Neurosymbolic AI: A bird's eye view



Luis Lamb, 2021

INSTITUTO
DE INFORMÁTICA
UFRGS

Summary

A bit of history: symbolism and connectionism

Cognition, AI and Computer Science

Neural-symbolic computing: origins and perspective

Afterword

"The tension between reasoning and learning has a long history, reaching back at least as far as Aristotle".

Leslie Gabriel **Valiant**, Probably Approximately Correct: Basic Books, NY, 2013.

Acknowledgments

Thanks to

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

https://mitibmwatsonailab.mit.edu/category/neuro-symbolic-ai/

Most Visited Getting Started l01c01_introduction...

Neuro-Symbolic AI

AI Today

MIT-IBM Watson AI Lab

As far back as the 1980s, researchers anticipated the role that deep neural networks could one day play in automatic image recognition and natural language processing. It took decades to amass the data and processing power required to catch up to that vision – but we're finally here. Similarly, scientists have long anticipated the potential for symbolic AI systems to achieve human-style comprehension. And we're just hitting the point where our neural networks are powerful enough to make it happen. We're working on new AI methods that combine neural networks, which extract statistical structures from raw data files – context about image and sound files, for example – with symbolic representations of problems and logic. By fusing these two approaches, we're building a new class of AI that will be far more powerful than the sum of its parts. These neuro-symbolic hybrid systems require less training data and track the steps required to make inferences and draw conclusions. They also have an easier time transferring knowledge across domains. We believe these systems will usher in a new era of AI where machines can learn more like the way humans do, by connecting words with images and mastering abstract concepts.

Categories

Deep Learning

Microsoft Research Blog

Next-generation architectures bridge gap between neural and symbolic representations with neural symbols

Published December 12, 2019

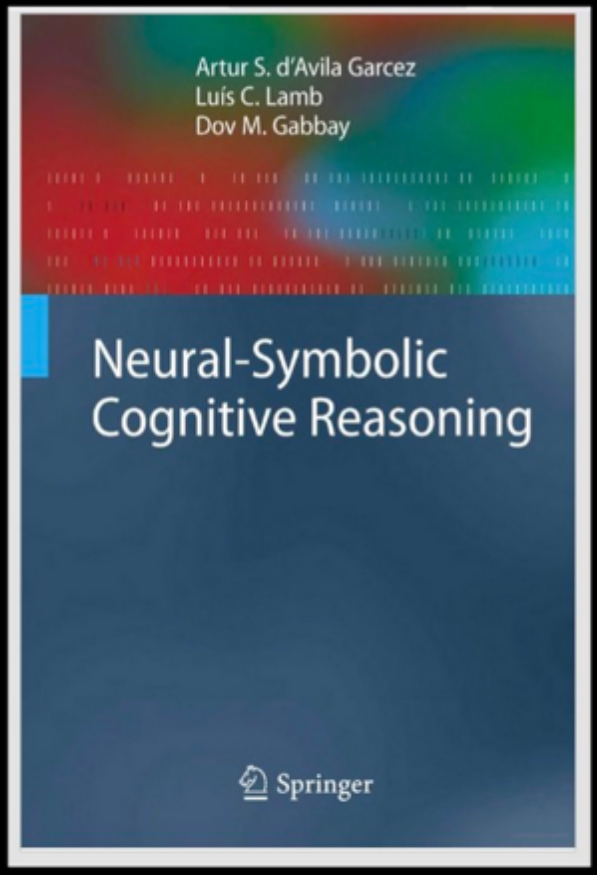
By Paul Smolensky, Partner Researcher

Research Area

Algorithms

Artificial intelligence

- Until roughly 2018, mainstream ML largely ignored *The Algebraic Mind*
- But *The Algebraic Mind* inspired the seminal book on neurosymbolic approaches
- And, as we will see, *Algebraic* also anticipated much of Yoshua's current argument



2008

Neural-Symbolic Cognitive Reasoning

Gary Marcus in the Great AI Debate, Montreal 23 december 2019.

Neurosymbolic AI: The 3rd 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

A key question: identifying the necessary and sufficient building blocks of AI, and how systems that evolve automatically based on machine learning can be developed in effective ways that make AI trustworthy.

What is the best representation for Neurosymbolic AI?

Challenges for principled integration of Learning and Reasoning.

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

Two high impact technologies in a generation







Hypes & Hopes: AI + Machine Learning

Machine learning: the power and promise of computers that learn by example
Issued: April 2017 DES4702
ISBN: 978-1-78252-259-1

Source: **ROYAL SOCIETY report: April 2017**
"Machine learning: the power and promise of computers that learn by example"

These can be related to Neurosymbolic AI

A new wave of research

Machine learning is a vibrant field of research, with a range of exciting areas for further development across different methods and applications. There is a collection of specific research questions where progress would directly address potential public concerns around machine learning, or constraints on its wider use:

- **Interpretability:** Can we create powerful machine learning systems where the reasons for particular decisions or recommendations can be understood or interrogated?
- **Verification and validation:** Can we create more advanced, and accurate, methods of verifying machine learning systems?
- **Privacy:** What are the technical solutions that can maintain the privacy of datasets, while allowing them to be used in new ways?
- **Fairness and dealing with real-world data:** How can real-world data be curated into usable forms, addressing 'real-world' messiness, and systemic – or social – bias?
- **Causality:** How can machine learning methods discover cause-effect relationships?
- **Human-machine interaction:** How do we design machine learning systems so they can work with humans safely and effectively?
- **Security:** How do we ensure machine learning systems are not vulnerable to cyber-attack?

Support for research in these areas is needed to help ensure continued public confidence in the deployment of machine learning systems.

Hypes & Hopes: AI + Machine Learning

BOX 5

Why are some machine learning systems 'black boxes'?

A neural network is an approach to machine learning in which small computational units are connected in a way that is inspired by connections in the brain. These systems may consist of many layers of neurons: the base layer receives an input from an external source, then each layer beyond it detects patterns in activity from the neurons in the layer beneath, integrates these inputs, and then passes a signal to the next layer[181]. In this way, signals can be passed through many layers, before reaching a top layer where a decision about the input is made. So, if the initial input is an image, the initial signals might come from the pixels of the image, and the top-level decision might be what object is in the image. As such systems learn from data, they strengthen or weaken synaptic connections to make their outputs more accurate.

This approach to processing data means that information across the neural network is highly dispersed, with complicated patterns of connection strength between units, and with potentially many thousands of layers of units. The result is the so-called 'black box' issue: these systems can create highly accurate results, but it is difficult to explain why a result has been obtained.

However, not all machine learning methods use this approach, and alternative approaches can be more readily interpreted.

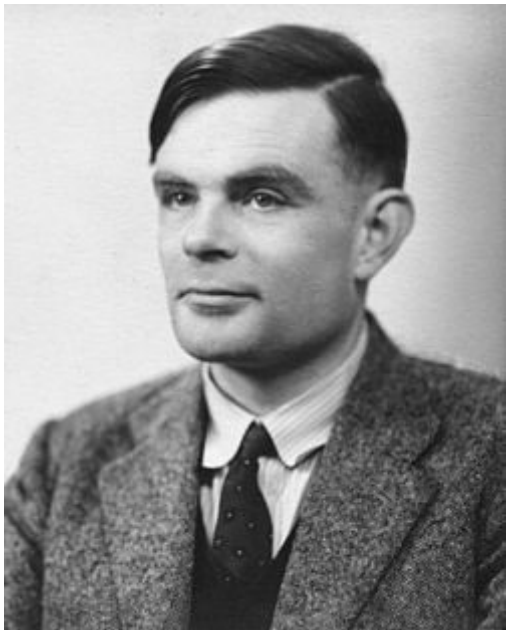
This is directly related to this topic

Source: **ROYAL SOCIETY report: April 2017**
"Machine learning: the power and promise of computers that learn by example".
ISBN: 978-1-78252-259-1

Enters Turing

The whole thinking process is still rather mysterious to us, but I believe that the attempt to make a thinking machine will help us greatly in finding out how we think ourselves.

Alan Turing, 15 May 1951, “Can Digital Machines Think” BBC.



.inf UFRGS

Can machines think?

How much has been done?

What happened from 2006?

*At **AAAI-07**, arguably the flagship, top AI Conference **there were virtually no papers which made any kind of use of Artificial Neural Networks***
(by the way: there was a paper on Neural-Symbolic Learning: Lamb, Borges & d'Avila Garcez)

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

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

"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 – President, Weizmann Institute of Science

What happened from 2006-2020?

We stress that it's hard to present a fair review of the literature, at this time.

Back to the 1990's:

Long short-term memory (LSTM): Sepp Hochreiter & Jürgen Schmidhuber: *Long Short-Term Memory*, Neural Computation, 9(8):1735–1780, 1997.
and many others.

Groundbreaking 2000's papers:

Imagenet classification with deep convolutional neural networks. A. Krizhevsky, I. Sutskever, GE Hinton, **NIPS 2012**.
A fast learning algorithm for deep belief nets. GE Hinton, S Osindero, YW Teh
Neural computation 18 (7), 1527-1554, 2006.

Top surveys:

*Bengio, Yoshua; LeCun, Yann; Hinton, Geoffrey (2015). "Deep Learning". **Nature**.* 521: *436–444.*

*Bengio, Y.; Courville, A.; Vincent, P. (2013). "Representation Learning: A Review and New Perspectives". **IEEE Transactions on Pattern Analysis and Machine Intelligence**. **35** (8): 1798–1828*

*Schmidhuber, J. (2015). "Deep Learning in Neural Networks: An Overview". Neural Networks. **61**: 85–117*

Alan Newell and Herbert Simon: Symbolists?

- Numerous contributions to AI, problem solving, psychology of cognition; decision making, list processing before LISP, since the 1950s
- **Simon is the only person to have won** both **the Turing Award and the Nobel Prize in Economic Sciences**.
- AI programs (Newell, Simon, Shaw): the Logic Theorist (1956: proved 38 theorems of Whitehead & Russell's *Principia Mathematica*).
- General Problem Solver (1957): theorems, geometric problems and chess playing. Separated knowledge from strategy.

http://diva.library.cmu.edu/Newell/biography.html



“Recent” AI Turing Award Winners: “Connectionists”

“For conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing.” 2018





Y. Bengio, G. Hinton, Y. LeCun

Photos from the Association for Computing Machinery
http://amturing.acm.org/

Recent AI Turing Award Winners

2010: LESLIE **VALIANT** (HARVARD): PAC ***LEARNING***; COMPUTATIONAL COMPLEXITY; PARALLEL/DISTRIBUTED COMPUTING;

"A fundamental question for artificial intelligence is to ***characterize the computational building blocks that are necessary for cognition.*** *A specific challenge is to build on the success of machine learning so as to cover broader issues in intelligence. This requires, in particular a* ***reconciliation between two contradictory characteristics – the apparent logical nature of reasoning and the statistical nature of learning****."*



Photos from the Association for Computing Machinery
http://amturing.acm.org/

A Small Step: Neural-symbolic Computation

Combines logical reasoning and neural learning:

Computer Science Logic + Neural Computation

Neural-symbolic computation: Learning from experience and reasoning about what has been learned from an uncertain environment in a computationally efficient way.



Neural-symbolic computation

Motivation: there is a need for systems that:

-Learn from changes in the environment.

-Reason about commonsense knowledge.

-Integrate reasoning (computation) and learning: we combine the logical nature of reasoning and the statistical nature of learning, see Valiant.

-Reasoning is hard:

Gertrude Stein, "Autobiography of Alice B. Toklas" ... that only three times in my life have I met a genius ... the three geniuses of whom I wish to speak are Gertrude Stein, Pablo Picasso and Alfred Whitehead.

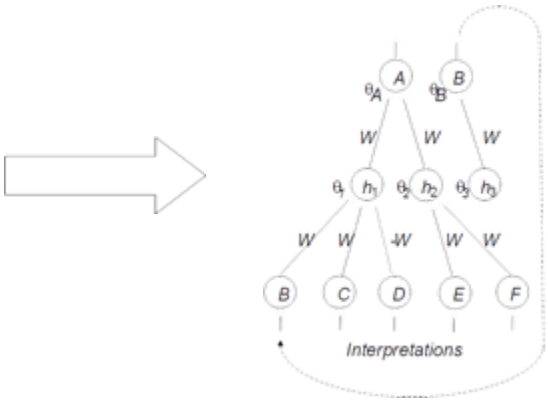
A History of Neural-Symbolic Computation

The CILP System: Artur Garcez & Gerson Zaverucha, 1999.

r_1: $A \leftarrow B, C, \sim D;$

r_2 : $A \leftarrow E, F;$

r_3 : $B \leftarrow$



KBANN: first system first to allow for learning with background knowledge in neural networks and knowledge extraction.

G. Towell and J. Shavlik. Knowledge-based artificial neural networks. Artif. Intel.,70:119–165, 1994.

S. Hoelldobler & Y. Kalinke: Toward a New Massively Parallel Computational Model for Logic Programming; *Workshop on Combining Symbolic and Connectionist Processing*, ECAI-1994.

CILP: allowed for the integration of learning, reasoning and knowledge extraction in recurrent networks.

CILP: neural networks endowed with semi-linear neurons approximate the fixed-point operator of propositional logic programs with negation.

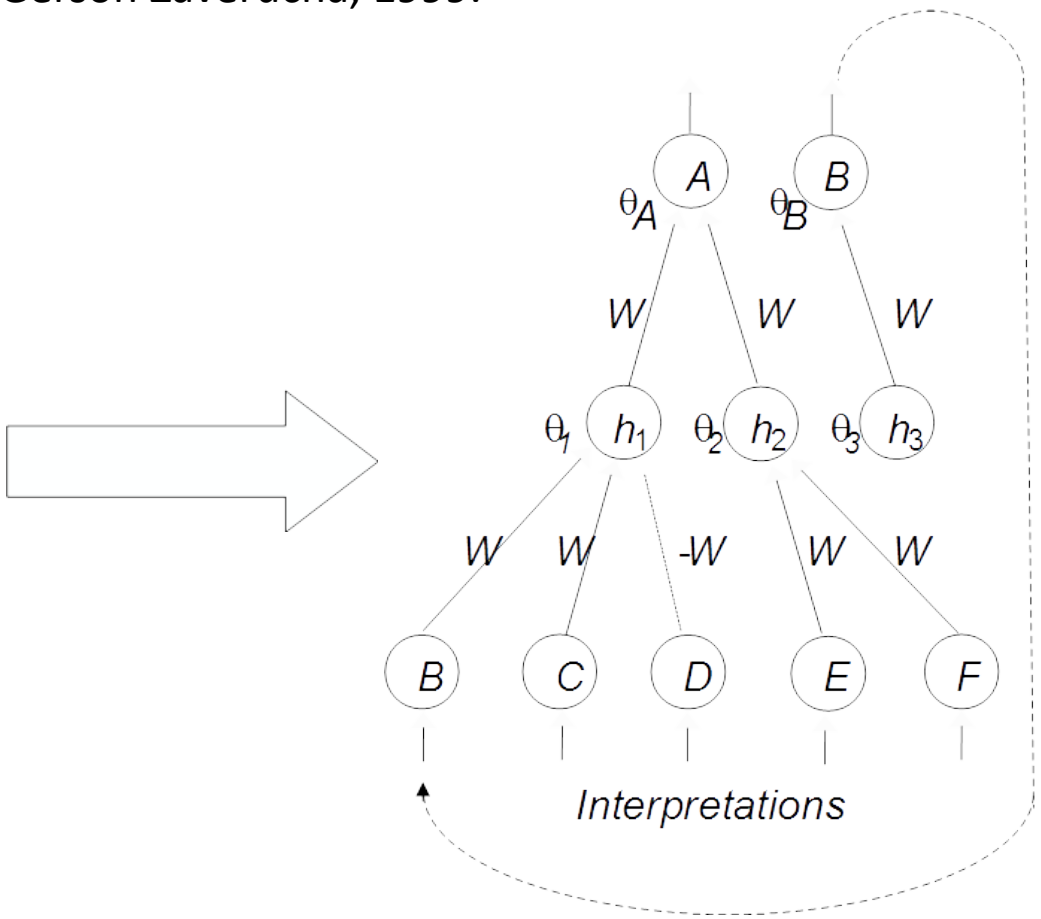
A History of Neural-Symbolic Computation

The CILP System: Artur Garcez & Gerson Zaverucha, 1999.

r_1: $A \leftarrow B,C,\sim D;$

r_2 : $A \leftarrow E,F;$

r_3 : $B \leftarrow$



THEOREM 1: For any logic program P there exists a neural network N such that N computes P. *Garcez, Zaverucha, Appl. Intell J, 1999.*

Conectionist Modal Logics

Modal logic goes beyond propositional reasoning: Moshe Vardi, 1997.

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 if it is true in at least one world which is possible in relation to that same world.

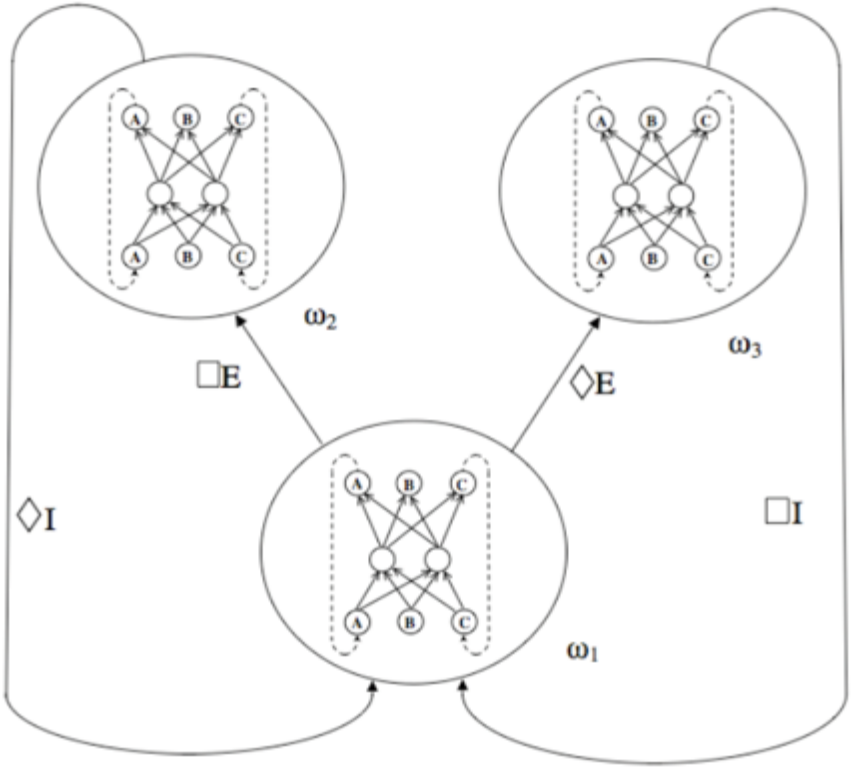
Modalities also used for reasoning about uncertainty (following J. Halpern).

Relational learning/reasoning is notoriously hard:

NeSy workshop series; StarAI workshops, Dagstuhl Seminars, AAAI Spring symposium KRR+ML.

Learning to reason in connectionist models

- Insight: *assume that neurons are possible worlds.*
- Propositional Modal Logic = decidable fragment of FOL with two variables.
- Full solution of Muddy Children puzzle and other testbeds.
- Ensembles of NNs that are seen as possible world.



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

Table 1
Rules for modality operators

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

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

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

Connectionist Modal LogicsRules and Reasoning

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

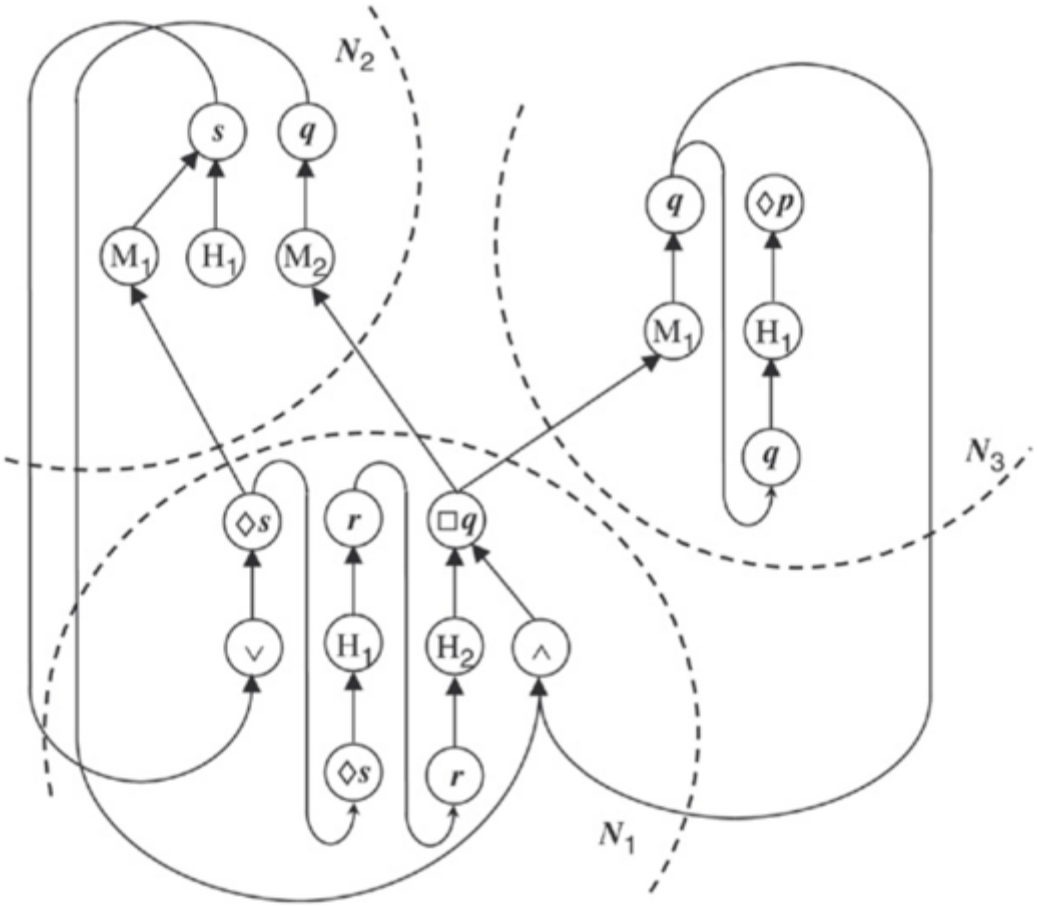


Fig. 5. The ensemble of networks $\{\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3\}$ that represents $\mathcal{P}$.

Connectionist Modal Logics – translation or embedding

Algorithm 2: Translation of ●-based programs

●-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**
    - $j \leftarrow maximum(i)$;
    - $AddDelayedLink(\mathcal{N}, n-j, out_{\bullet^j \beta}, in_\alpha)$;
  - **else**
    - $AddInputDelay(\mathcal{N}, n, in_\alpha)$

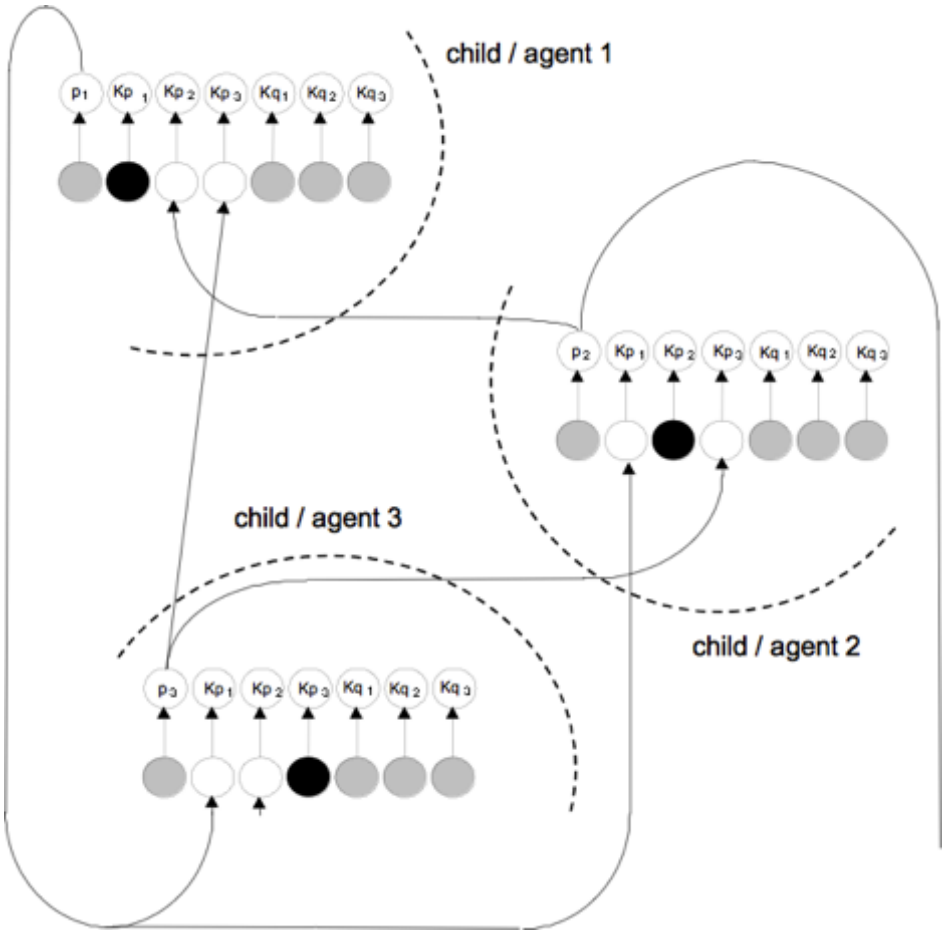
end

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 Logics

Definition 9 (*temporal model*). *A model M is a tuple $M = (T, R_1, \ldots, R_n, \pi)$, where (i) T is a (linear) time line; (ii) R_i ($i \in \mathcal{A}$) is an agent accessibility relation over T; and (iii) $\pi : T \mapsto \varphi$ is a map associating with each propositional variable of CTLK a set $\pi(p)$ of time points in T.*

Definition 10 (*satisfiability of temporal formulas*). *Let $M = \langle T, R_i, \pi \rangle$ be a temporal model for CTLK. The satisfiability relation $\models$ is uniquely defined as follows:*

i. $(M, t) \models p$ *iff* $t \in \pi(p)$

ii. $(M, t) \models \neg\alpha$ *iff* $(M, t) \not\models \alpha$

iii. $(M, t) \models \alpha \wedge \beta$ *iff* $(M, t) \models \alpha$ *and* $(M, t) \models \beta$

iv. $(M, t) \models \alpha \vee \beta$ *iff* $(M, t) \models \alpha$ *or* $(M, t) \models \beta$

v. $(M, t) \models \alpha \rightarrow \beta$ *iff* $(M, t) \not\models \alpha$ *or* $(M, t) \models \beta$

vi. $(M, t) \models \Box\alpha$ *iff for all* $u \in T$, *if* $R(t, u)$ *then* $(M, u) \models \alpha$

vii. $(M, t) \models \Diamond\alpha$ *iff there exists a* u *such that* $R(t, u)$ *and* $(M, u) \models \alpha$

viii. $(M, t) \models \bigcirc\alpha$ *iff* $(M, t + 1) \models \alpha$

ix. $(M, t) \models K_i\alpha$ *iff for all* $u \in T$, *if* $R_i(t, u)$ *then* $(M, u) \models \alpha$

Connectionist Modal/Temporal Logics

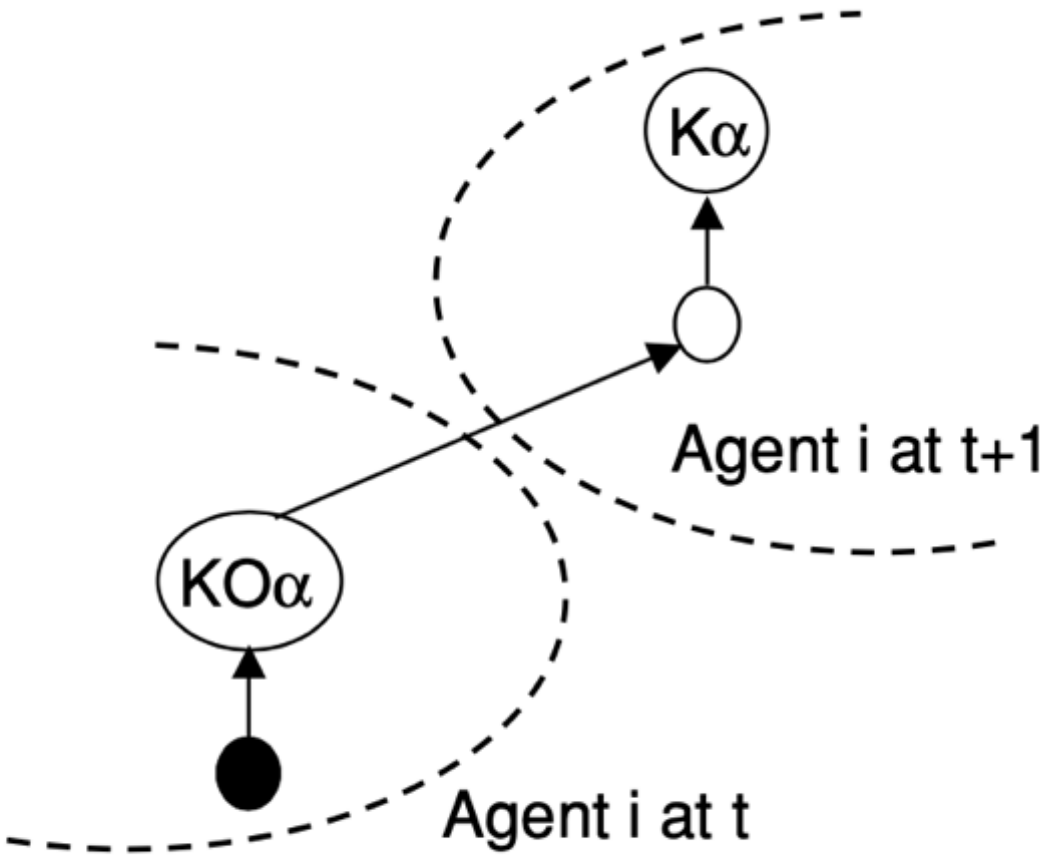


Figure 3: Simple example of connectionist temporal reasoning.

Connectionist Modal Logics

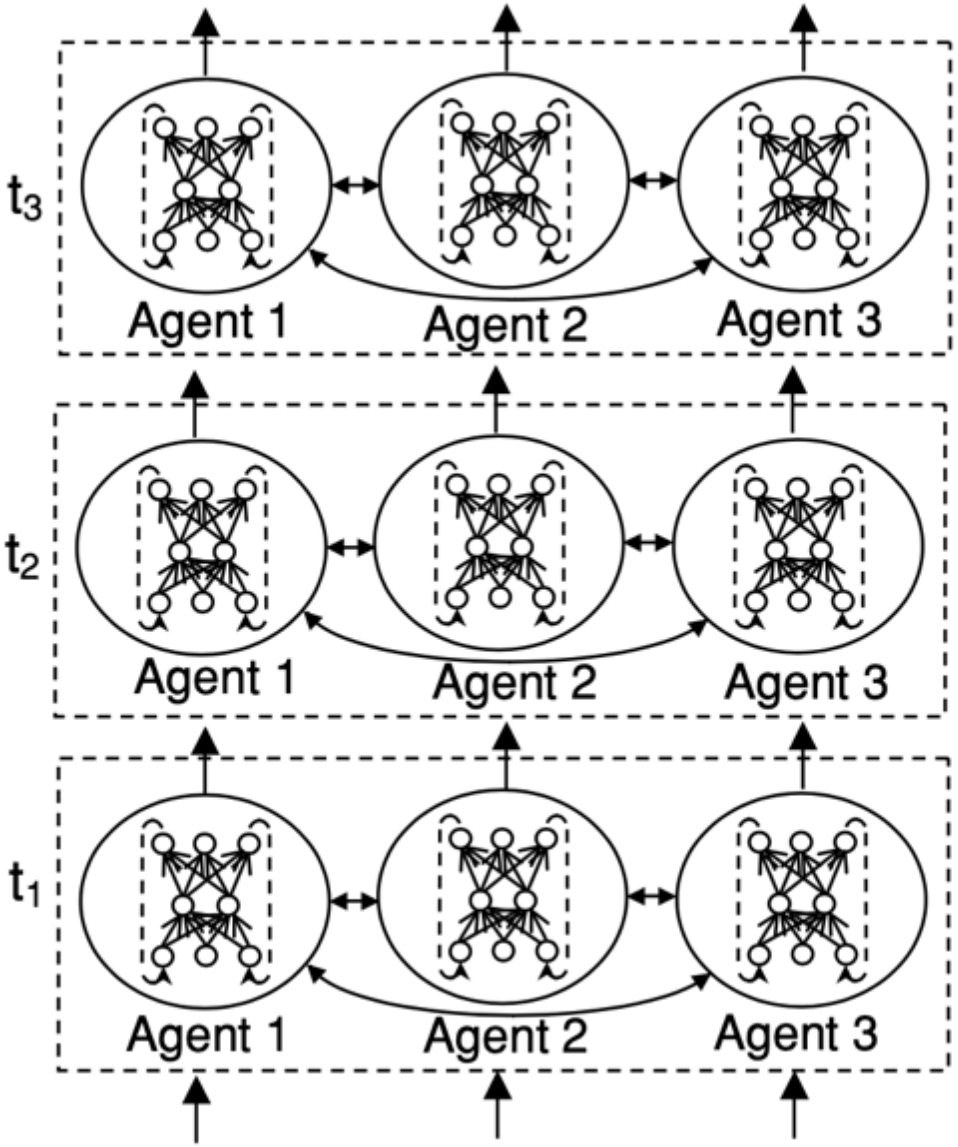


Figure 4: Evolving knowledge through time.

Connectionist Modal Logics

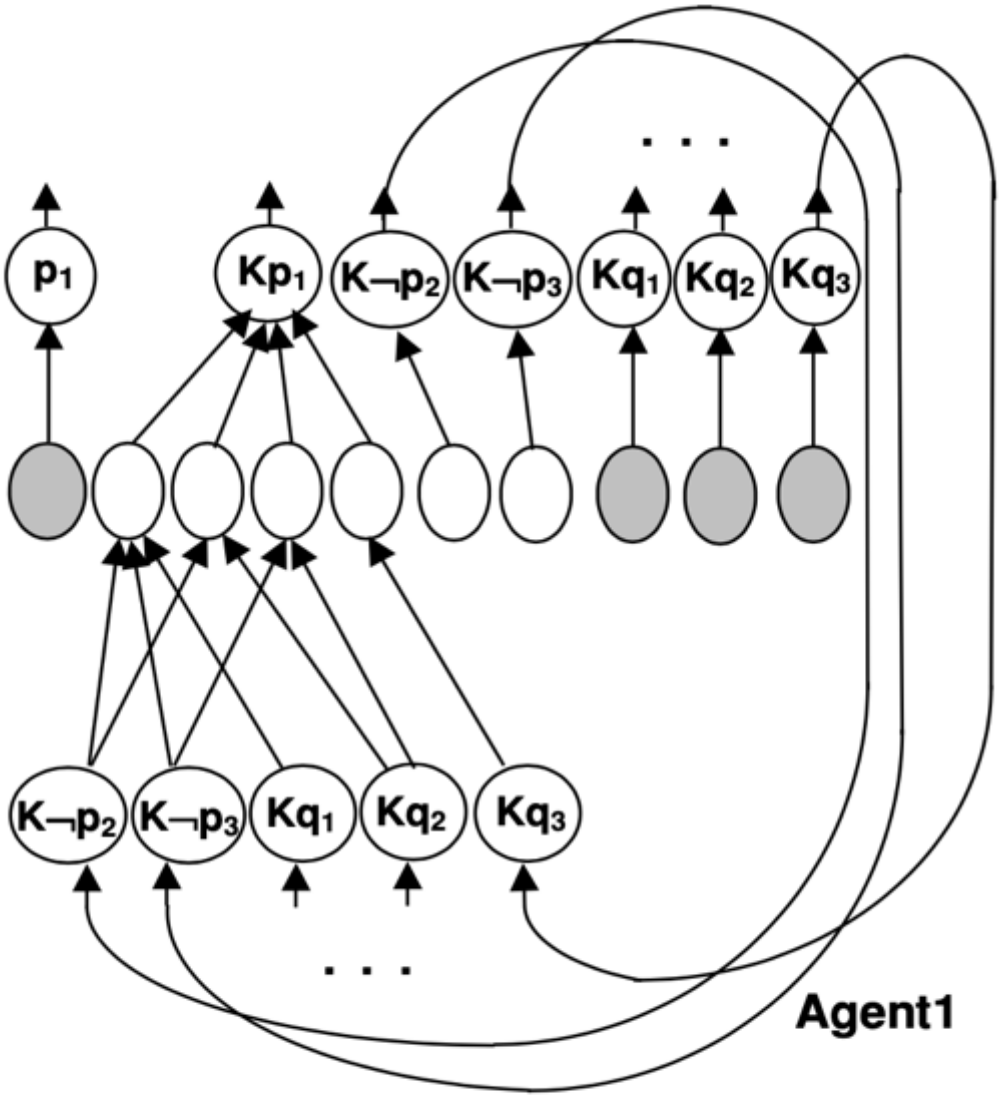


Figure 5: The implementation of rules $\{r_1^1, \ldots, r_4^1\}$.

Snapshot Rules for Agent (Child) 1

r_1^1: $\mathbf{K}_1 q_1 \wedge \mathbf{K}_1 \neg p_2 \wedge \mathbf{K}_1 \neg p_3 \rightarrow \mathbf{K}_1 p_1$

r_2^1: $\mathbf{K}_1 q_2 \wedge \mathbf{K}_1 \neg p_2 \rightarrow \mathbf{K}_1 p_1$

r_3^1: $\mathbf{K}_1 q_2 \wedge \mathbf{K}_1 \neg p_3 \rightarrow \mathbf{K}_1 p_1$

r_4^1: $\mathbf{K}_1 q_3 \rightarrow \mathbf{K}_1 p_1$

Connectionist Modal Logics

Training and Assessment in Simulators: Learning new information from observation of experts. Leo de Penning, TNO

Trainees at task execution and reasoning about this information online to provide feedback to the user



Neurosymbolic AI: Evolution and Events

- We are constructing a principled approach to combine learning and reasoning.
- Reasoning is hard... as one can see around the world.
- Relational, full predicate reasoning is undecidable: modal logics offer a decidable, powerful alternative.
- Scientists still don't know how to reason and learn relations, see **deep learning + knowledge representation**
- NeSy workshop series since 2005
- KR2ML (Knowledge Representation & Reasoning Meets Machine Learning at NeuRIPS 2019, 2020
- AAAI Spring Symposium, Stanford, 2015.

Neural-symbolic computation

“We need some language for describing the alternative algorithms that a network of neurons may be implementing”. (Valiant, CACM 2011)

Cognition is Power

- Turing, Simon, Newell: AI, cognition, ambition.

- **Watson/Deep Blue**: gigantic amounts of ~~computing~~ (reasoning on) data; 1997: Gary Kasparov walks away from Deep Blue

"I could feel – I could smell – a new kind of intelligence across the table... *Although I think I did see some signs of intelligence, it's a weird kind, an inneficient, inflexible kind that maked me feel I have a few years left."*
Gary Kasparov, 1997.



A History of Neural-Symbolic Computation

by A. d'Avila Garcez et al –AAAI Spring Symposium 2015

1988: P Smolensky, On the proper treatment of connectionism, BBS:11(1); J McCarthy (commentary), Epistemological challenges for connectionism.
1990: G Hinton, Preface to the special issue on connectionist symbol processing, Artificial Intelligence 46,1-4

1993: L Shastri and V Ajjanagadde: From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony, BBS:16 (SHRUTI)
1994: G Towell, J Shavlik: Knowledge-Based Artificial Neural Networks. Artif. Intell. 70(1-2): 119-165 (KBANN)
1994: S. Hoelldobler & Y. Kalinke: Toward a New Massively Parallel Computational Model for Logic Programming; Workshop on Combining Symbolic and Connectionist Processing, ECAI.
1997: M Craven, J Shavlik, Understanding Time-Series Networks: A Case Study in Rule Extraction. Int. J. Neural Syst. 8(4): 373-384

2001: A Browne, R Sun, Connectionist inference models. Neural Networks 14(10): 1331-1355
2002: A d'Avila Garcez, K. Broda and D Gabbay, Neural-Symbolic Learning Systems: Foundations and Applications, Springer. (CILP)

A history of Neural-Symbolic Computation

2006: A d'Avila Garcez, L 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.

2007: S Bader, P Hitzler, S Hölldobler, A. Witzel, A Fully Connectionist Model Generator for Covered First-Order Logic Programs. IJCAI 2007: 666-671

2007: Y Bengio, Y LeCun, Scaling learning algorithms towards AI. Large-scale kernel machines 34(5) (Representation Learning)

2009: A. d'Avila Garcez, L. Lamb, D. Gabbay, Neural-Symbolic Cognitive Reasoning. Cognitive Technologies, Springer

2010: D. Endres, P. Foldiak, U. Priss, An application of formal concept analysis to semantic neural decoding. Annals of Mathematics and Artificial Intelligence 57(3-4): 233-248

2014: M. Franca, G. Zaverucha, A d'Avila Garcez, Fast relational learning using bottom clause propositionalization with artificial neural networks. Machine Learning 94(1): 81-104 (CILP++)

A history of Neural-Symbolic Computation

Artur S. d'Avila Garcez, Marco Gori, Luís C. Lamb, Luciano Serafini, Michael Spranger, Son N. Tran: **Neural-symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning.** FLAP 6(4): 611-632 (2019)

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

Books

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

•Barbara Hammer, Pascal Hitzler (Eds), Perspectives of Neural-Symbolic Integration. Studies in Computational Intelligence, Vol. 77. Springer, 2007.

•Artur S. d'Avila Garcez, Krysia Broda and Dov M. Gabbay. Neural-Symbolic Learning Systems: Foundations and Applications, Springer-Verlag, 2002.

And much, much more.

VENUES

NeSy Workshops and Seminars:

NeSy'20, 15th International Workshop on Neural-Symbolic Learning and Reasoning at IJCLR-20
NeSy'19, 14th International Workshop on Neural-Symbolic Learning and Reasoning at IJCAI-19
NeSy'18, 13th International Workshop on Neural-Symbolic Learning and Reasoning at HLAI-18
[…]
NeSy'05, First International Workshop on Neural-Symbolic Learning and Reasoning at IJCAI-05

Dagstuhl Seminars:

Dagstuhl Seminar on Logic and Learning, September 2019.
Dagstuhl Seminar on Human-Like Neural-Symbolic Computing, May 2017.
Dagstuhl Seminar on Neural-Symbolic Learning and Reasoning, September 2014.
Dagstuhl Seminar on Perspectives of Recurrent Neural Networks - Models, Capacities, and Applications, January 2008.

Recent Developments, AAAI 2020, 2021

More papers on neural networks and symbolic AI,

Keynotes and debates.

(1) Fireside conversation between 2019 Turing award winners, Francesca Rossi, and Nobel Laureate Daniel Kahneman on thinking fast and slow and its relation to neural networks and symbol manipulation.

(2) Henry Kautz's The Third AI Summer - Robert S. Engelmore Memorial Award Lecture, which introduced a taxonomy for neurosymbolic computing

(3) David Cox IAAI2020 Invited talk - Director of MIT/IBM Watson AI laboratory On neurosymbolic AI and its implications on vision and language understanding, machine commonsense, question answering, argumentation and XAI.

Henry Kautz

Neurosymbolic Taxonomy - AAAI 2020

In "The Third AI Summer"

AAAI Robert S. Engelmore Memorial Lecture

Thirty-fourth AAAI Conference on Artificial Intelligence, New York, Feb. 10, 2020



Kautz's taxonomy answers a question:

How do neuro(al) and symbolic systems interact?

- Type 1.symbolic Neuro symbolic
- Type 2.Symbolic[Neuro]
- Type 3.Neuro;Symbolic
- Type 4.Neuro:Symbolic → Neuro
- Type 5.$\text{Neuro}_{\text{Symbolic}}$
- Type 6.Neuro[Symbolic] - Kautz's favourite category

1.symbolic Neuro symbolic

Standard deep learning; is a stretch to refer to as neural-symbolic, but is included to note that the input and output of a neural network can be made of symbols

e.g. language translation or question answering applications.

2.Symbolic[Neuro]

hybrid systems such as DeepMind's AlphaGo and other systems where the core neural network is coupled with a symbolic problem solver such as Monte Carlo tree search. E.g. AlphaGo (2016); AlphaZero (2017)

Type 3. Neuro;Symbolic

hybrid system where a neural network focus one task (e.g. object detection) & interacts via input/output with a symbolic system specialising in a complementary task (e.g. query answering).

E.g. neuro-symbolic concept learner.

Type 4. Neuro:Symbolic → Neuro

symbolic knowledge is compiled into the training set of a neural network.

- Rule A → B becomes an input-output training pair (A, B)
- Deep Learning for Symbolic Mathematics (Lample & Charton 2020 - **Deep Learning for Symbolic Mathematics)**

Type 5. Neuro$_{\text{Symbolic}}$

neural-symbolic systems where a symbolic logic rule is mapped onto a distributed representation (an embedding) and acts as a soft-constraint (a regularizer) on the network's loss function. Such systems are referred to as tensorization; e.g. Logic tensor networks.

Symbolic rules are used as *templates* for structures within the neural network

Logic Tensor Networks - Serafini & D'Avila Garcez, 2016.

Type 6. Neuro[Symbolic]

True symbolic reasoning inside a neural engine.
Fully-integrated system.

Early work in neural-symbolic computing has achieved this (CML, etc), but in a localist rather than a distributed architecture.

Kautz adds that such a system should be capable of combinatorial reasoning.

"Enable super-human and super-neuro COMBINATORIAL reasoning." Towards **THINKING, FAST AND SLOW?**

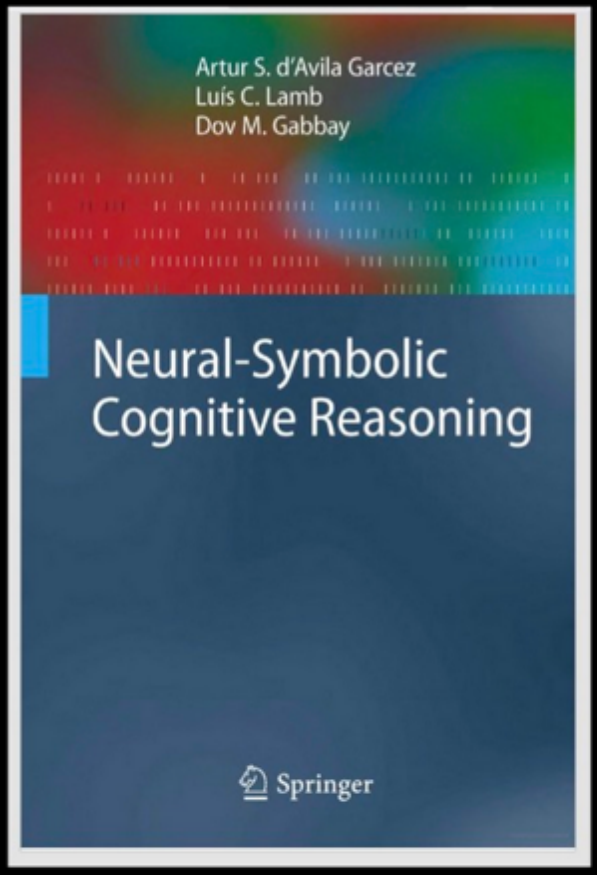
Henry Kautz's conclusions at AAAI2020:

- Each AI summer has led to enduring scientific insights.
- Today's third summer is different
  - It might not be followed by a winter
  - It enables powerful applications for good and bad
- **The next steps in AI are tighter symbolic-neuro integration**

Valiant, 1984, 2003, 2013

"...[the] remaining design choice that has to be made is that of **knowledge representation.** Since our declared **aim is to represent general knowledge**, it seems almost unavoidable that we use **some kind of logic** rather than, for example, formal grammars or geometrical constructs. [...] we shall **represent concepts as Boolean functions of a set of propositional variables**. The recognition algorithms that we attempt to deduce will be therefore Boolean circuits or expressions."

- Until roughly 2018, mainstream ML largely ignored *The Algebraic Mind*
- But *The Algebraic Mind* inspired the seminal book on neurosymbolic approaches
- And, as we will see, *Algebraic* also anticipated much of Yoshua's current argument



2008

Neural-Symbolic Cognitive Reasoning

Gary Marcus in the Great AI Debate, Montreal 23 december 2019.

Kahneman made the point that system 1 and system 2 have a long history in psychology research

He argued that system 1 - as the intuitive, fast parallel system - is capable of understanding language in contradiction with DL current account of system 1 and system 2.

System 2, is sequential, deliberative, reflexive system and is most probably performing symbol manipulation as argued by Gary Marcus.

System 1: that operates automatically and quickly with little or no effort and no sense of voluntary control;

System 2: that allocates attention to the effortful mental activities that demand it, including complex computations.

Kahneman "...so far as I'm concerned, System 1 certainly knows language... System 2... does involve certain manipulation of symbols."

AI and public perception

- Deep learning + AI have recently impacted **people's perception** of Computer Science. A few groundbreaking results contributed to this.
- There have been many accomplishments, but the follwing are clearly noteworthy:

(0) 1997: Deepblue beats chess world champion Kasparov.
(1) 2011: Watson wins Jeopardy!
(2) 2012: ImageNet Classification, by Hinton et al - groundbreaking result from deep learning in image recognition.
(3) 2016: Google DeepMind's AlphaGo beats Lee Sedol at the ancient Chinese game of Go.
(4) 2017: Poker playing Texas Hold'Em at human level ability: CMU's Libratus + U of Alberta's DeepStack

Explaining AI

(1) 2012: ImageNet Classification, by Hinton et al - groundbreaking result from deep learning in image recognition.

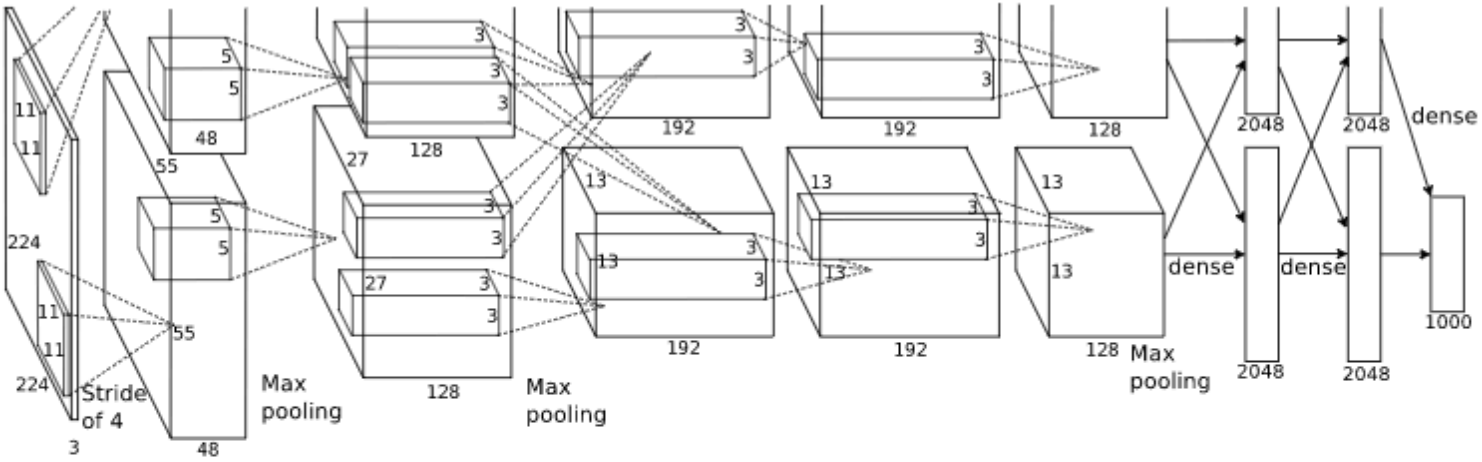


Figure 2: An illustration of the architecture of our CNN, explicitly showing the delineation of responsibilities between the two GPUs. One GPU runs the layer-parts at the top of the figure while the other runs the layer-parts at the bottom. The GPUs communicate only at certain layers. The network's input is 150,528-dimensional, and the number of neurons in the network's remaining layers is given by 253,440–186,624–64,896–64,896–43,264–4096–4096–1000.

Imagenet classification with deep convolutional neural networks. A. Krizhevsky, I. Sutskever, GE Hinton, **NIPS 2012**.

From Dagstuhl 2017:

In summary, the main research directions identified in the seminar are:

The recent success of deep learning in computer vision and language processing [10],associated with the growing complexity of big data applications demands the development of explainable AI models.

Neural-symbolic computation offers several methods andprinciples that contribute towards this aim.

For instance, rule extraction mechanisms and the integrated methodology for the development of neural-symbolic cognitive systems offers principled alternatives.

The concept of modularity in deep networks can be explained by using neural-symbolic methods. Modularity is suitable to knowledge extraction, which can result in the development of effective knowledge extraction methods towards explainable AI.

The development of Cognitive Architectures capable of simulating and explaining aspects of human cognition also remains an important research endeavour.

Some cognitive architectures typically consider symbolic representations, whereas others employ neural simulations. **The integration of these models remains a challenge and there are benefits on integrating the accomplishments of both paradigms.**

There remains a number of challenges in integrating, explaining and computing symbolic knowledge in deep networks. Both LTNs and Connectionist Modal and Temporal Logics offer effective alternatives towards these research challenges

Understanding AI

“If you want the computer to have general intelligence, the outer structure has to be commonsense knowledge and reasoning”, John McCarthy, in Shasha&Lazere, 1995**.**

“Our field is still in its embryonic stage. It's great that we haven't been around for 2000 years. We are at a stage where very, very important results occur in front of our eyes.”

M.O. Rabin, *in*
Denis Shasha and Cathy Lazere: ***Out of Their Minds: The Lives and Discoveries of 15 Great Computer Scientists****, 1995.*

References

- Artur S. d'Avila Garcez, Marco Gori, Pascal Hitzler, Luís C. Lamb: **Neural-Symbolic Learning and Reasoning** (Dagstuhl Seminar 14381), 2015.

- A. S. d'Avila Garcez, L. de Raedt, Luis. C. Lamb, R. Miikkulainen, P. Hitzler, T. Icard, T. Besold, P. Foldiak, D. Silver and K.U. Kuehnberger. *Neural-Symbolic Learning and Reasoning: Contributions and Challenges*, Proceedings of the **AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches**, Stanford Univ., CA, March 2015, pp. 18-21, AAAI Press, 2015

- A.S. d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. A neural cognitive model of argumentation with application to legal inference and decision making**. Journal of Applied Logic,** 2014.

References

- 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.

References

- Roni Khardon, Dan Roth: Learning to reason. J. ACM 44(5): 697-725 (1997).

- Leslie G. Valiant: Knowledge Infusion: In Pursuit of Robustness in Artificial Intelligence. FSTTCS 2008: 415-422

- Leslie G. Valiant: Three problems in computer science. J. ACM 50(1): 96-99 (2003)