

for sure; i am working on an essay that mentions this. the unfortunate division between schools didn't seem to arise until some time the 50s

2 1 12

2. A SCHEMATIC VIEW OF AUTOMATA

2.1 Logics and Automata

It has been pointed out by A. M. Turing [5] in 1937 and by W. S. McCulloch and W. Pitts [2] in 1943 that effectively constructive logics, that is, intuitionistic logics, can be best studied in terms of automata. Thus logical propositions can be represented as electrical networks or (idealized) nervous systems. Whereas logical propositions are built up by combining certain primitive symbols, networks are formed by connecting basic components, such as relays in electrical circuits and neurons in the nervous system. A logical proposition is then represented as a "black box" which has a finite number of inputs (wires or nerve bundles) and a finite number of outputs. The operation performed by the box is determined by the rules defining which inputs, when stimulated, cause responses in which outputs, just as a propositional function is determined by its values for all possible assignments of values to its variables.

There is one important difference between ordinary logic and the automata which represent it. Time never occurs in logic, but every network or nervous system has a definite time lag between the input signal and the output response. A definite temporal sequence is always inherent in the operation of such a real system. This is not entirely a disadvantage. For example, it prevents the occurence of various kinds of more or less overt vicious circles (related to "non-constructivity", "impredicativity", and the like) which represent a major class of dangers in modern logical systems. It should be emphasized again, however, that the representative automaton contains more than the content of the logical proposition which it symbolizes — to be precise, it embodies a definite time lag.

2. A SCHEMATIC VIEW OF AUTOMATA

2.1 Logics and Automata

It has been pointed out by A. M. Turing [5] in 1937 and by W. S. McCulloch and W. Pitts [2] in 1943 that effectively constructive logics, that is, intuitionistic logics, can be best studied in terms of automata. Thus logical propositions can be represented as electrical networks or (idealized) nervous systems. Whereas logical propositions are built up by combining certain primitive symbols, networks are formed by connecting basic components, such as relays in electrical circuits and neurons in the nervous system. A logical proposition is then represented as a "black box" which has a finite number of inputs (wires or nerve bundles) and a finite number of outputs. The operation performed by the box is determined by the rules defining which inputs, when stimulated, cause responses in which outputs, just as a propositional function is determined by its values for all possible assignments of values to its variables.

There is one important difference between ordinary logic and the automata which represent it. Time never occurs in logic, but every network or nervous system has a definite time lag between the input signal and the output response. A definite temporal sequence is always inherent in the operation of such a real system. This is not entirely a disadvantage. For example, it prevents the occurence of various kinds of more or less overt vicious circles (related to "non-constructivity", "impredicativity", and the like) which represent a major class of dangers in modern logical systems. It should be emphasized again, however, that the representative automaton contains more than the content of the logical proposition which it symbolizes — to be precise, it embodies a definite time lag.

Vardi, Wolper, 1980s...

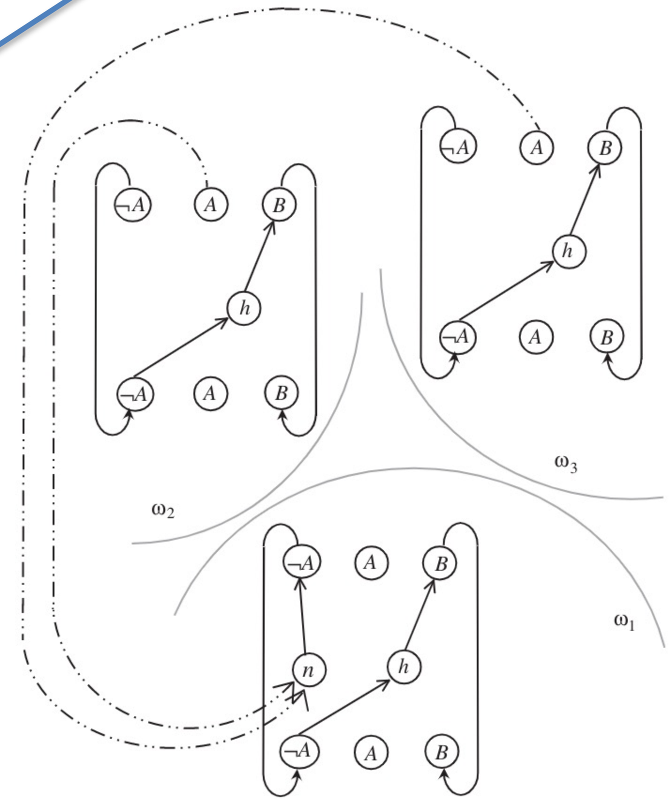

Fig. 5. Representing intuitionistic negation.

Available online at www.sciencedirect.com

SCIENCE DIRECT

Theoretical Computer Science 358 (2006) 34–55

Theoretical Computer Science

www.elsevier.com/locate/tcs

Connectionist computations of intuitionistic reasoning

Artur S. d'Avila Garcez[a,*], Luís C. Lamb[b], Dov M. Gabbay[c]

Neurosymbolic AI Headlines

FEATURES

The future of A.I.: 4 big things to watch for in the next few years

By Luke Dormehl
May 21, 2021

CONTENTS

In a December 2020 publication, researchers Artur d'Avila Garcez and Luis Lamb described neuro-symbolic A.I. as the "third wave" of artificial intelligence. Neuro-symbolic A.I. is not, strictly speaking, totally new. It's more like getting two of the world's greatest rock stars, who once battled at the top of the charts, together to create a supergroup. In this case, the supergroup consists of self-learning neural networks and rule-based symbolic A.I.

"Neural networks and symbolic ideas are really wonderfully complementary to each other," David Cox, director of the MIT-IBM Watson A.I. Lab in Cambridge, Massachusetts, previously told Digital Trends. "Because neural networks give you the answers for getting from the messiness of the real world to a symbolic representation of the world, finding all the correlations within images. Once you've got that symbolic representation, you can do some pretty magical things in terms of reasoning."

The Machine
Making sense of AI

Leading computer scientists debate the next steps for AI in 2021

Written by **Tiernan Ray**, Contributing Writer
on December 14, 2020 | Topic: Artificial Intelligence

A panel talk Friday afternoon brought together AI scholars Gary Marcus, Yoshua Bengio, Daniel Kahneman, Luis Lamb, and moderator Francesca Rossi, for a spirited discussion of where machines and humans differ in their processing of abstract thought, logic, reason and many, many related questions.

NeurIPS 2020

Hybrid artificial intelligence

Cognitive scientist Gary Marcus, who cohosted the debate, reiterated some of the key shortcomings of deep learning, including excessive data requirements, low capacity for transferring knowledge to other domains, opacity, and a lack of reasoning and knowledge representation.

"One of the key questions is to identify the building blocks of AI and how to make AI more trustworthy, explainable, and interpretable," computer scientist Luis Lamb said.

Neurosymbolic AI in Industry

IBM Neuro-Symbolic AI Workshop 2022

feedback

Overview Speakers Organizers

IBM Neuro-Symbolic AI Workshop 2022
Unifying Statistical and Symbolic AI

https://researcher.watson.ibm.com/researcher/view_group.php?id=10897

IBM Research Research Areas Labs Disciplines Blog

Day 1, Session 1: Introduction -- (Replay)

18 January 2022 (08:30 - 10:40 ET)

Time	Topic	Speaker
08:30 ET	**Workshop Introduction** (10 mins) • Opening Words • Motivation and overview	**Lead:** Alexander Gray (IBM) **Speakers:** Francesca Rossi (IBM), Murray Campbell (IBM), Lior Horesh (IBM)
08:40 ET	**Invited talk 1: A Short on the History and Evolution of Neurosymbolic AI** (30 mins)	Luis Lamb (Universidade Federal do Rio Grande do Sul)
09:10 ET	**Neuro-symbolic AI overview** (1 hour + 5 mins QA)	Alexander Gray (IBM)
10:15 ET	**General AI and Interactive fiction** (30 mins + 5 mins QA)	Murray Campbell (IBM)

Day 1, Session 2: Learnable Reasoning -- (Replay)

18 January 2022 (11:30 - 13:40 ET)

Time	Topic	Speaker
11:30 ET	**Learnable Reasoning** (1 hour + 5 mins QA)	Ndivhuwo Makondo (IBM), Hima Karanam (IBM)
12:40 ET	**Invited talk 2: Theory of real-valued logics** (30 mins)	Ron Fagin (IBM)

IBM

IBM Research Research Areas Labs Disciplines Blog

Neuro-Symbolic AI

feedback

Overview Selected Publications Blogs & News Awards & Recognitions

320+ Publications 2018-2020

50+ Projects

10+ University Collaborations

The Neuro-Symbolic AI (NS) initiative aims to conceive a fundamental new methodology for AI, to address the gaps remaining between today's state-of-the-art and the full goals of AI, including AGI. In particular it is aimed at augmenting (and retaining) the strengths of statistical AI (machine learning) with the complementary capabilities of symbolic or classical AI (knowledge and reasoning). It is aimed at a construction of new paradigms rather than superficial synthesis of existing paradigms, and revolution rather than evolution.

Microsoft | **Research** Our research Programs & events Blogs & podcasts

Neuro-Symbolic Program Synthesis

Emilio Parisotto, Abdelrahman Mohamed, Rishabh Singh, Lihong Li, Denny Zhou, Pushmeet Kohli

5th International Conference on Learning Representations (ICLR 2017) | February 2017

What happened from 2006?

In the early 2000s: top ML/AI Conferences - virtually **no papers** which made any kind of use of Artificial Neural Networks

Exceptions: there was a paper on Neural-Symbolic Learning:
Lamb, Borges & d'Avila Garcez, AAAI 2007.

So, how deep learning, a.k.a. deep artificial neural networks developed?

GE Hinton, S Osindero, YW Teh
A fast learning algorithm for deep belief nets.
Neural computation 18 (7), 1527-1554, 2006.
In the same issue:

A. d'Avila Garcez, Luís C. Lamb:
A Connectionist Computational Model for Epistemic and Temporal Reasoning. Neural Computation 18(7): 1711-1738 (2006)

Recent Developments

More attention to neural networks and symbolic AI

2019 - 2022

(1) Fireside conversation: Hinton, LeCun, Bengio; Francesca Rossi & Nobel Laureate Daniel Kahneman on thinking fast and slow and its relation to neural networks and symbol manipulation.

(1) Henry Kautz's The Third AI Summer - Robert S. Engelmore Memorial Award Lecture: taxonomy for neurosymbolic computing

(1) David Cox's, IAAI2020 Invited talk - neurosymbolic AI and its implications on vision, language understanding, machine commonsense, QA, argumentation and XAI.

2020/2022:

(1) AI DEBATE #2: Moving AI Forward.

(2) AAAI2021 Panel on Neurosymbolic AI

(3) IBM Neuro-Symbolic Workshop, Jan. 2022

(4) Dagstuhl Seminars: 2014, 2017, 2019, 2022.

Luis Lamb
@luislamb

@vardi on thinking fast and slow, 2017. Att @jeublanc @AvilaGarcez @frossi_t @pascalhitzler @erichorvitz

 **Moshe Vardi** @vardi · Feb 11, 2020

A slide that I introduced in 2017. @luislamb @GaryMarcus

Logic vs. Machine Learning

Daniel Kahneman, *Thinking, Fast and Slow*, 2011:

- **Machine Learning**: fast thinking, e.g., "Is this a stop sign?"
- **Logic**: slow thinking, e.g., "Do you stop at a stop sign?"

Grand Challenge: Combine logic with machine learning!

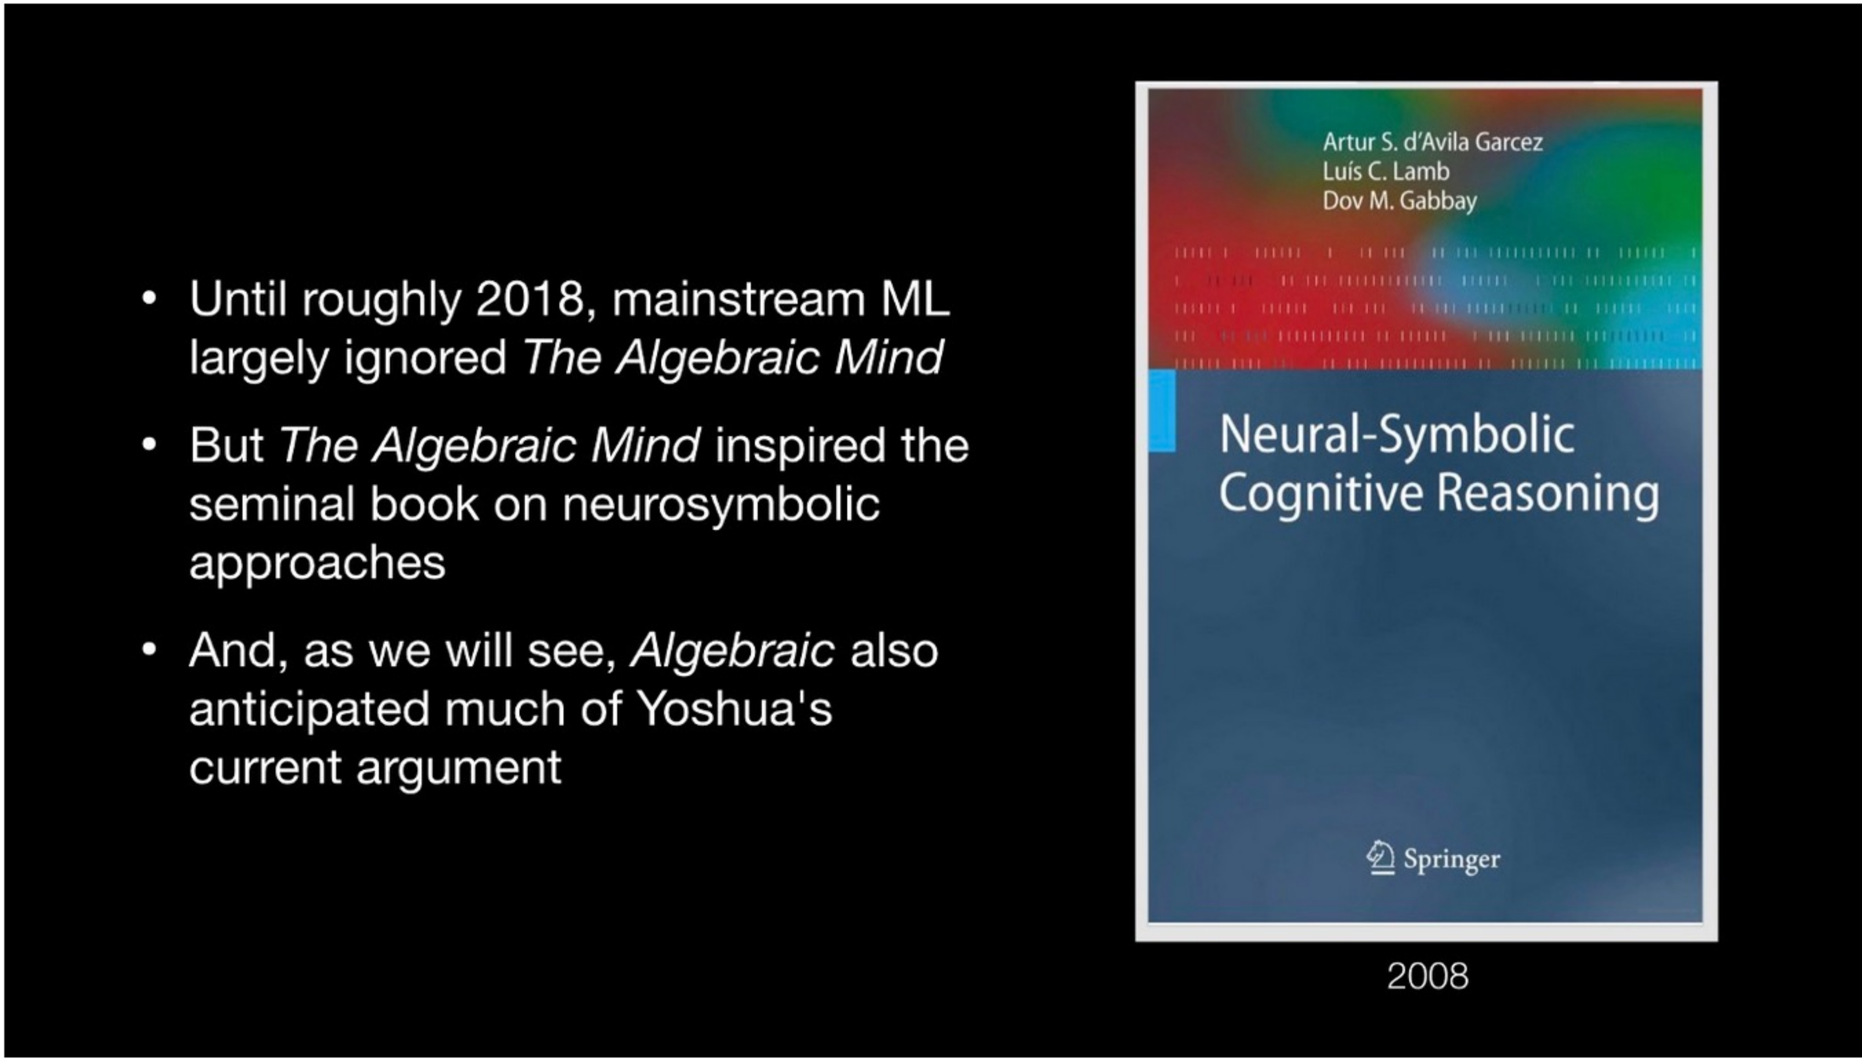

Neural-Symbolic Cognitive Reasoning

GARY MARCUS IN THE GREAT AI DEBATE, MONTREAL
23 DECEMBER 2019.

A Small Step Towards integration: NSAI

Combines logical reasoning and neural learning:

Computer Science Logic + Neural Computation

Neurosymbolic AI: Learning from experience and reasoning about what has been learned from an uncertain environment in a computationally efficient way.

LEARN in order to REASON
(& conversely)

They test hypotheses.

They use abduction, deduction and induction.

Why nonclassical logics?

expressing several reasoning features, allowing for the representation of temporal, epistemic and probabilistic abstractions in computer science and AI.

(Fagin, Halpern, Moses, Vardi, 1995; Halpern 2005)

Classic (and most valuable) Literature:

Ronald Fagin, Joseph Y. Halpern, Yoram Moses, Moshe Y. Vardi: **Reasoning About Knowledge.** MIT Press 1995. **(all here today!)**

Joseph Y. Halpern: **Reasoning about uncertainty.** MIT Press 2005

Conectionist Modal Logics

Modal logic goes beyond propositional reasoning: Gabbay 70s, Vardi, 1996.

A proposition is necessary (*box*) in a possible world (state of affairs) if it is true in all worlds which are possible in **relation** to that world.

A proposition is possible (*diamond*) in a possible world (state of affairs) if it is true in at least one world which is possible in relation to that same world (reference state).

Modalities also used for reasoning about uncertainty (Halpern).

Relational learning/reasoning is notoriously hard.

There is continuous research on the **need for modularity.**

Learning to reason in connectionist models

- **Insight: *assume that (ensembles of) neurons are seen as possible worlds.***
- Propositional Modal Logic => decidable fragment of FOL with two variables (see Gabbay 1970s, **Vardi 1996**).
- Full solution of Muddy Children puzzle and other testbeds.

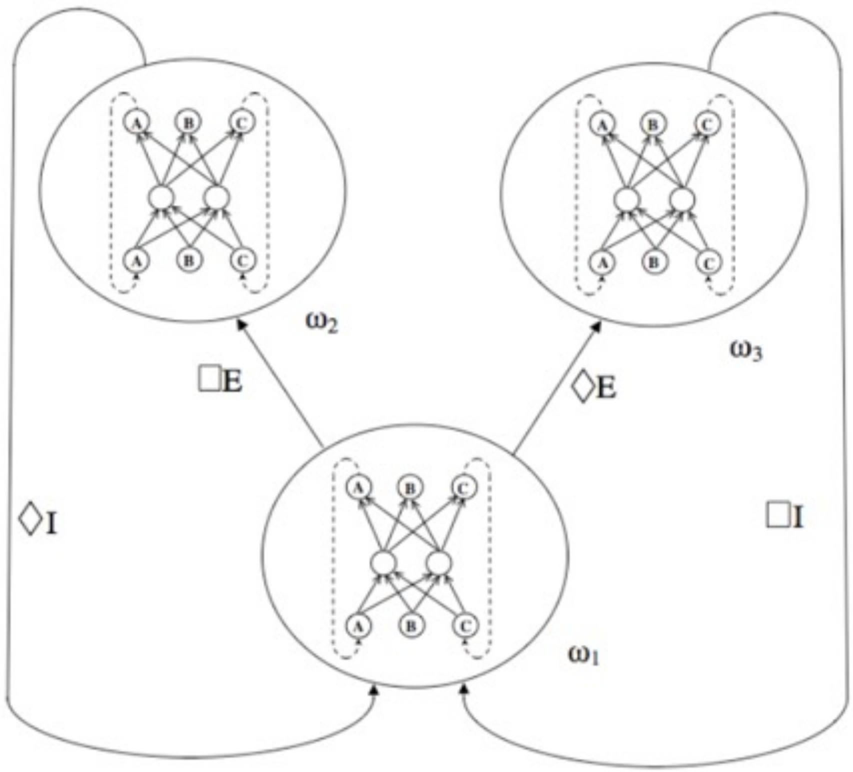

Garcez, Lamb, Gabbay. Connectionist Modal Logic. *Theoretical Computer Science*, 371: 34-53, 2007.
Garcez, Lamb. Connectionist Model for Epistemic and Temporal Reasoning. *Neural Computation,* 18:1711-1738, July 2006

Connectionist Modal Logics: Inference Rules

Say you want to learn the following:

Let $\mathcal{P} = \{\omega_1 : r \rightarrow \Box q;\ \omega_1 : \Diamond s \rightarrow r;\ \omega_2 : s;\ \omega_3 : q \rightarrow \Diamond p;\ \mathcal{R}(\omega_1,\omega_2),\ \mathcal{R}(\omega_1,\omega_3)\}$.

From logic inference to connectionist reasoning.

Table 1
Rules for modality operators

$$\dfrac{\begin{array}{c}[R(\omega, g_\varphi(\omega))] \\ \vdots \\ g_\varphi(\omega) : \varphi\end{array}}{\omega : \Box\varphi}\ \Box I \qquad \dfrac{\omega_1 : \Box\varphi,\ R(\omega_1, \omega_2)}{\omega_2 : \varphi}\ \Box E$$

$$\dfrac{\omega : \Diamond\varphi}{f_\varphi(\omega) : \varphi,\ R(\omega, f_\varphi(\omega))}\ \Diamond E \qquad \dfrac{\omega_2 : \varphi,\ R(\omega_1, \omega_2)}{\omega_1 : \Diamond\varphi}\ \Diamond I$$

Let $\mathcal{P} = \{\omega_1 : r \rightarrow \Box q;\ \omega_1 : \Diamond s \rightarrow r;\ \omega_2 : s;\ \omega_3 : q \rightarrow \Diamond p;\ \mathcal{R}(\omega_1,\omega_2),\ \mathcal{R}(\omega_1,\omega_3)\}$.

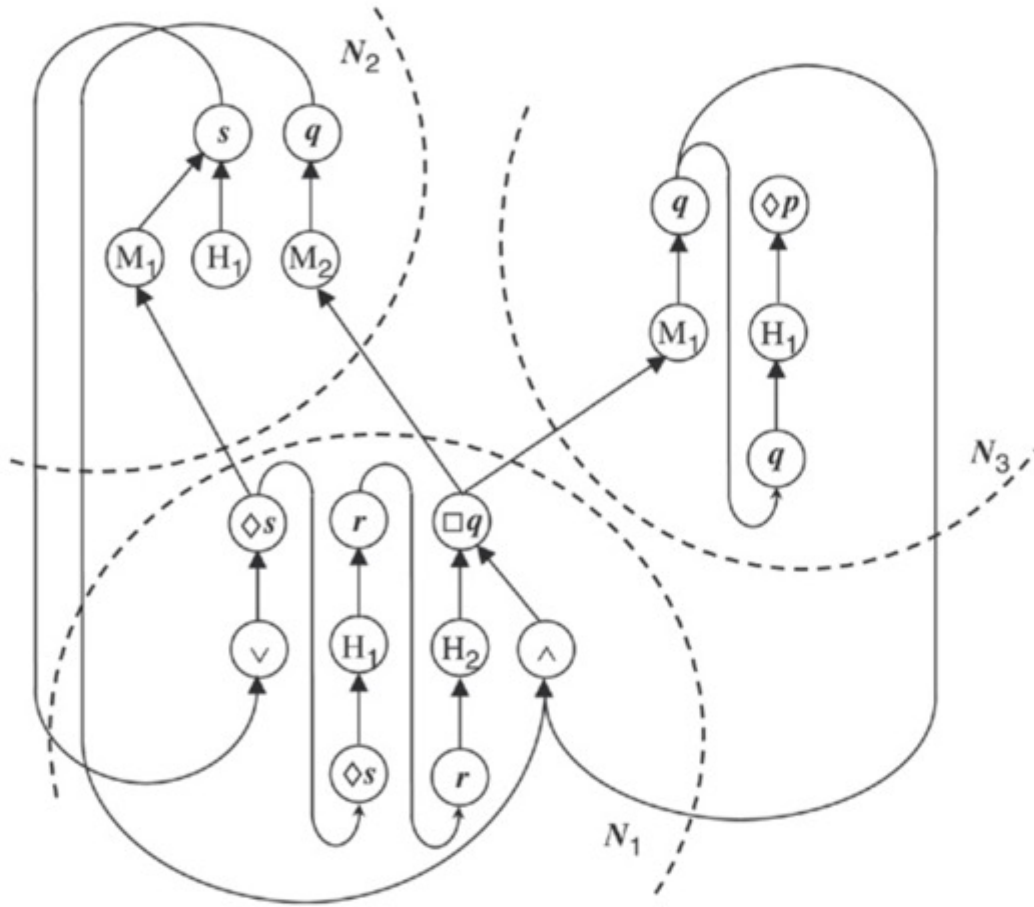

Fig. 5. The ensemble of networks $\{\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3\}$ that represents $\mathcal{P}$.

Algorithm 2: Translation of $\bullet$-based programs

$\bullet$-based_Translation$(\mathcal{P})$

Define $\frac{max_{\mathcal{P}}(k,\mu)-1}{max_{\mathcal{P}}(k,\mu)+1} \leq A_{min} < 1$;

Define $W \geq \frac{ln(1+A_{min})-ln(1-A_{min})}{max_{\mathcal{P}}(k,\mu)(A_{min}-1)+A_{min}+1} \cdot \frac{2}{\beta}$;

foreach $C_l \in Clauses(\mathcal{P})$ **do**

$AddHiddenNeuron(\mathcal{N}, h_l)$;

foreach $\alpha \in body(C_l)$ **do**

if $in_\alpha \notin Neurons(\mathcal{N})$ **then**

$AddInputNeuron(\mathcal{N}, in_\alpha)$;

$ActivationFunction(in_\alpha) \leftarrow g(x)$;

$AddLink(\mathcal{N}, in_\alpha, h_l, W)$;

end

foreach $\sim\alpha \in body(C_l)$ **do**

if $in_\alpha \notin Neurons(\mathcal{N})$ **then**

$AddInputNeuron(\mathcal{N}, in_\alpha)$;

$ActivationFunction(in_\alpha) \leftarrow g(x)$;

$AddLink(\mathcal{N}, in_\alpha, h_l, -W)$;

end

$\alpha \leftarrow head(C_l)$;

if $out_\alpha \notin Neurons(\mathcal{N})$ **then**

$AddOutputNeuron(\mathcal{N}, out_\alpha)$;

$AddLink(\mathcal{N}, h_l, out_\alpha, W)$;

$Threshold(h_l) \leftarrow \frac{(1+A_{min})(k_l-1)}{2} W$;

$Threshold(out_\alpha) \leftarrow \frac{(1+A_{min})(1-\mu_l)}{2} W$;

$ActivationFunction(h_l) \leftarrow h(x)$;

$ActivationFunction(out_\alpha) \leftarrow h(x)$;

end

foreach $\alpha \in atoms(\mathcal{P})$ **do**

if $(in_\alpha \in neurons(\mathcal{N})) \wedge (out_\alpha \in neurons(\mathcal{N}))$ **then**

$AddLink(\mathcal{N}, out_\alpha, in_\alpha, 1)$

end

foreach $in_\alpha \in neurons(\mathcal{N})$ **do**

if $(\alpha = \bullet^n \beta)$ **then**

if $\exists i < n(out_{\bullet^i \beta} \in neurons(\mathcal{N}))$ **then**

$j \leftarrow maximum(i)$;

$AddDelayedLink(\mathcal{N}, n-j, out_{\bullet^i \beta}, in_\alpha)$;

else

$AddInputDelay(\mathcal{N}, n, in_\alpha)$

end

return $\mathcal{N}$;

end

Connectionist Modal and Temporal Logics

Neural network ensembles correspond to possible worlds/states; modularity for learning; accessibility relations, disjunctive information.

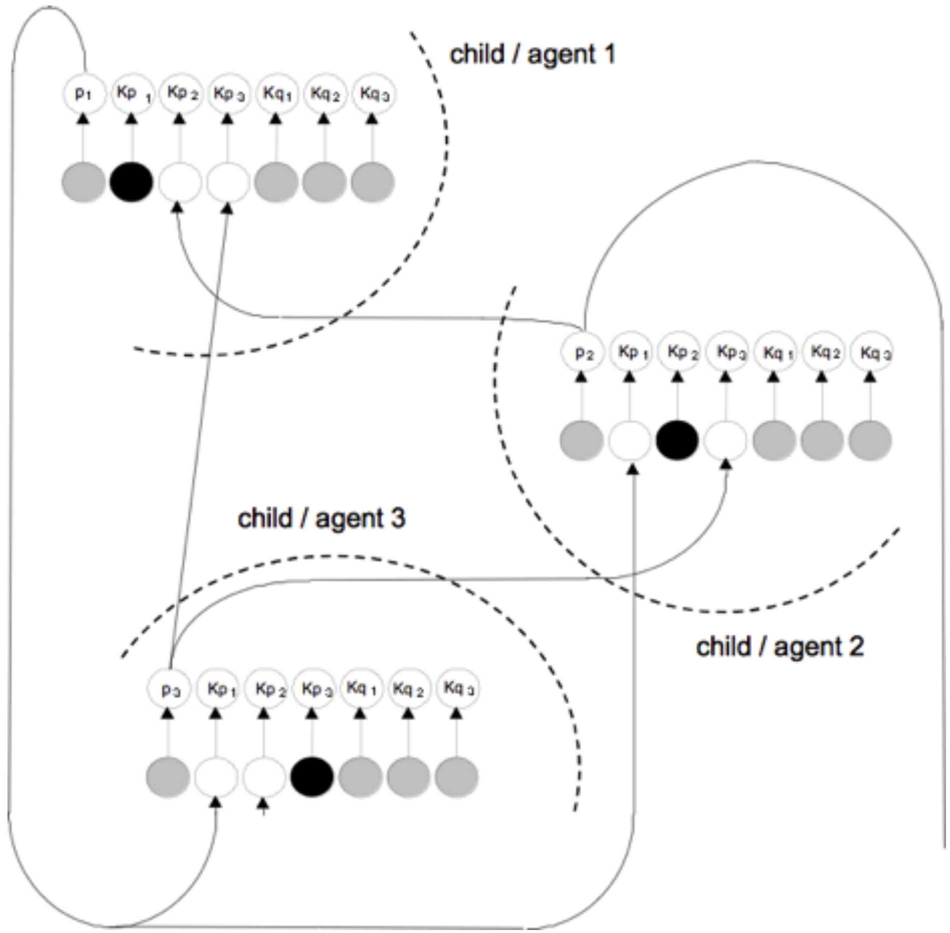

THEOREM 2: For any modal/temporal logic program *P* there exists an ensemble of neural networks *N* such that *N* computes *P*. *Garcez, Lamb, 2006.*

Connectionist Modal/Temporal Logics

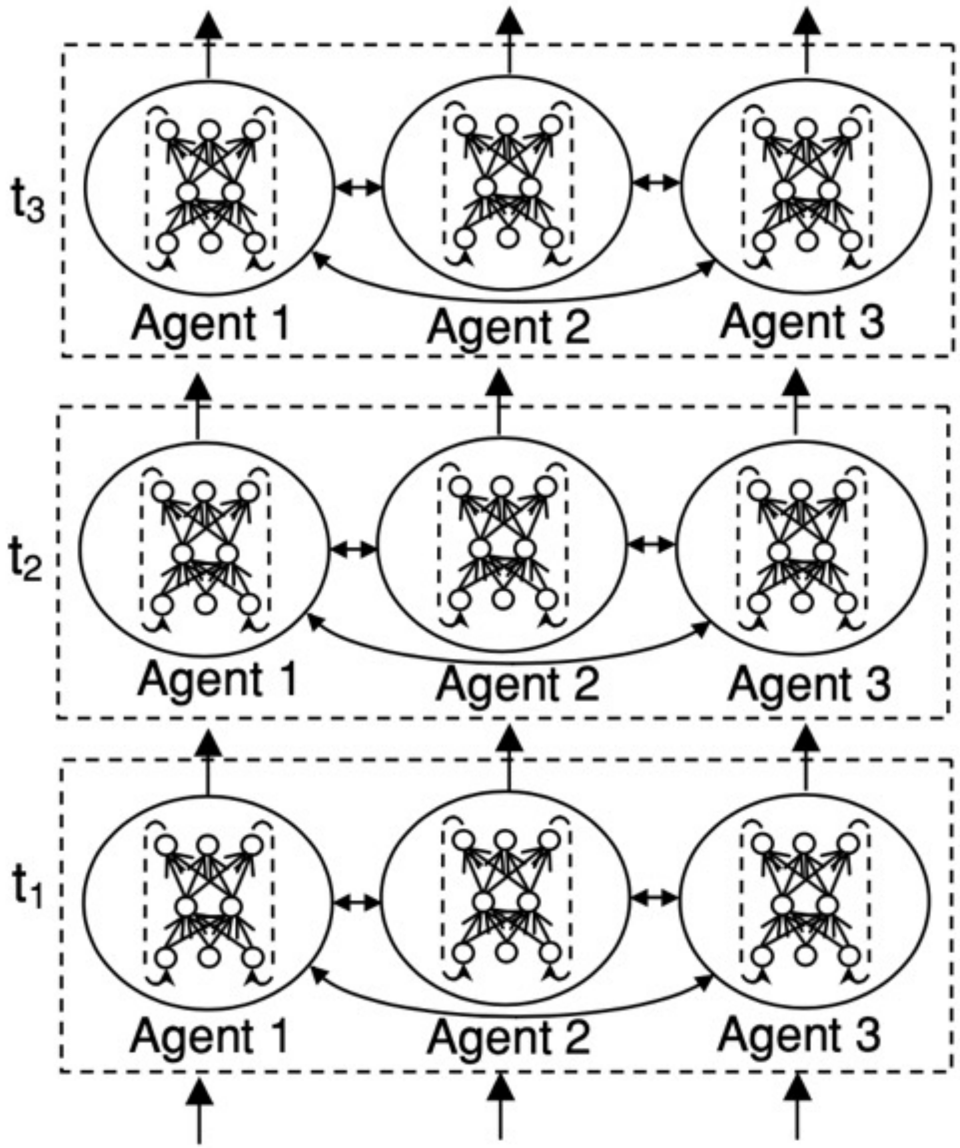

Figure 4: Evolving knowledge through time.

Knowledge evolves and is learned through time.
See NIPS 2003, AAAI 2007, NeCo 2006, IEEE TNN 2011

Learning to Adapt Requirements Specifications of Evolving Systems

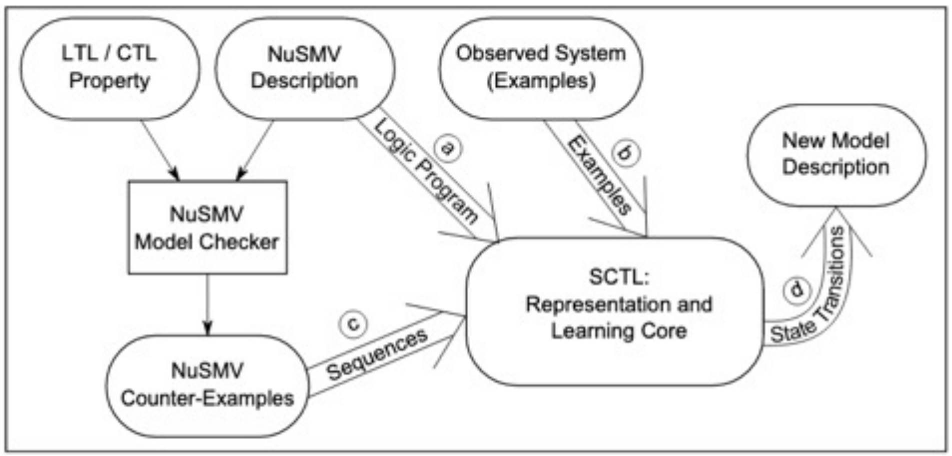

Figure 1: General diagram of the proposed framework

We consider the NuSMV model checker, and a neural network-based system (SCTL) to perform adaptation.

An initial description of a model can be expressed in NuSMV or as a temporal logic program. Also, it can be generalized from examples of the observed behaviour of an existing system

Borges, Garcez, Lamb,
IEEE Trans. NN 2011
ICSE2011, ASE2011 and more

Another relation to Vardi

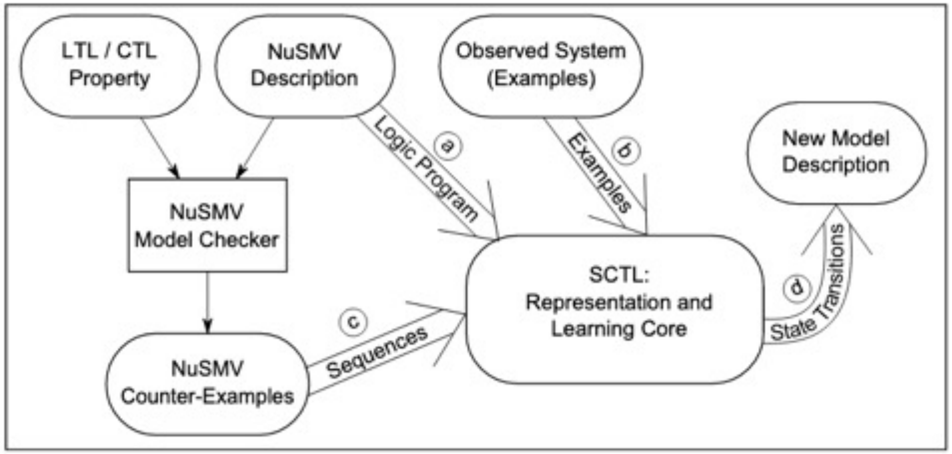

Figure 1: General diagram of the proposed framework

This model can then be subject to verification by the NuSMV model checker. If the model does not satisfy the given properties, a set of counter-examples is returned.

These counter-examples can be used as input of the adaptation engine in order to obtain a new, improved model. This new model can be subject to the same process until the properties are satisfied.

See also:
Labor Division with Movable Walls: Composing Executable Specifications with Machine Learning and Search.
David Harel, Assaf Marron, Ariel Rosenfeld, Moshe Y. Vardi, Gera Weiss: **AAAI 2019**

Sepp Hochreiter - Defined LSTMs with Schmidhuber

S. Hochreiter: Towards a Broad AI, CACM, April 2022.
"A broad AI [...] performs any cognitive task by virtue of its sensory perception, previous experience, and learned skills."

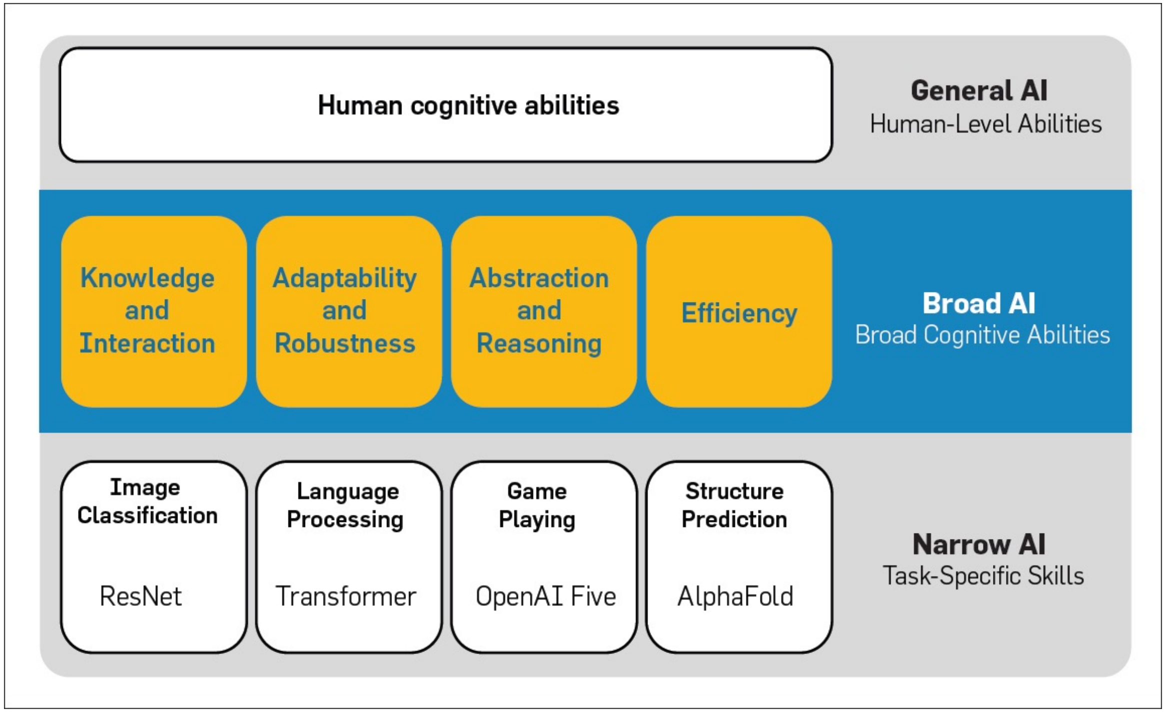

"The most promising approach to a broad AI is a neuro-symbolic AI, that is, a *bilateral AI* that combines methods from symbolic and sub-symbolic AI.

"GNNs are the predominant models of neural-symbolic computing.[6]"

[6] Lamb, L.C., Garcez, A., Gori, M., Prates, M., Avelar, P. and **Vardi, M**. ***Graph Neural Networks Meet Neural-Symbolic Computing: A Survey and Perspective.*** IJCAI (2020)

Neurosymbolic AI Open Challenges

(1) First-order logic and higher-order knowledge extraction from very large networks that is provably sound and efficient, explains the entire model and local network interactions and accounts for different levels of abstraction.

(2) Goal-directed commonsense and efficient combinatorial reasoning about what has been learned by a complex deep network trained on large amounts of multimodal data.

- E.g. Learning to Solve NP-Complete Problems: A Graph Neural Network for Decision TSP. M. Prates, P. Avelar, H. Lemos, L. Lamb, Moshe Vardi. AAAI-2019.

(3) Human-network communication: think of a multiagent system that promotes communication/argumentation protocols between the user and an agent that can ask questions and check her understanding.

Constructing an AI system that truly understands what it does is a recurring theme in the current debate.

The combination of learning and reasoning should offer an important alternative to the problem of combinatorial reasoning by learning **to reduce the number of effective combinations,** thus producing simpler symbolic descriptions as part of the neurosymbolic cycle.

Evolution of Neurosymbolic AI

2003: A. D’Avila Garcez & Luís C. Lamb: **Reasoning about Time and Knowledge in Neural Symbolic Learning Systems.** NIPS 2003: 921-928 - First Neurosymbolic System involving temporal logics of knowledge

2005: S. Bader and P. Hitzler, Dimensions of Neural-symbolic Integration - A Structured Survey, in: We Will Show Them! Essays in Honour of Dov Gabbay, Vol. 1., S.N. Artëmov, H. Barringer, A.S. d’Avila Garcez, L.C. Lamb, J. Woods, eds, College Pub., 2005.

2006: A d'Avila Garcez, LC Lamb, A Connectionist Computational Model for Epistemic and Temporal Reasoning. Neural Computation 18(7): 1711-1738

2007: A. d’Avila Garcez, L.C. Lamb, D.M. Gabbay, Connectionist Modal Logic: Theoretical Computer Science, 371: 34-53.

References

- **Learning to Solve NP-Complete Problems: A Graph Neural Network for Decision TSP. M. Prates, P. Avelar, H. Lemos, L. Lamb, M. Vardi. AAAI-2019.**
- L.C. Lamb, R. Borges, A. Garcez. A Connectionist Cognitive Model for **Temporal Synchonisation and Learning.** Proc. **AAAI-2007**.
- L. de Penning; A.S. d'Avila Garcez, Luis C. Lamb and J.J. Ch. Meyer. *A Neural-Symbolic Cognitive Agent for Online Learning and Reasoning*. Proc. **IJCAI-11**.
- R. V. Borges, Artur d'Avila Garcez, Luis C. Lamb. *Learning and Representing Temporal Knowledge in Recurrent Networks*. **IEEE T. Neural Networks**, Dec. 2011.
- R. V. Borges, Artur S. d'Avila Garcez, Luis C. Lamb and Bashar Nuseibeh. *Learning to Adapt Requirements Specifications of Evolving Systems*. In **ICSE 2011**.
- A.S. d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. ***Neural-Symbolic Cognitive Reasoning***, Springer 2009, 198pp.
- A.S. d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. Connectionist Modal Logic. ***Theoretical Computer Science***, 2007.

Evolution of Neurosymbolic AI

2019: A.D. Garcez, M. Gori, L. C. Lamb, L. Serafini, M. Spranger, S.Tran: **Neural-symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning.** FLAP 6(4): 611-632 (2019)

2020: Luís C. Lamb, Artur S. d'Avila Garcez, Marco Gori, Marcelo O. R. Prates, Pedro H. C. Avelar, **Moshe Y. Vardi**: **Graph Neural Networks** Meet **Neural-Symbolic Computing: A Survey and Perspective. IJCAI 2020**

Artur S. d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. Neural-Symbolic Cognitive Reasoning. Cognitive Technologies. Springer, 2009.

Thanks to Gary Marcus: "The Seminal Book on Neurosymbolic AI"

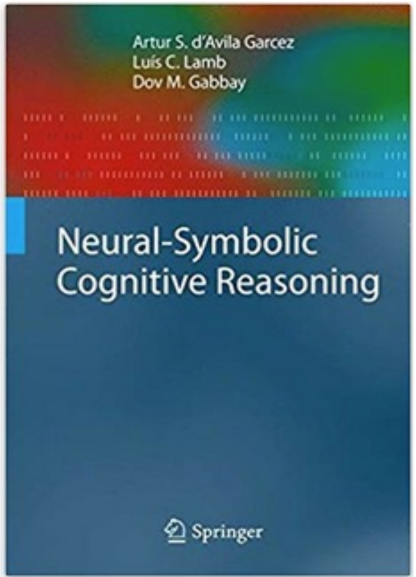

Evolution of Neurosymbolic AI- Recent Pointers

Artur Garcez, Luís C. Lamb: **Neurosymbolic AI: The 3rd Wave**. CoRR abs/2012.05876 (2020)

A. R. Tavares, P. Avelar, J. Flach, M. Nicolau, Luís C. Lamb, Moshe Y. Vardi: **Understanding Boolean Function Learnability on Deep Neural Networks**. CoRR abs/2009.05908 (2020)

Pascal Hitzler, Md. Kamruzzaman Sarker (Editors). **Neuro-Symbolic Artificial Intelligence.** IOS, 2021.

R.Riegel, A. Gray, F. Luus, N.Khan, N. Makondo, I. Akhalwaya, H.Qian, R.Fagin, F. Barahona, U.Sharma, S.Ikbal, H.Karanam, S.Neelam, A.Likhyani, S.Srivastava: **Logical Neural Networks**. CoRR 2006.13155 (2020)

Gabriele Ciravegna, Pietro Barbiero, Francesco Giannini, Marco Gori, Pietro Lió, Marco Maggini, Stefano Melacci: **Logic Explained Networks**. CoRR abs/2108.05149 (2021)

S. Badreddine, A. Garcez, L, Serafini, M. Spranger: **Logic Tensor Networks**. CoRR abs/2012.13635 (2020)

V. Belle: **Symbolic Logic meets Machine Learning:** A Brief Survey in Infinite Domains. CoRR abs/2006.08480 (2020)

Md. Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, Pascal Hitzler:
Neuro-Symbolic Artificial Intelligence: Current Trends. CoRR abs/2105.05330 (2021)

Afterword

*"Discoveries typically aren't made by people trying to solve a problem or invent something. Major discoveries are **not made in the lab**. They are **made in the minds of scientists**. Scientific research is what you do when you don't know what you are doing."*

Daniel Zajfman - Former President, Weizmann Institute of Science.

I wonder what's going on in Moshe's mind at the moment...

I still remember year 2000 (there is a story here)

Thank you, Moshe and Pam