Neurosymbolic AI – a selective history

Pascal Hitzler
Data Semantics Laboratory (DaSe Lab)
Kansas State University

http://www.daselab.org

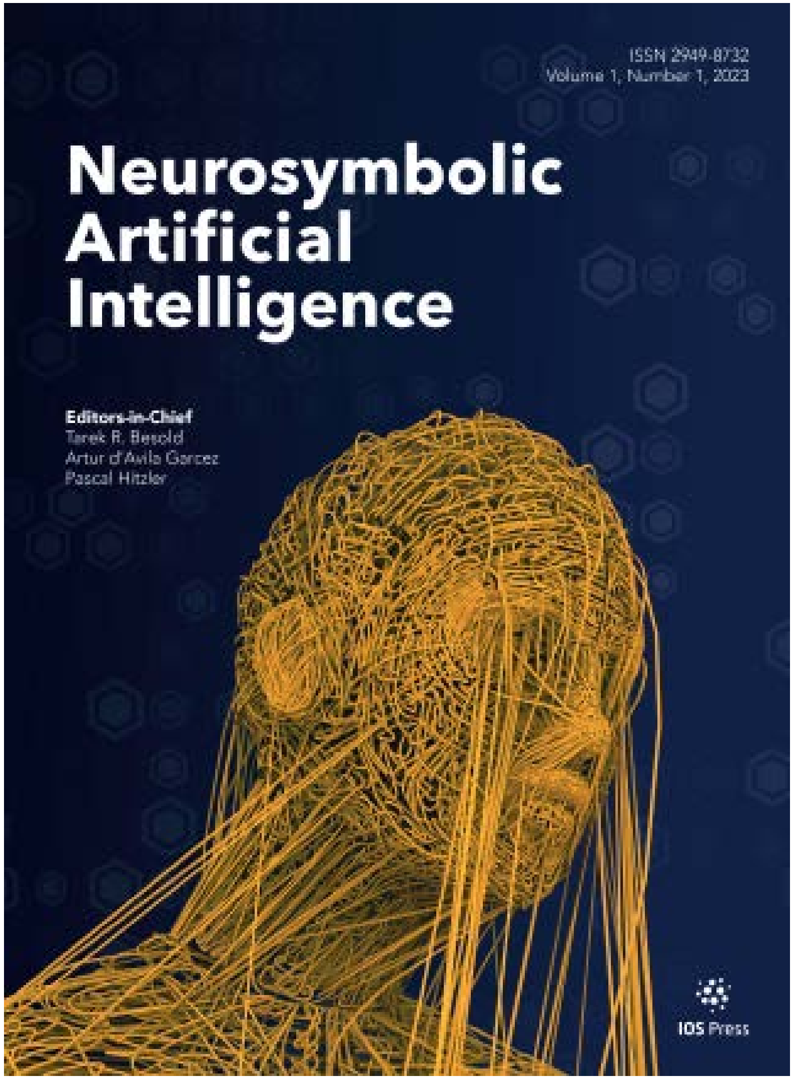

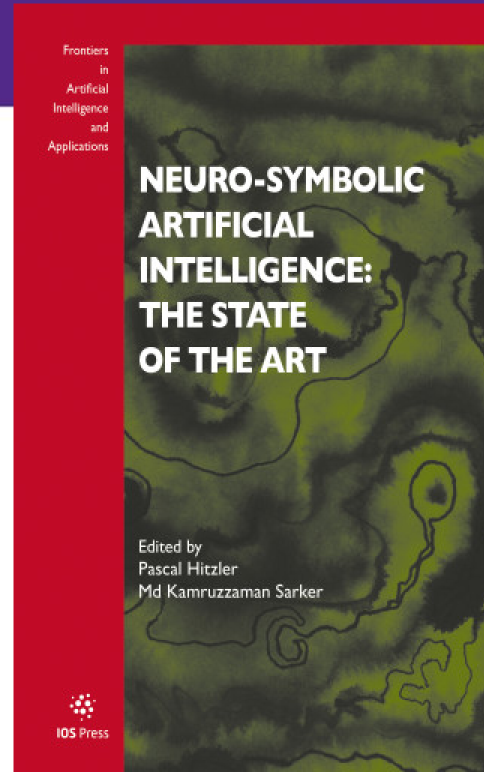

2022, 17 chapters

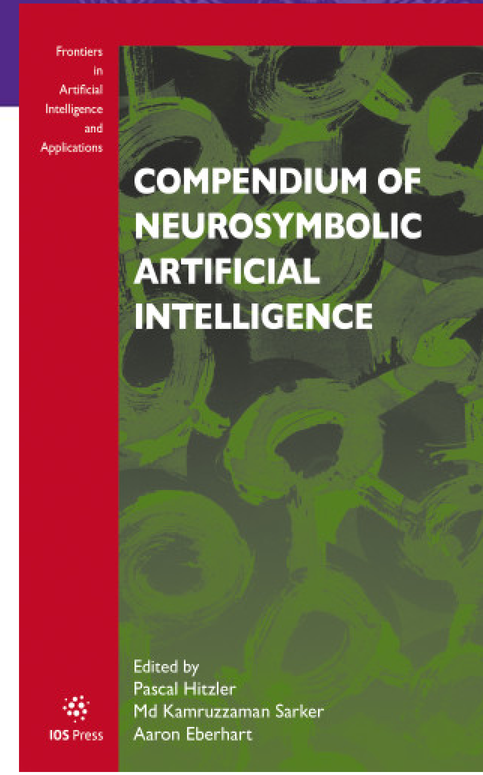

2023, 30 chapters

Neurosymbolic AI community slack
currently over 800 members
email hitzler@ksu.edu to get an invite

Neuro-symbolic AI

Publications on neuro-symbolic AI in major conferences (research papers only):

conference	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	total
ICML	0	0	0	0	0	1	3	2	5	6	17
NeurIPS	0	0	0	0	0	0	0	4	2	4	10
AAAI	0	0	0	0	0	1	0	1	1	1	4
IJCAI	1	0	0	0	0	0	2	2	0	2	7
ICLR	N/A	N/A	0	0	0	0	1	1	1	3	6
total	1	0	0	0	0	2	6	10	9	16	44

See

Md Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, Pascal Hitzler

Neuro-Symbolic Artificial Integration: Current Trends

AI Communications 34 (3), 197-209, 2022.

Neurosymbolic? Neuro-symbolic? Neural-Symbolic? Symbolic-Subsymbolic?

Neural

- **Refers to computational abstractions of (natural) neural network systems.**
- **Prominently includes Artificial Neural Networks and Deep Learning as machine learning paradigms.**
- **More generally sometimes referred to as *connectionist systems*.**

- **Prominent applications come from the machine learning world.**

- **And of course, there is the current deep learning hype.**

Symbolic

- Refers to (computational) symbol manipulations of all kind.

- Graphs and trees, traversal, data structure operations.
- Knowledge representation in explicit symbolic form (data base, ontology, knowledge graph)
- Inductive and statistical inference.
- Formal logical (deductive or abductive) reasoning.

- Prominent applications all over computer science, including expert systems (and their modern versions), information systems, data management, added value of data annotation, etc.

- Semantic Web data is inherently symbolic.

Neural-Symbolic

Computer Science perspective:

- **Connectionist machine learning systems are**
  - **very powerful for some machine learning problems**
  - **robust to data noise**
  - **very hard to understand or explain**
  - **really poor at symbol manipulation**
  - **unclear how to effectively use background (domain) knowledge**
- **Symbolic systems are**
  - **Usually rather poor regarding machine learning problems**
  - **Intolerant to data noise**
  - **Relatively easy to analyse and understand**
  - **Really good at symbol manipulation**
  - **Designed to work with other (background) knowledge**

Neural-Symbolic

Computer Science perspective:

- **Let's try to get the best of both worlds:**
  - very powerful machine learning paradigm
  - robust to data noise
  - easy to understand and assess by humans
  - good at symbol manipulation
  - work seamlessly with background (domain) knowledge

- **How to do that?**
  - Endow connectionist systems with symbolic components?
  - Add connectionist learning to symbolic reasoners?

Neural-Symbolic

Note:

- **Deep Learning systems are a far cry from how natural neural networks work.**
- **There are things that our brain can do, and things that symbolic approaches can do, where we do not have the faintest idea how to solve them through deep learning (or any other connectionist learning approach).**
- **The argument that we "just don't have enough training data" speaks (understandably) to the current hype, but is a hope that is unfounded: While this may be the case in some cases, there is no rationale to overgeneralize.**
  [Note: if we had "enough computational power," we could also solve all reasonable-size NP-complete problems in an instant.]

The Interface Issue

- Symbolic knowledge comes as logical theories (sets of formulas over a logic)
- Subsymbolic systems process tuples of real/float numbers (vectors, matrices, tensors)

- How do you interface?
- How do you map between the symbolic world and the subsymbolic world?

Some key problems that need to be overcome:

- Logic is full of highly structured objects, how to represent them in Real Space?
- How to represent variable bindings in a distributed setting?
- The required length of logical deduction chain is not known up front.

Representation Issues

McCulloch & Pitts, 1943

- McCulloch & Pitts 1943
  - Neurons with binary activation functions.
  - Modelling of propositional connectives.
  - Networks equivalent to finite automata.

Values 0 („false“) and 1 („true“) being propagated.

Simultaneous update of all nodes in network.

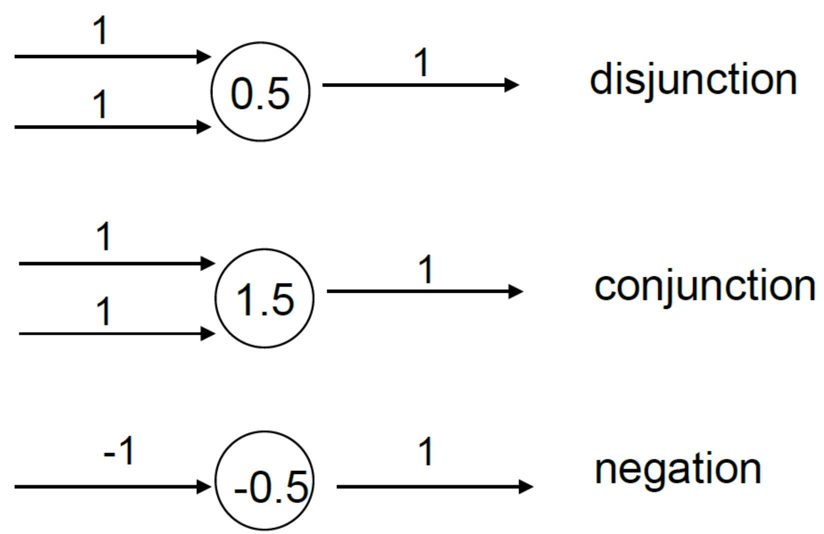

McCulloch & Pitts follow-on

- Hölldobler & Kalinke 1994
  - Extends the approach by McCulloch & Pitts.
  - Representation of propositional logic programs and their semantics.
  - „Massively parallel reasoning.“

logic program ⟶ core net ⟶ recurrent net

$a \leftarrow$
$b \leftarrow a$
$c \leftarrow a \wedge b$
$d \leftarrow e$
$e \leftarrow d$

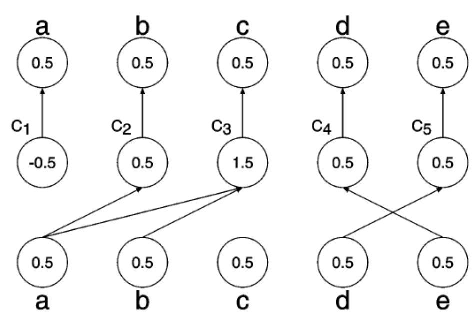

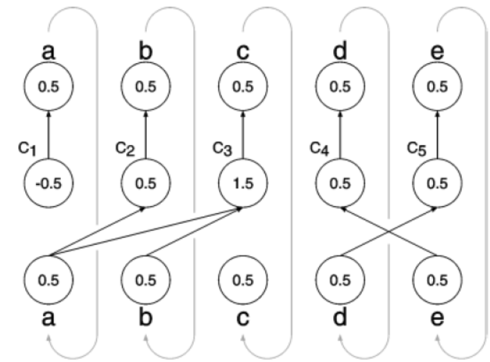

McCulloch & Pitts follow-on

Logic program P → core net

$a \leftarrow$
$b \leftarrow a$
$c \leftarrow a \wedge b$
$d \leftarrow e$
$e \leftarrow d$

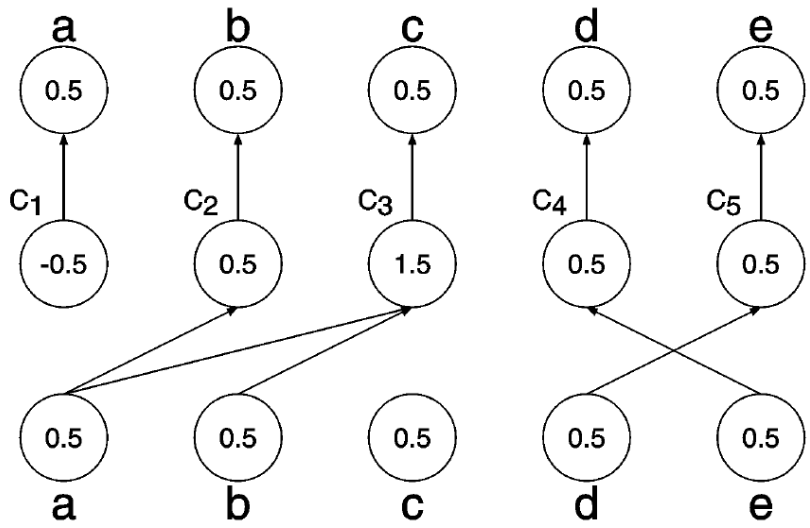

- Update „along implication“.
- Corresponds to computing the semantic operator T_P.
- T_P represents meaning (semantics) of P through its fixed points.

McCulloch & Pitts follow-on

DaSe Lab

core net

recurrent net

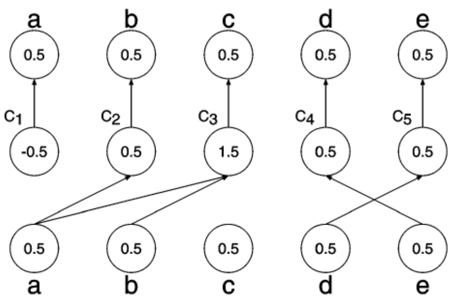

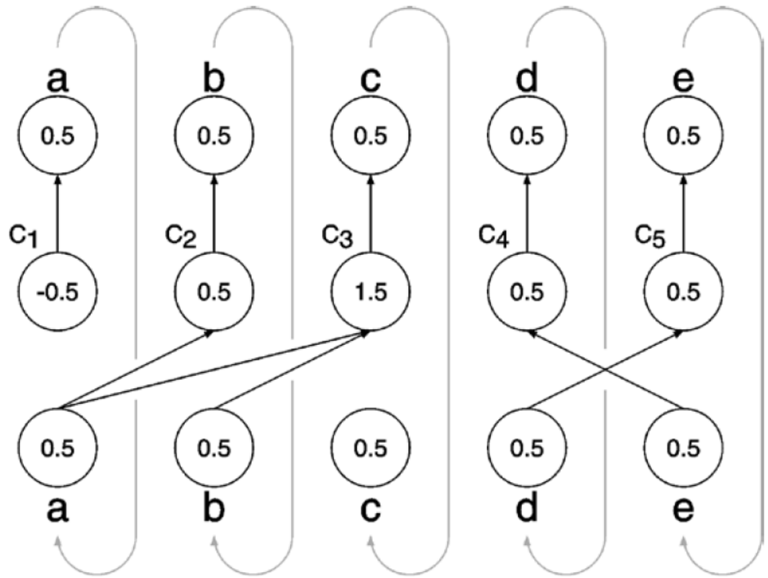

- Repeated updates along layers corresponds to iterations of the semantic operator.
- Semantics of the program (= fixed point of the operator) can be computed in a parallel manner.

McCulloch & Pitts follow-on

- Garcez & Zaverucha 1999
  Garcez, Broda & Gabbay 2001
- Development of a learning paradigm from the Core Method.
- Required: differentiable activation function.
  – Allows learning with standard methods.
  – Backpropagation algorithm.

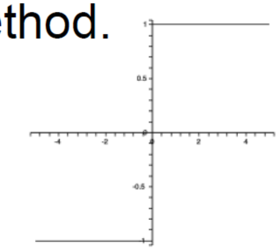

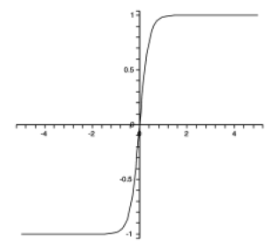

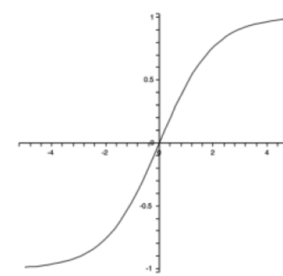

- Establishing the *neural-symbolic learning cycle*.

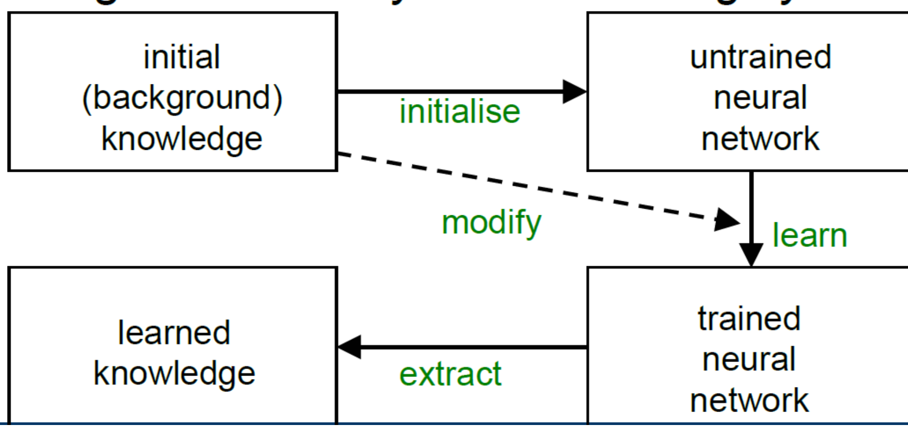

The catch

- **This is all propositional.**
- **There's only that much you can do with propositional logic.**
- **In particular, in terms of knowledge representation and reasoning, propositional logic doesn't really get you anything useful.**

E.g.

- **RDF (knowledge graphs) is already much closer to datalog than to propositional logic.**
- **OWL (knowledge graph schemas) is a fragment of first-order predicate logic.**

Variable Binding

SHRUTI

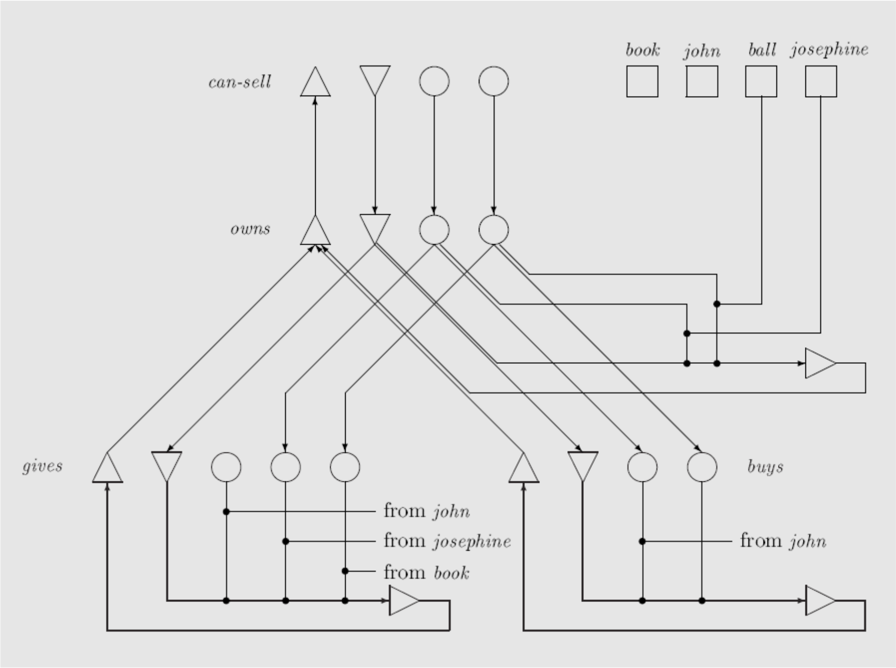

Shastri & Ajjanagadde 1993

Variable binding
via time synchronization.

Reflexive (i.e. fast)
reasoning possible.

Picture: Hölldobler,
Introduction to
Computational Logic, 2001

gives(X,Y,Z) $\longrightarrow$ owns(Y,Z)	gives(john,josephine,book)
buys(X,Y) $\longrightarrow$ owns(X,Y)	$(\exists X)$ buys(john,X)
owns(X,Y) $\longrightarrow$ can-sell(X,Y)	owns(josephine,ball)

Problems:

- **It's still essentially datalog.** *** It doesn't really learn.**
- **It has a globally bounded reasoning depth.**

Logic in Real Space

DaSe Lab

Interpretations with Cantor topology

Homeomorphism

Cantor space as compact subspace of $\mathcal{R}$

$$
\begin{tikzcd}
(I_P, Q) \arrow[r, "T_P"] \arrow[d, "i"'] & (I_P, Q) \arrow[d, "i"] \\
\text{Cantor} \arrow[r, "i(T_P)"] & \text{Cantor}
\end{tikzcd}
$$

[Hitzler, Dissertation, 2001]

Logic in Real Space

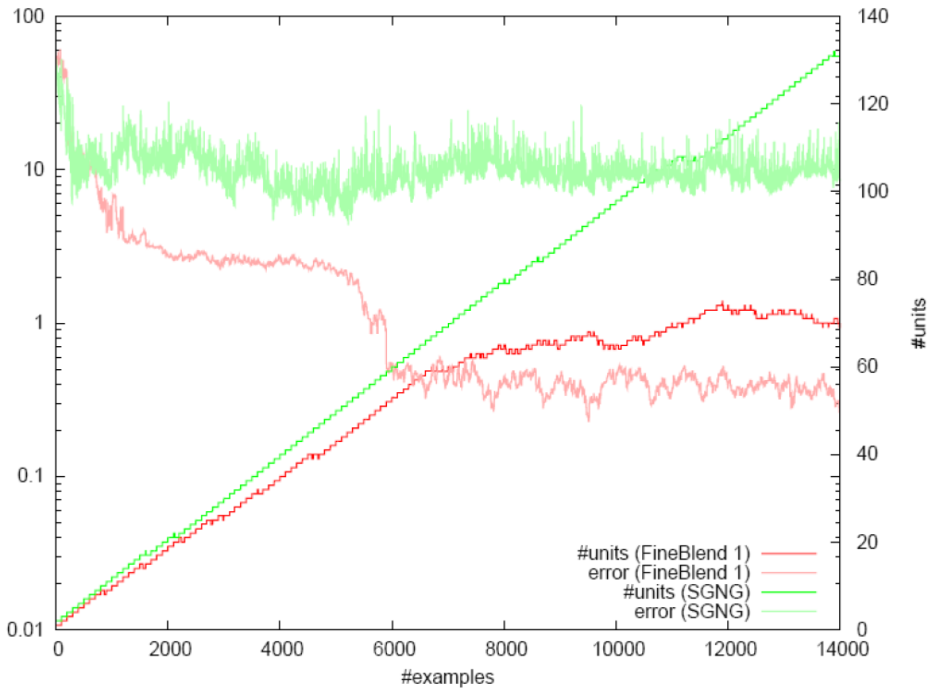

Bader, Hitzler, Hölldobler, Witzel, IJCAI-07

Architecture is mix of radial basis function network and neural gas approach.

target:	$e(0).$		
	$e(s(X))$	$\leftarrow$	$o(X).$
	$o(X)$	$\leftarrow$	$\neg e(X)$
initial:	$e(s(X))$	$\leftarrow$	$\neg o(X)$
	$e(X)$	$\leftarrow$	$e(X)$

unit failure

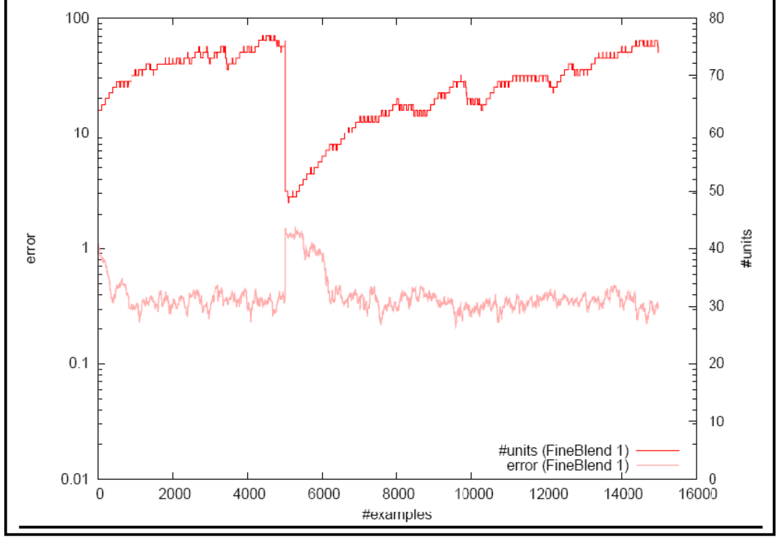

Logic in Real Space

We observe convergence to unique supported model of the program.

Bader, Hitzler, Hölldobler, Witzel, IJCAI-07

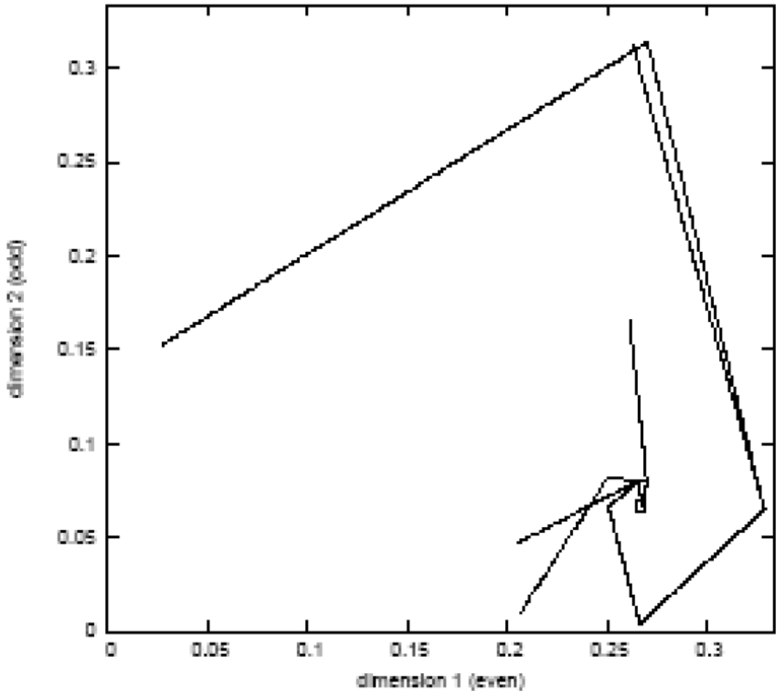

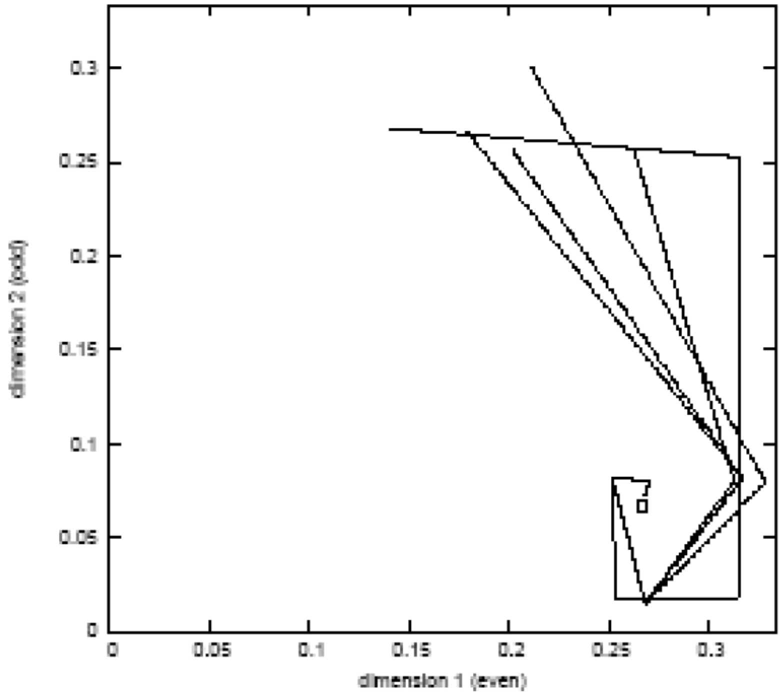

But it works only for toy size problems.
The theoretically required embedding into real numbers cannot scale.

Some other early contributions

[Bader, Hitzler, 2004]

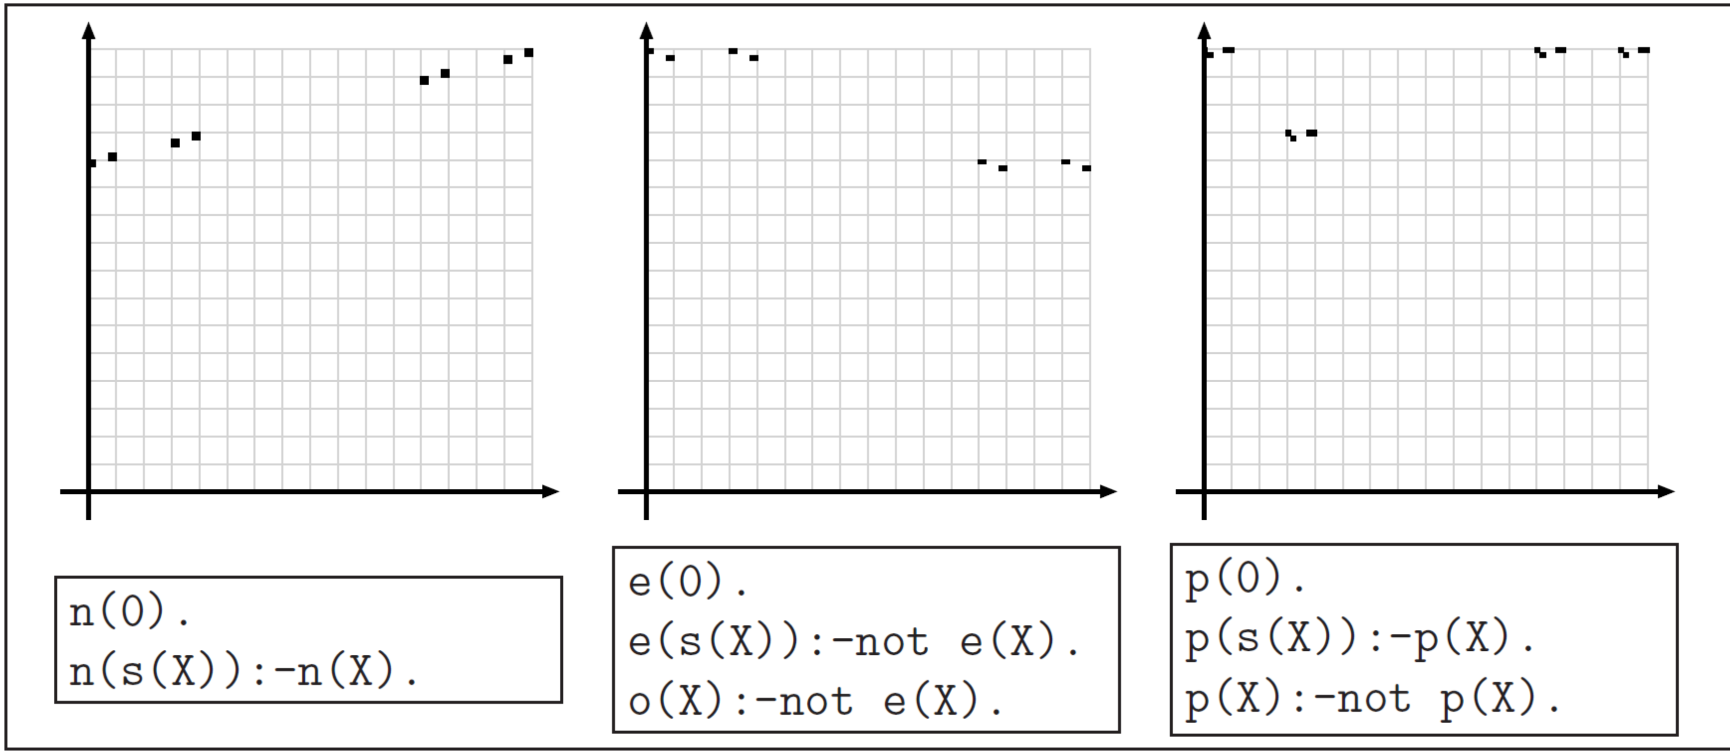

Consequence operators of logic programs mapped to the Cantor set in the real numbers are fractals (self-similar) – and can formally described by means of Iterated Function Systems.

Rule Extraction

[Lehmann, Bader, Hitzler 2010; Labaf, Hitzler, Evans 2017]

- There is always a unique reduced/minimal (definite) propositional logic program that captures the input-output behavior of a given ANN (3-layer feedforward neural network with threshold activation functions).

- Corresponding programs always exist and are usually smaller if additional background knowledge is taken into account, but uniqueness is not guaranteed then.

Current Trends

Neuro-symbolic AI

Publications on neuro-symbolic AI in major conferences (research papers only):

conference	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	total
ICML	0	0	0	0	0	1	3	2	5	6	17
NeurIPS	0	0	0	0	0	0	0	4	2	4	10
AAAI	0	0	0	0	0	1	0	1	1	1	4
IJCAI	1	0	0	0	0	0	2	2	0	2	7
ICLR	N/A	N/A	0	0	0	0	1	1	1	3	6
total	1	0	0	0	0	2	6	10	9	16	44

See

Md Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, Pascal Hitzler

Neuro-Symbolic Artificial Integration: Current Trends

AI Communications 34 (3), 197-209, 2022.

Comparison with 2005 survey [Bader, Hitzler 2005]

	dimension	(a)	(b)	N/A
Interrelation	integrated (a) vs. hybrid (b)	43	0	0
	neuronal (a) vs. connectionist (b)	0	43	0
	local (a) vs. distributed (b)	2	42	0
	standard (a) vs. nonstandard (b)	43	0	0
Language	symbolic (a) vs. logical (b)	21	24	0
	propositional (a) vs. first-order (b)	3	22	18
Usage	extraction (a) vs. representation (b)	6	37	3
	learning (a) vs. reasoning (b)	19	29	0

Kautz 2020 Categories

category	number of papers
[symbolic Neuro symbolic]	13
[Symbolic[Neuro]]	9
[Neuro $\cup$ compile(Symbolic)]	10
[Neuro $\rightarrow$ Symbolic]	13
[Neuro[Symbolic]]	0

(6) We finally come to the approach to neuro-symbolic reasoning that I believe has the greatest potential to combine the strengths of logic-based and neural-based AI, namely the **Neuro[Symbolic]** architecture (Figure 15). The basic idea is to embed a symbolic reasoning engine inside a neural engine, with the goal of enabling superneuro and combinatorial reasoning. The architecture is based on Daniel Kahneman's theory of "thinking fast and

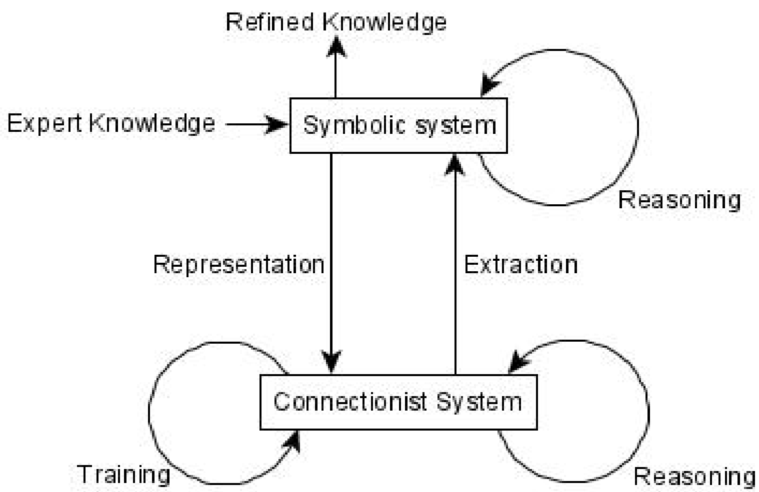

Neuro-symbolic Learning Cycle

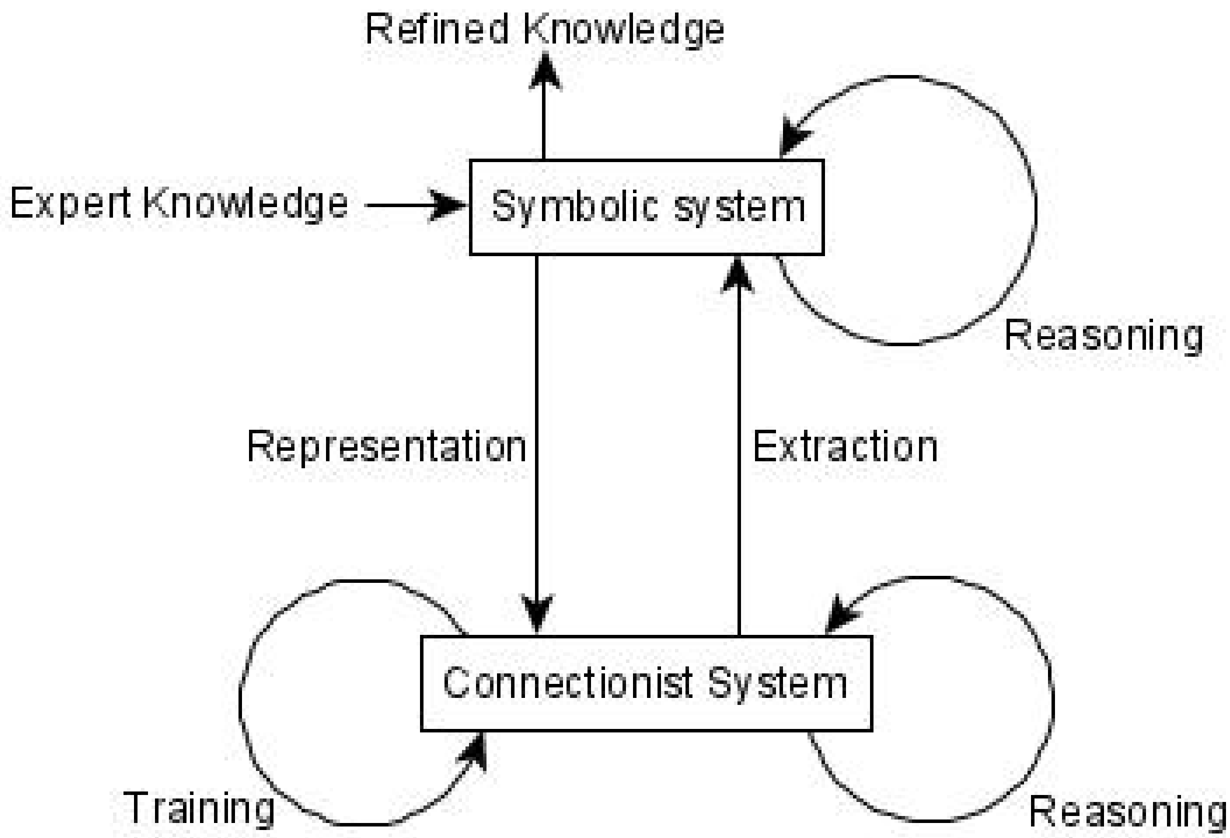

[Bader and Hitzler 2005]

Some Background

Workshop Series on Neural-Symbolic Learning and Reasoning, since 2005. Joint with Artur d’Avila Garcez.
http://neural-symbolic.org/

Barbara Hammer and Pascal Hitzler (eds), Perspectives of Neural-Symbolic Integration, Springer, 2007

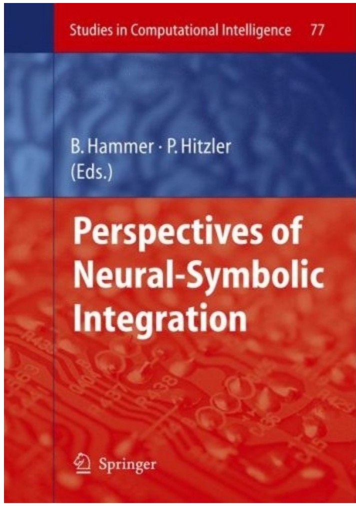

Neural-Symbolic Learning and Reasoning: A Survey and Interpretation

Tarek R. Besold, Artur d'Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pascal Hitzler, Kai-Uwe Kuehnberger, Luis C. Lamb, Daniel Lowd, Priscila Machado Vieira Lima, Leo de Penning, Gadi Pinkas, Hoifung Poon, Gerson Zaverucha

https://arxiv.org/abs/1711.03902 (2017)

Ilaria Tiddi, Freddy Lecue, Pascal Hitzler (eds.), Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges. Studies on the Semantic Web Vol. 47, IOS Press, 2020.

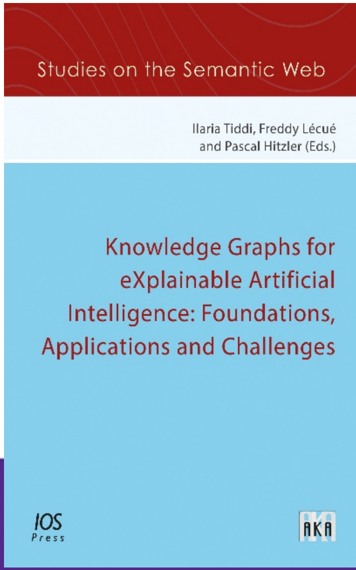

Thanks!

References

Md Kamruzzaman Saker, Lu Zhou, Aaron Eberhart, Pascal Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends. AI Communications 34 (3), 197-209, 2022.

Pascal Hitzler, Md Kamruzzaman Sarker (eds.), Neuro-Symbolic Artificial Intelligence - The State of the Art. Frontiers in Artificial Intelligence and Applications Vol. 342, IOS Press, Amsterdam, 2022.

Pascal Hitzler, Md Kamruzzaman Sarker, Aaron Eberhart (eds.), Compendium of Neurosymbolic Artificial Intelligence. Frontiers in Artificial Intelligence and Applications Vol. 369, IOS Press, Amsterdam, 2023.

Barbara Hammer and Pascal Hitzler (eds), Perspectives of Neural-Symbolic Integration, Springer, 2007

References

Tarek R. Besold, Artur d'Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pascal Hitzler, Kai-Uwe Kuehnberger, Luis C. Lamb, Priscila Machado Vieira Lima, Leo de Penning, Gadi Pinkas, Hoifung Poon, Gerson Zaverucha, Neural-Symbolic Learning and Reasoning: A Survey and Interpretation. In: Hitzler, Sarker, Eberhart (eds.), Neuro-Symbolic Artificial Intelligence: The State of the Art. Frontiers in Artificial Intelligence and Applications Vol. 342, IOS Press, Amsterdam, 2021, pp. 1-51.

Ilaria Tiddi, Freddy Lecue, Pascal Hitzler (eds.), Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges. Studies on the Semantic Web Vol. 47, IOS Press, 2020.

McCulloch, Pitts, A Logical Calculus of the Ideas Immanent in Nervous Activity. Bull of Mathematical Biophysics 5, 1943, 115-113.

References

Artur S. d'Avila Garcez, Gerson Zaverucha, The Connectionist Inductive Learning and Logic Programming System. Appl. Intell. 11(1): 59-77 (1999)

Artur S. d'Avila Garcez, Krysia Broda, Dov M. Gabbay, Symbolic knowledge extraction from trained neural networks: A sound approach. Artif. Intell. 125(1-2): 155-207 (2001)

Steffen Hölldobler, Yvonne Kalinke, Ein massiv paralleles Modell für die Logikprogrammierung. WLP 1994: 89-92

Pascal Hitzler, Steffen Hölldobler, Anthony Karel Seda, Logic programs and connectionist networks. J. Appl. Log. 2(3): 245-272 (2004)

References

Sebastian Bader, Pascal Hitzler, Logic programs, iterated function systems, and recurrent radial basis function networks. J. Appl. Log. 2(3): 273-300 (2004)

Sebastian Bader, Pascal Hitzler, Dimensions of Neural-symbolic Integration - A Structured Survey. We Will Show Them! (1) 2005: 167-194

Sebastian Bader, Pascal Hitzler, Steffen Hölldobler, Andreas Witzel, A Fully Connectionist Model Generator for Covered First-Order Logic Programs. IJCAI 2007: 666-671

Jens Lehmann, Sebastian Bader, Pascal Hitzler, Extracting reduced logic programs from artificial neural networks. Appl. Intell. 32(3): 249-266 (2010)

Maryam Labaf, Pascal Hitzler, Anthony B. Evans, Propositional Rule Extraction from Neural Networks under Background Knowledge. NeSy 2017