Dagstuhl seminar 17192
8 May 2017

Neural-Symbolic Systems for Human-like Computing

Artur d'Avila Garcez
City, University of London
a.garcez@city.ac.uk

Neural-Symbolic Systems

Cognitive Science

↓

Logic

Learning

Neural Computation

↑

Neuroscience

One Structure for Learning and Reasoning

(NSS = KR + ML)

Why Neurons and Symbols?

"We need a language for describing the alternative algorithms that a network of neurons may be implementing" L. Valiant

(New) Logic + Neural Computation

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

Neural-Symbolic Methodology

high-level symbolic representations
(abstraction, recursion, relations, modalities)

translations

low level, efficient neural structures
(with the same, simple architecture throughout)

Analogy: low-level implementation (machine code) of high-level representations (e.g. java, requirement specs)

A Foundational Approach

(as opposed to the neuroscience or the engineering approach)

One Structure for Learning and Reasoning:

Take different tasks, consider what they have in common, formalize, evaluate and repeat

KEY: controlling the inevitable accumulation of errors (robustness)

Applications: training in simulators, robotics, evolution of software models, bioinformatics, power systems fault diagnosis, semantic web (ontology learning), general game playing, visual intelligence, finance, compliance.

Neural-Symbolic Learning Cycle

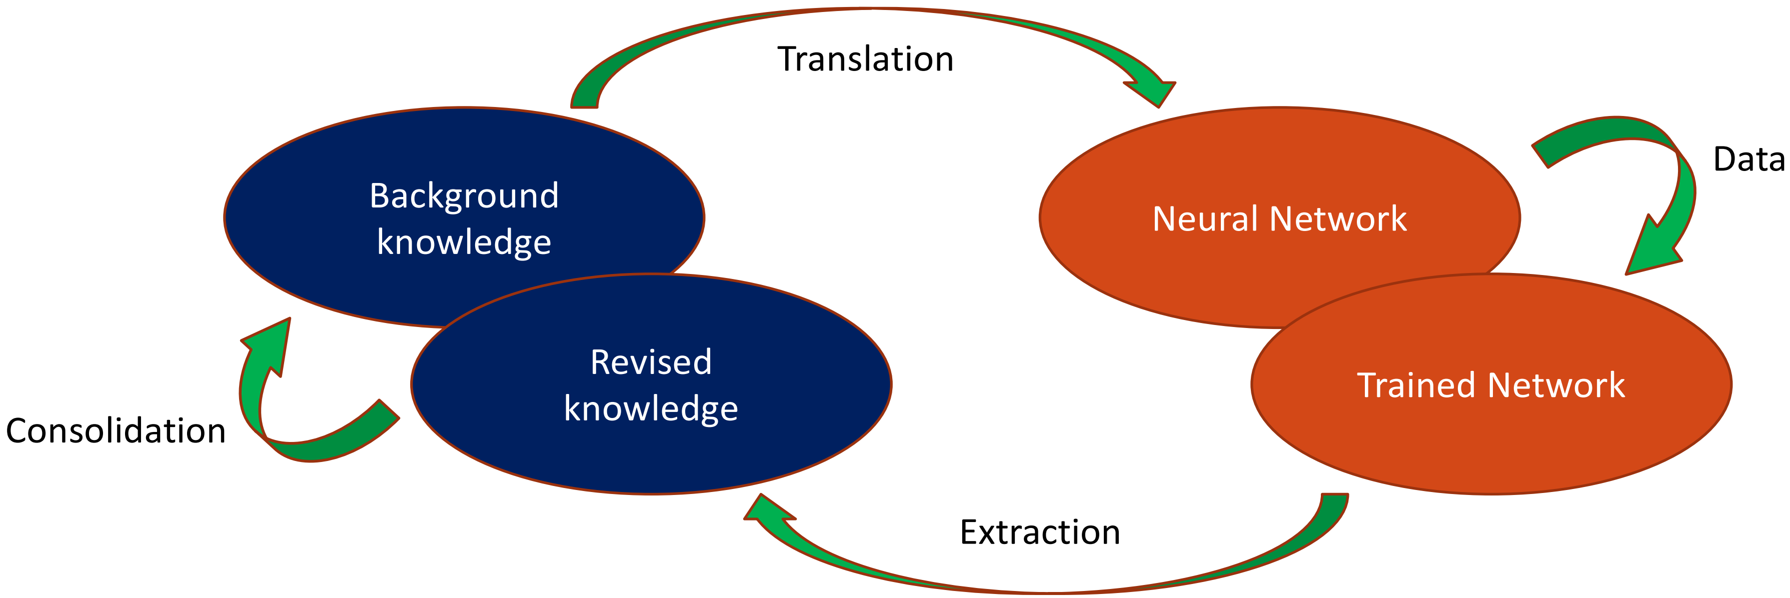

Connectionist Inductive Logic Programming (CILP System)

A Neural-Symbolic System for Integrated Reasoning and Learning (neural nets + logic programming)

Background Knowledge Insertion + Learning with Backpropagation + Knowledge Extraction

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

r_2 : $A \leftarrow E, F;$

r_3 : $B \leftarrow$

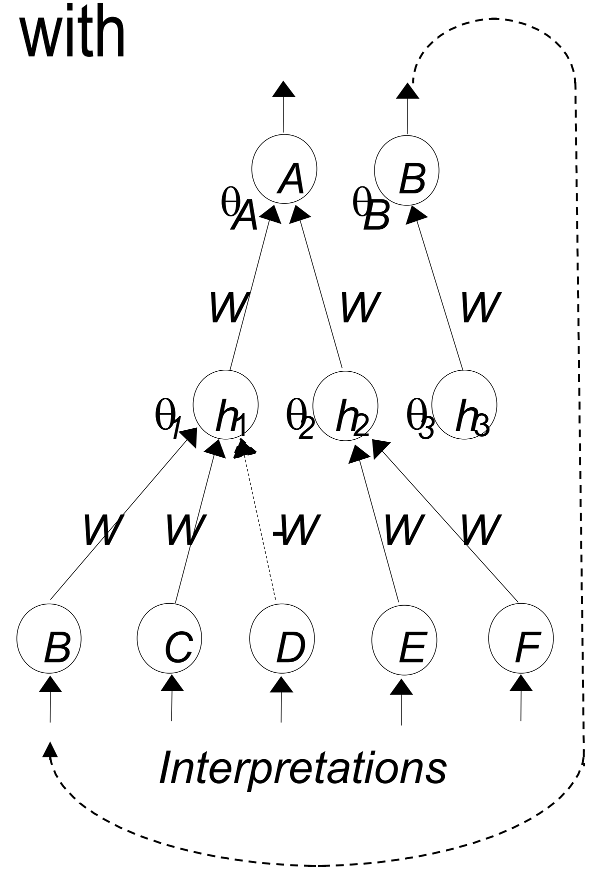

Power Plant Fault Diagnosis

Background Knowledge (35 rules with errors)

278 examples of single and multiple faults

Fault(ground,close-up,line01,no-bypass) IF
Alarm(instantaneous,line01) AND
Alarm(ground,line01)

There is a fault at transmission line 01, close to the power plant generator, due to an over-current in the ground line of transmission line 01, which occurred when the system was not using the bypass circuit.

Rule Extraction: Neural Net = Black Box?

- Extracted rules can be visualized in the form of a state transition diagram (to follow)

- Alternatively, use (un-sound but efficient) TREPAN-like rule extrac-tion and variations...

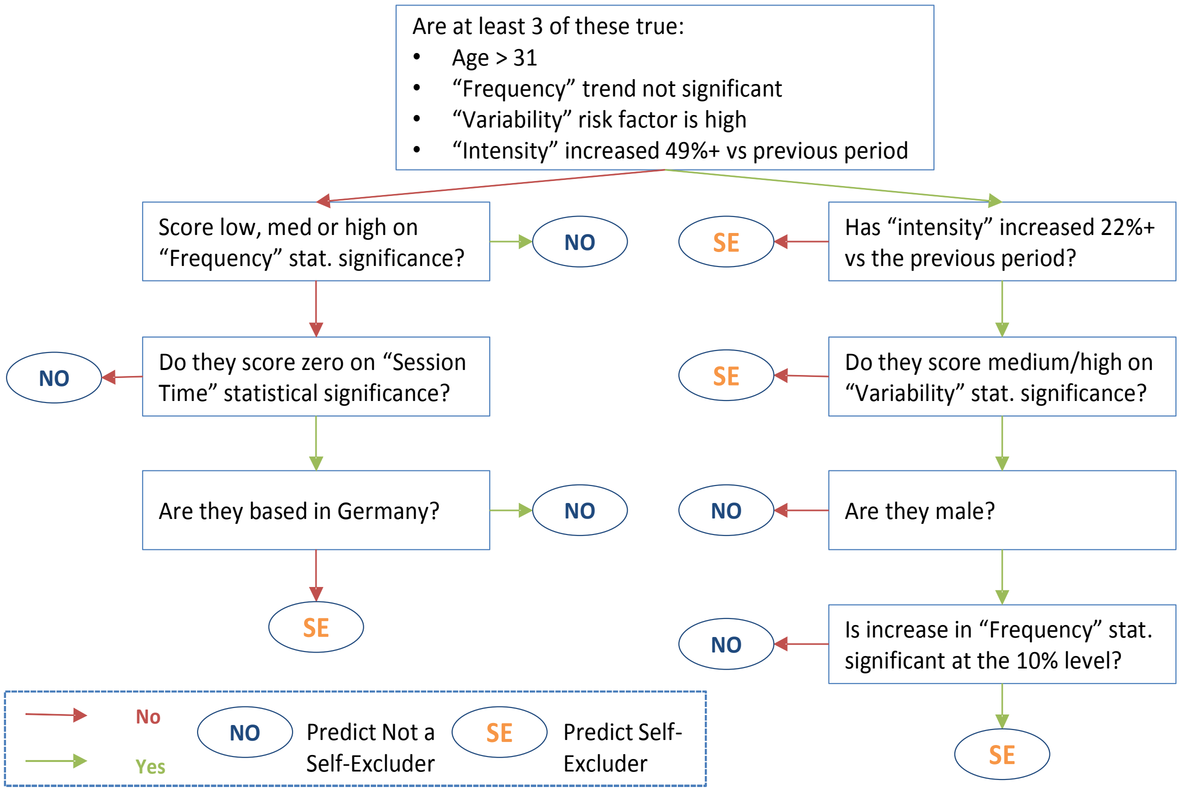

C. Percy, A. S. d'Avila Garcez, S. Dragicevic, M. Franca, G. Slabaugh and T. Weyde. The Need for Knowledge Extraction: Understanding Harmful Gambling Behavior with Neural Networks, In Proc. ECAI 2016, The Hague, September 2016.

Knowledge Consolidation

Challenge: efficient extraction of sound, comprehensible knowledge from large-scale networks (100's of neurons; 1000's of connections)

What makes knowledge comprehensible?

Transfer Learning

S. Tran and A. S. d'Avila Garcez. Deep Logic Networks: Inserting and Extracting Knowledge from Deep Belief Networks. IEEE TNNLS, Nov, 2016

CILP extensions (richer knowledge)

- The importance of non-classical reasoning: preferences, nonmonotonic, modal, temporal, epistemic, intuitionistic logic, abductive reasoning, value-based argumentation (dialogues).

- New applications: normative reasoning, temporal logic learning with model checking, software model adaptation (business process evolution), training and assessment in simulators (driving test), visual intelligence (action classification in video)...

CILP network ensembles (deep networks!)

Modularity for learning; accessibility relations for modal (temporal) reasoning; modelling uncertainty with disjunctive rules.

Connectionist Modal Logic = good trade-off between expressiveness and computational complexity

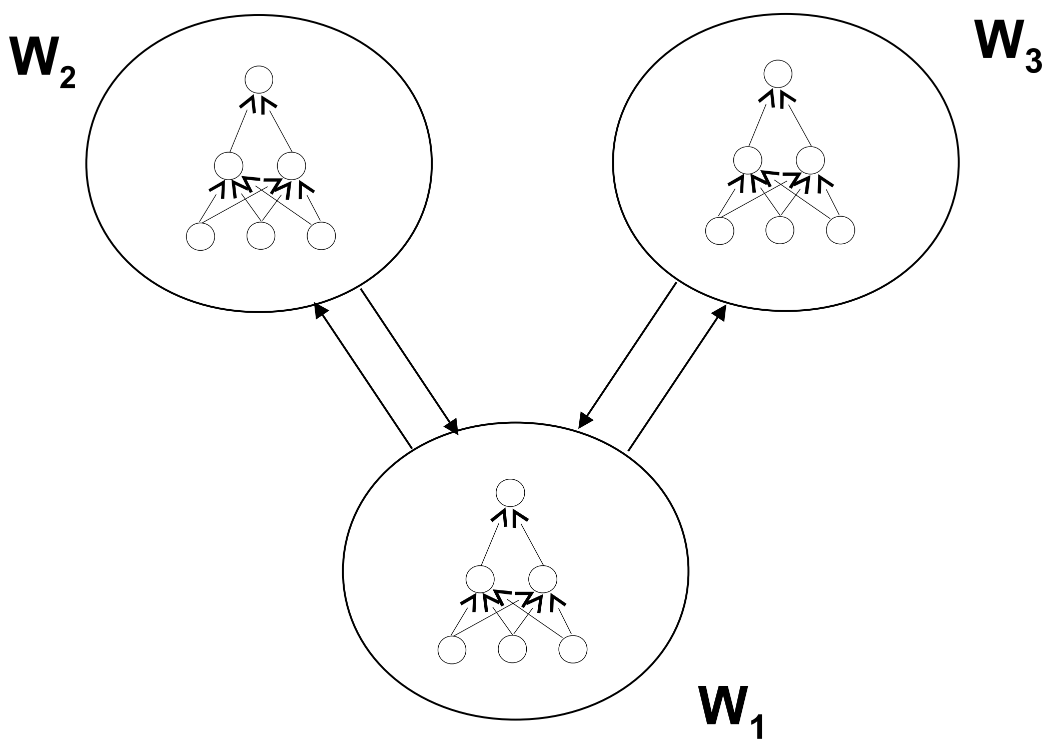

Connectionist Temporal Reasoning and Learning

The muddy children puzzle (children are playing in a garden; at least one of them is muddy; they can see if the others are muddy, but not themselves; a caretaker asks: do you know if you're muddy?). A full solution to the puzzle can only be given by a two-dimensional network ensemble.

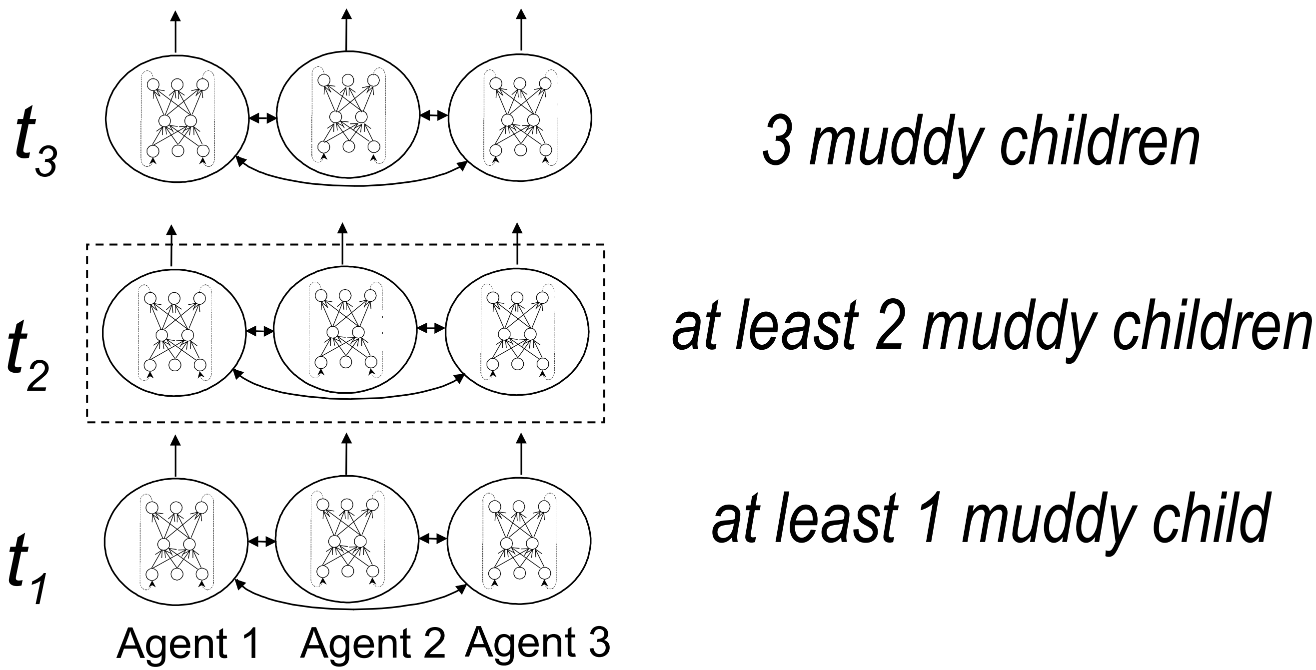

Learning with modal background knowledge is faster and offers better accuracy than learning by examples only (93% vs. 84% average test set accuracy)

Three wise men, kings and hats, etc.

- Various such logic puzzles and riddles can be useful at helping us understand the capabilities and limitations of neural models

For details: Garcez, Lamb and Gabbay, Neural-Symbolic Cognitive Reasoning, Springer, 2009.

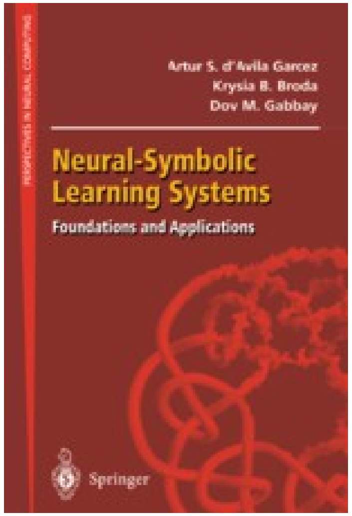

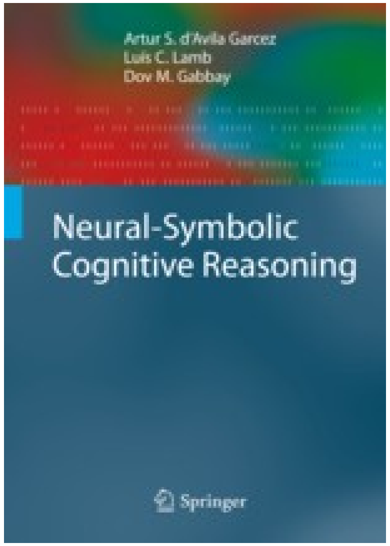

Applications

Software Model Verification and Adaptation

Verification: NuSMV
Adaptation: Neural-Symbolic System

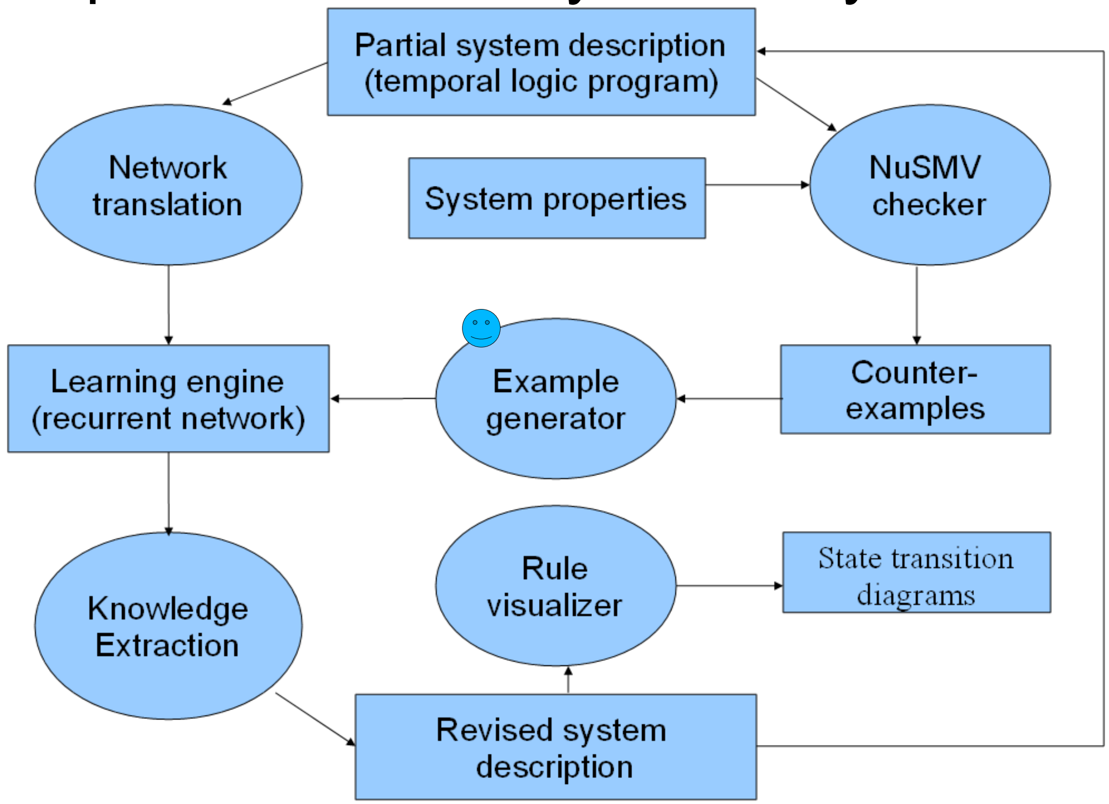

Borges, Garcez, Lamb. Learning and Representing Temporal Knowledge in Recurrent Networks. IEEE TNN 22(12):2409 - 2421, Dec 2011.

See also: F. Vaandrager, Model learning, CACM, Feb 2017.

V&A applied to Pump System

The pump system controls the levels of water in a mine to avoid the risk of overflow; an initial, partial system description is available.

State variables: *CrMeth* (level of methane is critical)
HiWat (level of water is high)
PumpOn (pump is turned on)

Safety property in LTL: *G¬(CrMeth ^ HiWat ^ PumpOn)*

Partial system spec (background knowledge; s = sensor):

$$\bigcirc CrMeth \leftarrow sCMOn.$$
$$\bigcirc CrMeth \leftarrow CrMeth, \sim sCMOff.$$
$$\bigcirc HiWat \leftarrow sHiW.$$
$$\bigcirc HiWat \leftarrow CrMeth, \sim sLoW.$$
$$\bigcirc PumpOn \leftarrow TurnPOn.$$
$$\bigcirc PumpOn \leftarrow CrMeth, \sim TurnPOff.$$

Verification (NuSMV) and example generation

New Counter-example		
t	State	Input
1	$\{\sim CrMeth, \sim HiWat \sim PumpOn\}$	$sCMOn$
2	$\{CrMeth, \sim HiWat, \sim PumpOn\}$	$TurnPOn$
3	$\{CrMeth, \sim HiWat, PumpOn\}$	$sHiW$
4	$\{CrMeth, HiWat, PumpOn\}$	–

A training example:
sCMOn → *TurnPOn* → *sHiW* → ¬*PumpOn*

Corresponding to a new rule: If methane is critical then turn the pump on, unless the water level is high, in which case turn the pump off...

Repeat the process until the property is (hopefully) satisfied (i.e. no counter-example is generated)

Neural network is three-valued {-1,0,1} CILP network, similar to NARX, trained with standard backprop.

Rule Extraction and Visualization

CrMeth = M (level of methane is critical)
HiWat = W (level of water is high)
PumpOn = P (pump is turned on)

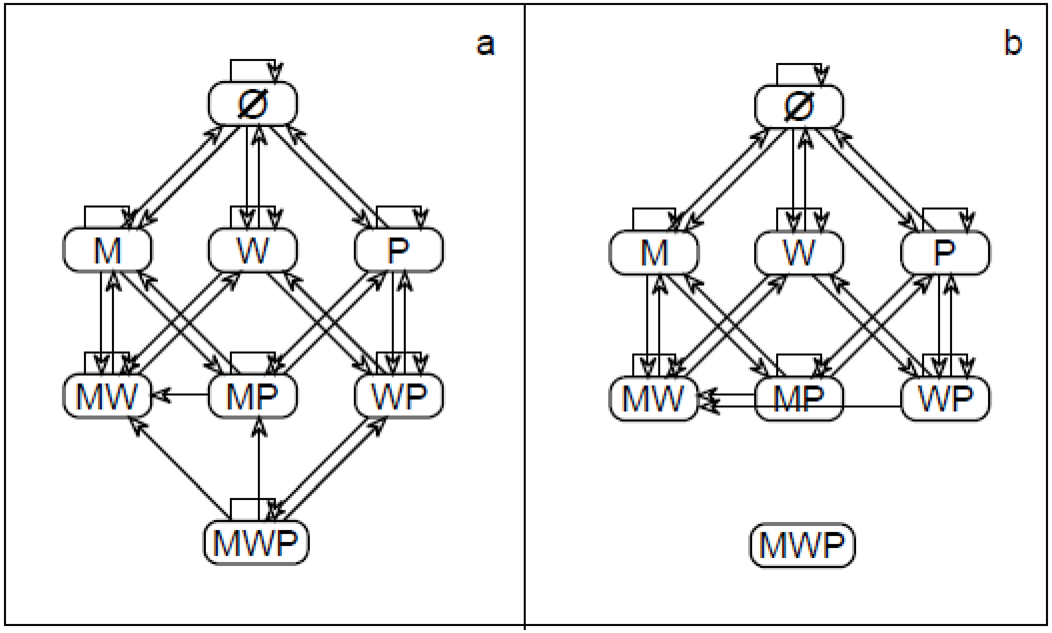

Power Plant Fault Diagnosis (real problem; ongoing validation)

Safety property: $G\neg(Fault(\_,\_,line1,bypass) \wedge Fault(\_,\_,line2,bypass))$

(diagrams are annotated with alarms which trigger derived faults)

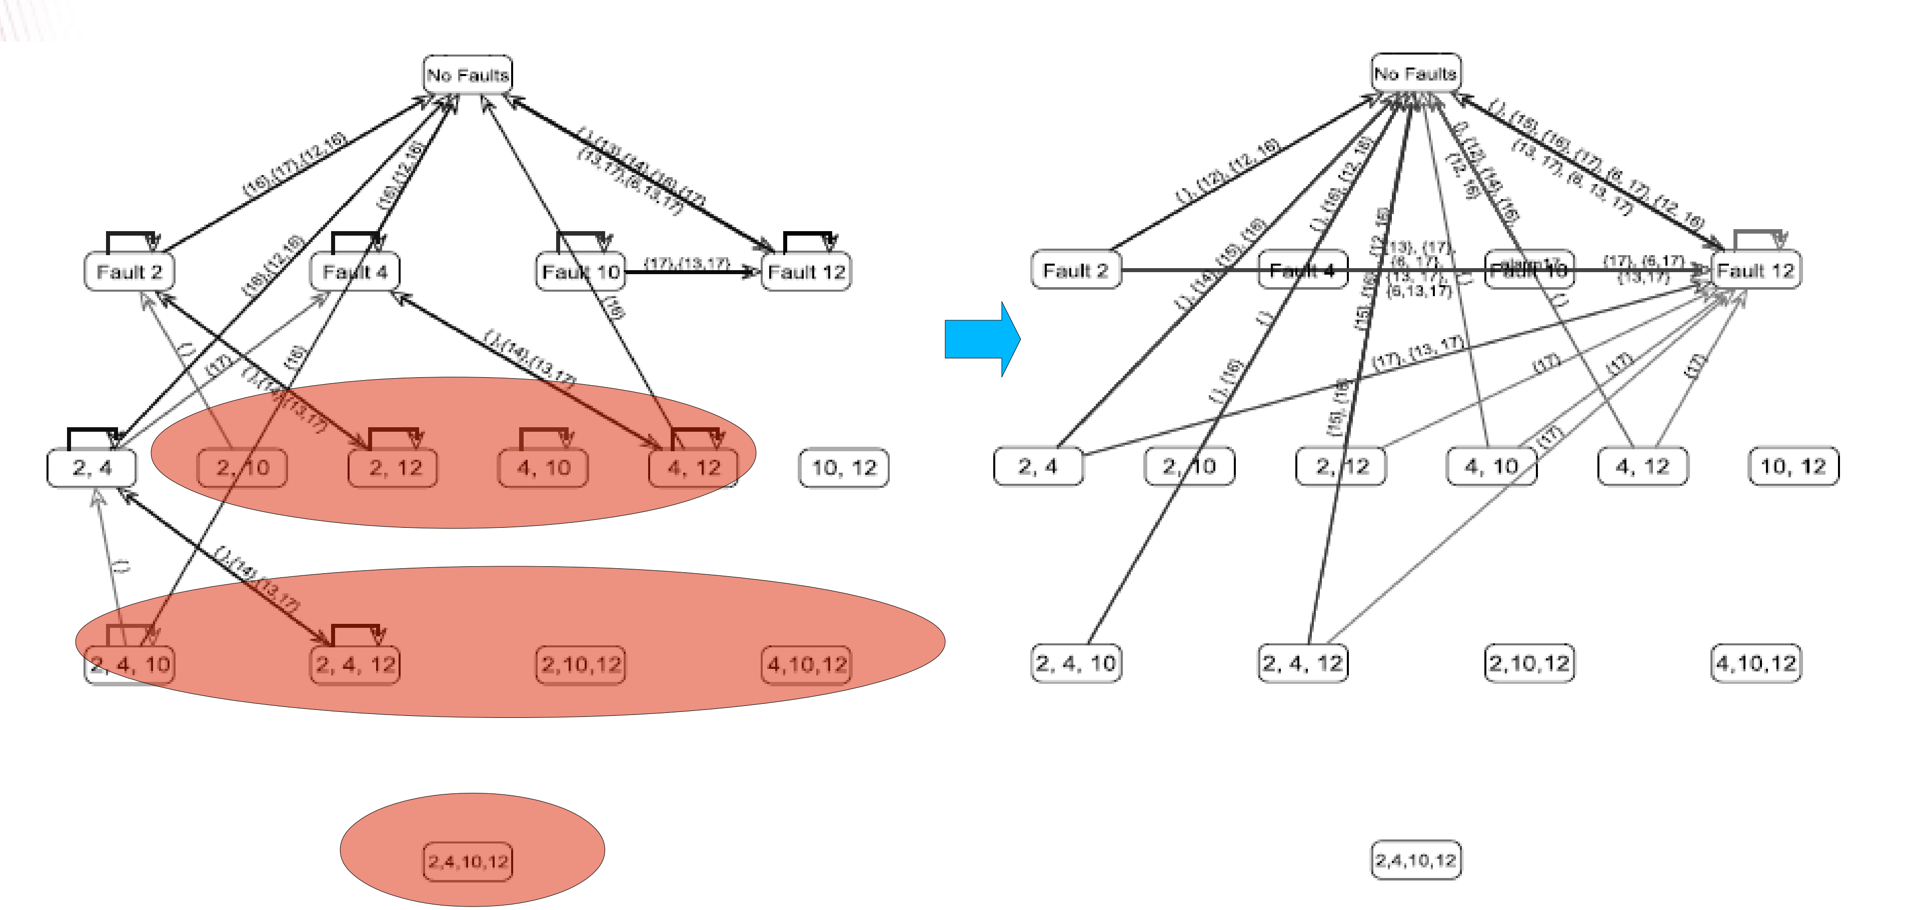

Run-time Monitoring

- So far, LTL property is outside the neural net
- Let's consider property adaptation next.

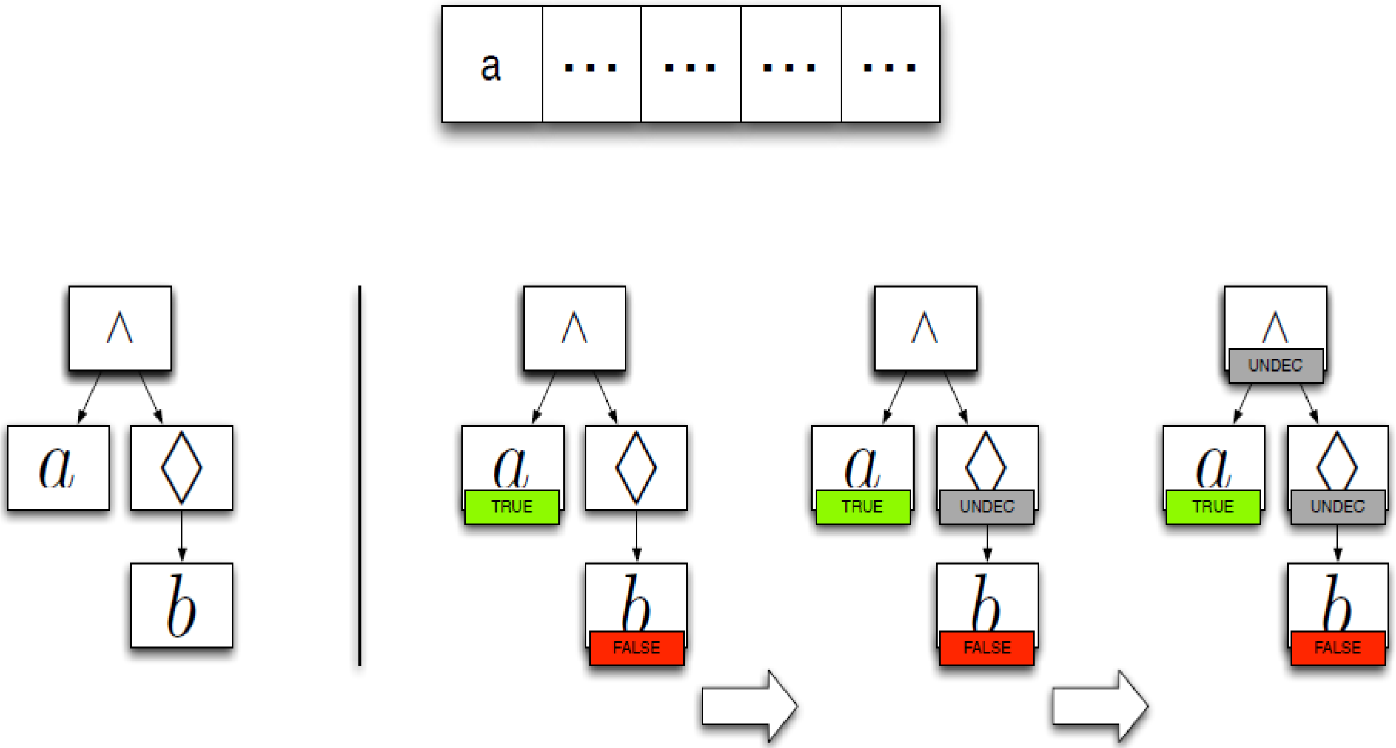

Neural Encoding

- Every tree node implements a truth-table for one operator
- Every truth-table can be represented in a CILP neural net

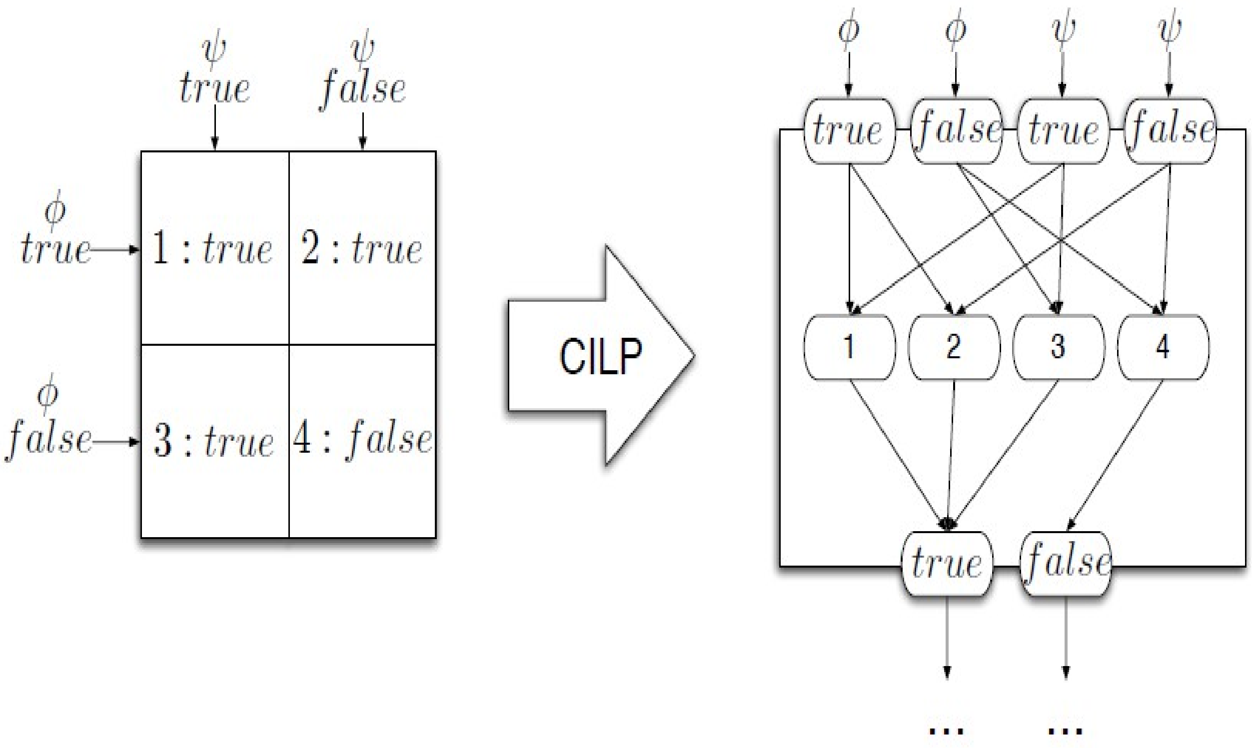

Run-Time Neural Monitor

- The tree structure is mapped onto an ensemble of CILP networks

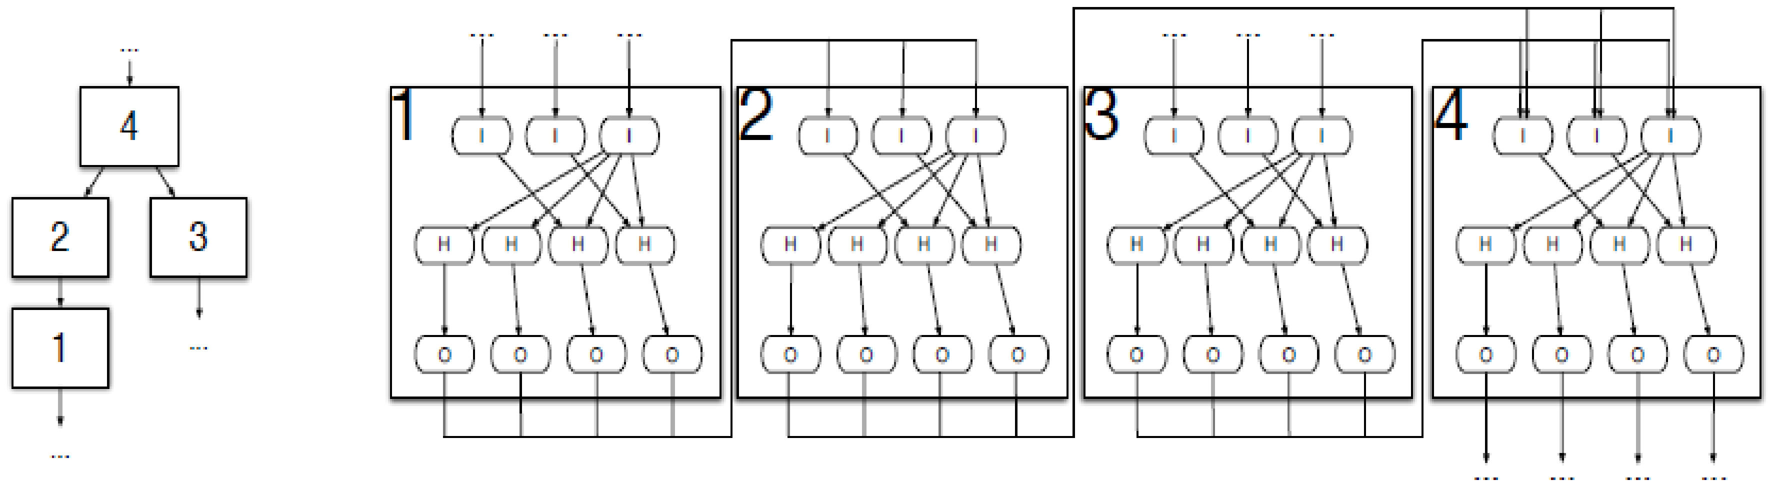

Learning = property adaptation

Local Training

Propagate from observations to verdict and backpropagate label to abduce local input-output patterns (e.g. for network 2).

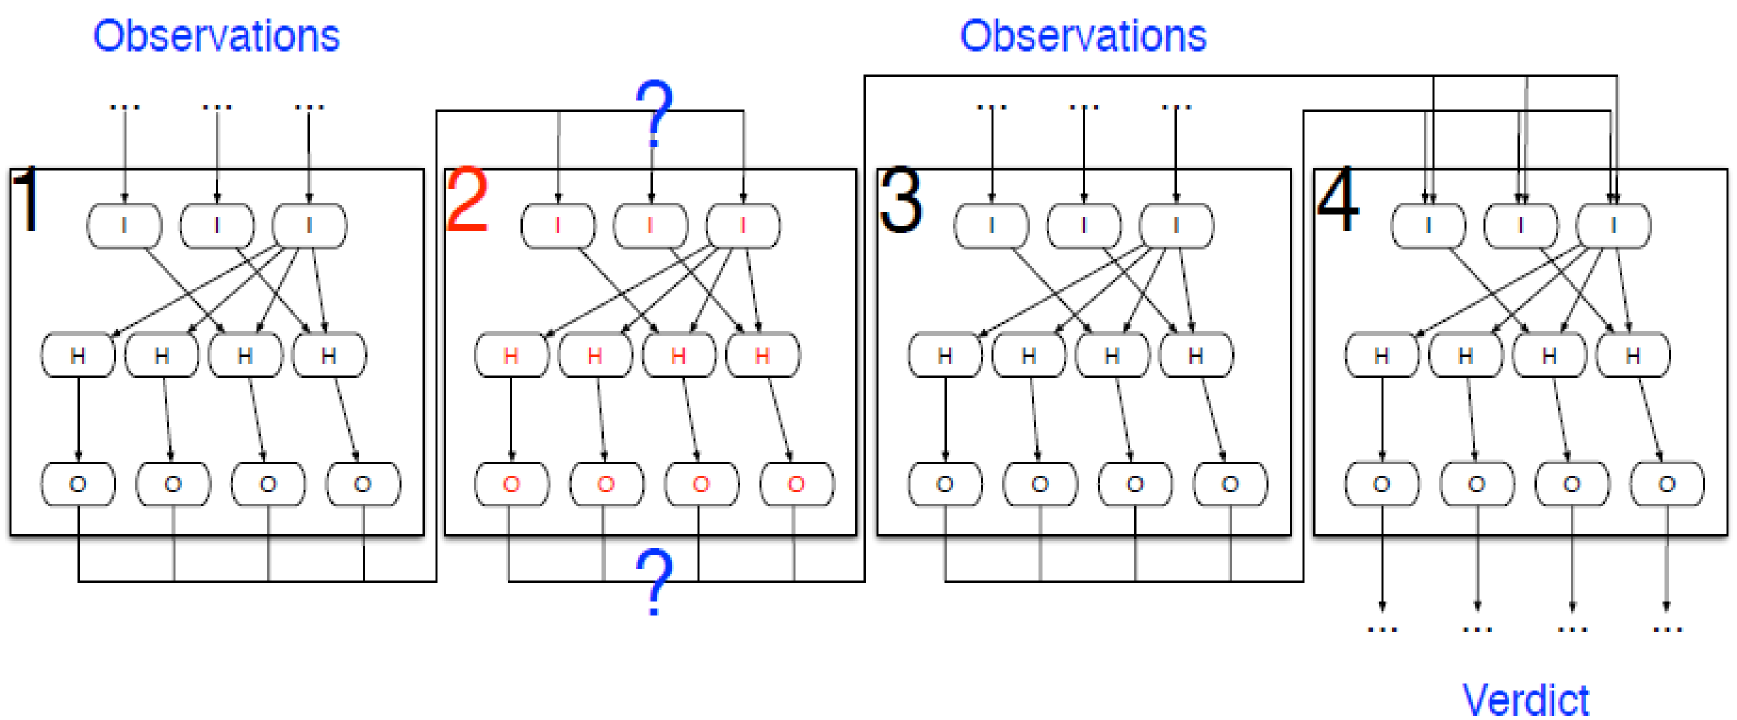

Adaptation: bending the rules

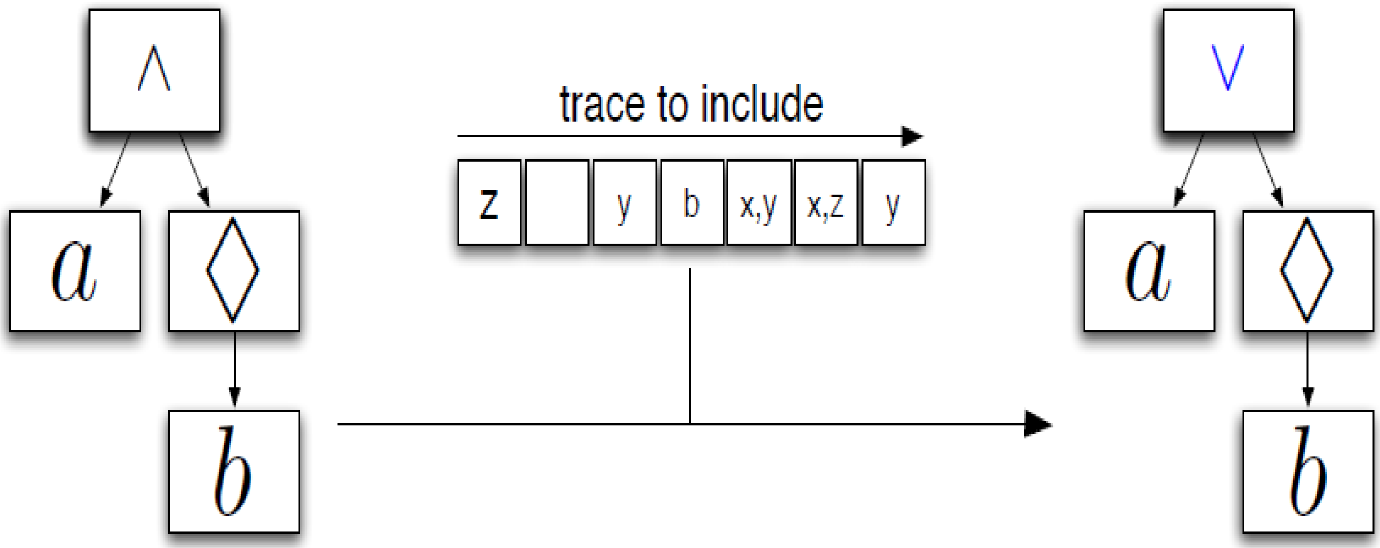

A. Perotti, G. Boella and A. S. d'Avila Garcez, Runtime Verification Through Forward Chaining. In Proc. RV'15, September 2015.

A. Perotti, A. S. d'Avila Garcez and Guido Boella. Neural-Symbolic Monitoring and Adaptation. In Proc. IJCNN 2015, July 2015.

Current/Future Work

- Verification of trained networks used e.g. for controller synthesis
  - c.f. Katz et al. Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks (Feb 2017) (Extension of Simplex to ReLUs)
- Knowledge Extraction from Deep Nets
- Relational (full FOL) Learning in Tensor Networks (with Tensorflow implementation):
  - I. Donadello, L. Serafini and A. S. d'Avila Garcez. Logic Tensor Networks for Semantic Image Interpretation, To appear IJCAI'17.
- Applications of knowledge extraction in industry: understanding pathways to harm in gambling (work with BetBuddy Ltd.)

EPSRC Human-like Computing Workshop, Cumberland Lodge, Windsor, October 2016

"Mind the Gap" HLC Desiderata

Representation Change: **neural and symbolic**, levels of abstraction

High to low-level learning: **bridging the gap**, coordinating multiple learning mechanisms

Memory and forgetting in people and AI systems: **RNNs**

Comprehensibility, language, explanation, accountability: Rule extraction

Small data learning: Using (defeasible) BK, Transfer learning

Verbal versus non-verbal communication

Social Robotics, Theories of mind, Sense of Self, Context cues, Spatial reasoning

Automated programming: psychology and application

Conclusion: Why Neurons and Symbols

To study the statistical nature of learning and the logical nature of reasoning.

To provide a unifying foundation for robust learning and efficient reasoning.

To develop effective computational systems for integrated reasoning and learning.

Thank you!